An Algorithmic Framework for Locally Constrained Homomorphisms*

Laurent Bulteau[1], Konrad K. Dabrowski[2], Noleen Köhler[3], Sebastian Ordyniak[4], and Daniël Paulusma[5]

[1] LIGM, CNRS, Université Gustave Eiffel, France
laurent.bulteau@univ-eiffel.fr[0000−0003−1645−9345]
[2] Newcastle University, UK konrad.dabrowski@newcastle.ac.uk[0000−0001−9515−6945]
[3] LAMSADE, CNRS, Université Paris-Dauphine, PSL University, France
noleen.kohler@dauphine.psl.eu[0000−0002−1023−6530]
[4] University of Leeds, UK sordyniak@gmail.com[0000−0002−1825−0097]
[5] University of Durham, UK daniel.paulusma@durham.ac.uk[0000−0001−5945−9287]

Abstract. A homomorphism ϕ from a guest graph G to a host graph H is locally bijective, injective or surjective if for every $u \in V(G)$, the restriction of ϕ to the neighbourhood of u is bijective, injective or surjective, respectively. We prove a number of new FPT, W[1]-hard and paraNP-complete results for the corresponding decision problems LBHom, LIHom and LSHom by considering a hierarchy of parameters of the guest graph G. In this way we strengthen several existing results. For our FPT results, we develop a new algorithmic framework that involves a general ILP model. We also use our framework to prove FPT results for the Role Assignment problem, which originates from social network theory and is closely related to locally surjective homomorphisms.

Keywords: (locally constrained) graph homomorphism · parameterized complexity · fracture number

1 Introduction

A *homomorphism* from a graph G to a graph H is a mapping $\phi : V(G) \to V(H)$ such that $\phi(u)\phi(v) \in E(H)$ for every $uv \in E(G)$. Graph homomorphisms generalise graph colourings (using a complete graph for H) and have been intensively studied over a long period of time, both from a structural and an algorithmic perspective. We refer to the textbook of Hell and Nešetřil [51] for a further introduction.

We write $G \to H$ if there exists a homomorphism from G to H; here, G is called the *guest graph* and H is the *host graph*. We denote the corresponding decision problem by Hom, and if H is fixed, that is, not part of the input, we write H-Hom. For graphs H without self-loops, the renowned Hell-Nešetřil dichotomy [49] states that H-Hom is

* An extended abstract of this paper appeared in the proceedings of WG2022 [14]. The second and fourth authors acknowledge support from the Engineering and Physical Sciences Research Council (EPSRC, project EP/V00252X/1).

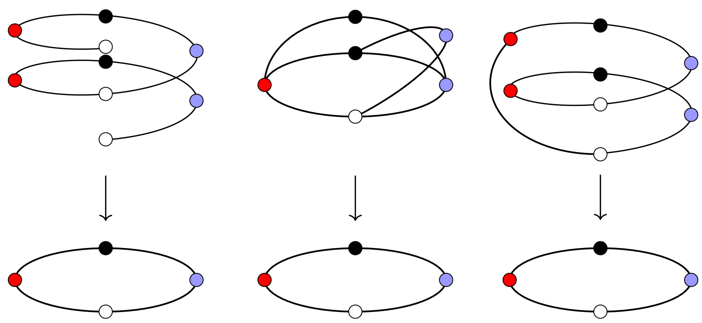

Fig. 1. Left: an example of a locally injective homomorphism which is not locally surjective. Middle: an example of a locally surjective homomorphism which is not locally injective. Right: an example of a locally bijective homomorphism.

polynomial-time solvable if H is bipartite, and NP-complete otherwise. We denote the vertices of H by $1, \dots, |V(H)|$ and call them *colours*. The reason for doing this is that graph homomorphisms generalise graph colourings: there exists a homomorphism from a graph G to the complete graph on k vertices if and only if G is k-colourable.

Instead of fixing the host graph H, one can also restrict the structure of the guest graph G by bounding some graph parameter. A classical result states that Hom is polynomial-time solvable when the guest graph G has bounded treewidth [20,42]. The *core* of a graph G is the subgraph F of G such that $G \to F$ and there is no proper subgraph F' of F with $G \to F'$ (the core is unique up to isomorphism [50]). Dalmau, Kolaitis and Vardi [23] proved that the Hom problem is polynomial-time solvable even if the core of the guest graph G has bounded treewidth. This result was strengthened by Grohe [47], who proved that if FPT $\neq$ W[1], then Hom can be solved in polynomial time if and only if this condition holds.

1.1 Locally Constrained Homomorphisms

We are interested in three well-studied variants of graph homomorphisms that occur after placing constraints on the neighbourhoods of the vertices of the guest graph G. Consider a homomorphism ϕ from a graph G to a graph H. We say that ϕ is *locally injective, locally bijective* or *locally surjective for a vertex* $u \in V(G)$ if the restriction $\phi_u : N_G(u) \to N_H(\phi(u))$ of ϕ is injective, bijective or surjective, respectively. Here, $N_G(u) = \{v \mid uv \in E(G)\}$ denotes the (open) neighbourhood of a vertex u in a graph G. We say that ϕ is *locally injective*, *locally bijective* or *locally surjective* if ϕ is locally injective, locally bijective or locally surjective for every $u \in V(G)$. We denote the existence of these *locally constrained* homomorphisms by $G \xrightarrow{B} H$, $G \xrightarrow{I} H$ and $G \xrightarrow{S} H$, respectively; see Figure 1 for some examples.

The three locally constrained variants have been well studied in several settings over a long period of time. For example, locally injective homomorphisms are also known as *partial graph coverings* and are used in telecommunications [34], in distance constrained labelling [33] and as indicators of the existence of homomorphisms of derivative graphs [66]. Locally bijective homomorphisms originate from topological graph theory [5,65] and are more commonly known as *graph coverings*. They are used in distributed computing [2,3,8] and in constructing highly transitive regular graphs [6]. Locally surjective homomorphisms are sometimes called *colour dominations* [60]. They have applications in distributed computing [16,17] and in social science [31,71,74,76]. In the latter context they are known as *role assignments*, as we will explain in more detail below. (see Section 1.2).

We study the following three decision problems that take two graphs G and H as input and ask if there exists a homomorphism from G to H of one of the three local kinds.

Locally Bijective Homomorphism (LBHom) **Input:** Graphs G and H. **Question:** Does $G \xrightarrow{B} H$ hold?

Locally Injective Homomorphism (LIHom) **Input:** Graphs G and H. **Question:** Does $G \xrightarrow{I} H$ hold?

Locally Surjective Homomorphism (LSHom) **Input:** Graphs G and H. **Question:** Does $G \xrightarrow{I} H$ hold?

As before, we use the notation H-LBHom, H-LIHom and H-LSHom in the case when the host graph H is fixed.

Out of the three problems, only the complexity of H-LSHom has been completely classified, both for general graphs and bipartite graphs [38]. We refer to a series of papers [1,7,34,36,57,58,63] for polynomial-time solvable and NP-complete cases of H-LBHom and H-LIHom; see also the survey by Fiala and Kratochvíl [35]. Some more recent results include sub-exponential algorithms for H-LBHom, H-LIHom and H-LSHom on string graphs [68], complexity results for the list version of H-LSHom[29] and complexity results for H-LBHom for host graphs H that are multigraphs and/or have semi-edges [10,11,12,59].

In our paper we assume that both G and H are part of the input. We note a fundamental difference between locally injective homomorphisms on the one hand and locally bijective and surjective homomorphisms on the other. Namely, for connected graphs G and H, we must have $|V(G)| \geq |V(H)|$ if $G \xrightarrow{B} H$ or $G \xrightarrow{S} H$ (see Lemma 1 in Section 2.3). In contrast, H might be arbitrarily larger than G if $G \xrightarrow{I} H$ holds. For example, if we let G be a complete graph and H be a graph without self-loops, then $G \xrightarrow{I} H$ holds if and only if H contains a clique on at least $|V(G)|$ vertices.

The above difference is also reflected in the complexity results for the three problems under input restrictions. In fact, LIHom is closely related to the Subgraph Isomorphism problem and is usually the hardest problem. For example, LBHom is Graph Isomorphism-complete on chordal guest graphs, but polynomial-time solvable on interval guest graphs and LSHom is NP-complete on chordal guest graphs, but polynomial-time solvable on proper interval guest graphs [48]. In contrast, LIHom is NP-complete even on complete guest graphs G, which follows from a reduction from the CLIQUE problem via the aforementioned equivalence: $G \xrightarrow{I} H$ holds if and only if H contains a clique on at least $|V(G)|$ vertices.

Finally, we emphasize that the aforementioned polynomial-time result on HOM for guest graphs G with a core of bounded treewidth [20,42] does not carry over to any of the three locally constrained homomorphism problems. Indeed, LBHOM, LSHOM and LIHOM are NP-complete for guest graphs G of path-width at most 5, 4 and 2, respectively [19] (all three problems are polynomial-time solvable if G is a tree [19,39]). It is also known that LBHOM [56], LSHOM [60] and LIHOM [34] are NP-complete even if G is cubic and H is the complete graph K_4 on four vertices, but polynomial-time solvable if G has bounded treewidth and one of the two graphs G or H has bounded maximum degree [19].

1.2 An Application: Role Assignments

Locally surjective homomorphisms from a graph G to a graph H are known as H-role assignments in social network theory. We will include this topic in our investigation and provide some brief context.

Suppose that we are given a social network of individuals whose properties we aim to characterise. Can we assign each individual a role such that individuals with the same role relate in the same way to other individuals with some role, using exactly h different roles in total?

To formalise the above question, we model the network as a graph G, where vertices represent individuals and edges represent the existence of a relationship between two individuals. We now ask whether G has an *h-role assignment*, that is, a function f that assigns each vertex $u \in V(G)$ a *role* $f(u) \in \{1, \ldots, h\}$, such that $f(V(G)) = \{1, \ldots, h\}$ and for every two vertices u and v, if $f(u) = f(v)$ then $f(N_G(u)) = f(N_G(v))$.

Role assignments were introduced by White and Reitz [76] as *regular equivalences* and were called *role colourings* by Everett and Borgatti [31]. We observe that two adjacent vertices u and v may have the same role, that is, $f(u) = f(v)$ is allowed (so role assignments are not proper colourings). Hence, a connected graph G has an h-role assignment if and only if $G \xrightarrow{S} H$ for some connected graph H with $|V(H)| = h$, as long as we allow H to have self-loops (while we assume that G is a graph with no self-loops). The corresponding decision problem is the following:

> ROLE ASSIGNMENT
> **Input:** A graph G and an integer h.
> **Question:** Does G have an h-role assignment?

If h is fixed, then we denote the problem h-ROLE ASSIGNMENT. Whereas 1-ROLE ASSIGNMENT is trivial, 2-ROLE ASSIGNMENT is NP-complete [74]. In fact, h-ROLE ASSIGNMENT is NP-complete even for the following classes of graphs: planar graphs ($h \geq 2$) [72], cubic graphs ($h \geq 2$) [73], bipartite graphs ($h \geq 3$) [70], chordal graphs ($h \geq 3$) [52] and split graphs ($h \geq 4$) [25]. Very recently, Pandey, Raman and Sahlo [69] gave an $n^{\mathcal{O}(h)}$-time algorithm for ROLE ASSIGNMENT on general graphs and an $f(h)n^{\mathcal{O}(1)}$-time algorithm on forests.

1.3 Our Focus

We continue the line of study in [19] and focus on the following research question:

For which parameters of the guest graph do LBHom, LSHom *and* LIHom *become fixed-parameter tractable?*

We will also apply our new techniques towards answering this question for the Role Assignment problem. We first introduce some additional terminology. A graph parameter p *dominates* a parameter q if there is a function f such that $p(G) \leq f(q(G))$ for every graph G. If p dominates q but q does not dominate p, then p is *more powerful* than q. We denote this by $p \rhd q$. If p dominates q and q dominates p, then p and q are *equivalent*. If neither p dominates q nor q dominates p, then p and q are *incomparable*.

Given the paraNP-hardness results on LBHom, LSHom and LIHom for graph classes of bounded path-width [19], we will naturally consider a range of graph parameters that are less powerful than path-width. In this way we aim to increase our understanding of the (parameterized) complexity of LBHom, LSHom and LIHom.

For an integer $c \geq 1$, a *c-deletion set* of a graph G is a subset $S \subseteq V(G)$ such that every connected component of the graph $G \setminus S$ has at most c vertices.[6] The *c-deletion set number* $\mathrm{ds}_c(G)$ of G is the minimum size of a c-deletion set in G. If $c = 1$ we obtain the *vertex cover number* $\mathrm{vc}(G)$ of G. The c-deletion set number is closely related to the *fracture number* $\mathrm{fr}(G)$, introduced by Dvořák et al. [28], which is the minimum k such that G has a k-deletion set of size at most k. For a graph G, it holds that $\mathrm{fr}(G) \leq \mathrm{ds}_c(G)$ if $c \leq \mathrm{fr}(G) - 1$, and $\mathrm{ds}_c(G) \leq \mathrm{fr}(G)$ if $c \geq \mathrm{fr}(G)$. Hence, in particular it holds for every integer $c \geq 1$ that $\mathrm{fr}(G) \leq \max\{c, \mathrm{ds}_c(G)\}$. As we can take the graph formed as the disjoint union of arbitrarily many complete graphs on $c+1$ vertices (which has an arbitrarily large c-deletion set number while its fracture number is $c+1$), this inequality shows that for fixed c, $\mathrm{fr} \rhd \mathrm{ds}_c$. However, if c is not fixed, then fr and $c + \mathrm{ds}_c$ are equivalent, as in that case $c + \mathrm{ds}_c(G) \leq 2\mathrm{fr}(G)$ and $\mathrm{fr}(G) \leq c + \mathrm{ds}_c(G)$ holds for every graph G.

The fracture number is equivalent to several other well-studied graph parameters, such as the *vertex integrity*, introduced by Barefoot, Entringer and Swart [4] or the *safe number*, introduced by Fujita, MacGillivray and Sakuma [44]. The vertex integrity of a graph G is the minimum value $|X| + n^{\mathrm{c}}(G \setminus X)$ over all sets $X \subseteq V(G)$, where $n^{\mathrm{c}}(G \setminus X)$ denotes the size of a largest connected component of $G \setminus X$. Hence, the equivalence between the fracture number and the vertex integrity follows directly from their definitions, whereas the equivalence between the safe number and vertex integrity (and thus fracture number) is shown by Fujita and Furuya [43].

The *feedback vertex set number* $\mathrm{fv}(G)$ of a graph G is the size of a smallest set S such that $G \setminus S$ is a forest. We write $\mathrm{tw}(G)$, $\mathrm{pw}(G)$, $\mathrm{td}(G)$ and $n(G)$ for the treewidth, path-width, tree-depth[7] and number of vertices of a graph G, respectively; see [67] for more information. It is known that

$$\mathrm{tw} \rhd \mathrm{pw} \rhd \mathrm{td} \rhd \mathrm{fr} \rhd \mathrm{ds}_c(\text{fixed } c) \rhd \mathrm{vc} \rhd n,$$

[6] The graph $G \setminus S$ is the graph obtained from G by deleting all vertices of S; see Section 2 for any undefined terminology in this section.

[7] See Section 2 for the definitions of treewidth, path-width and tree-depth.

where the second relationship is proven in [9] and the others follow immediately from their definitions (see also Section 2.2). It is readily seen that

$$\text{tw} \triangleright \text{fv} \triangleright \text{ds}_2$$

and moreover, that fv is incomparable with the parameters pw, td, fr and, for every fixed $c \geq 3$, ds_c; consider, for example, a tree T of arbitrarily large path-width, whilst $\text{fv}(T) = 0$, and consider also the disjoint union G of arbitrarily many triangles, which has an arbitrarily large feedback vertex set number, but $\text{ds}_c(G) = 0$ for every $c \geq 3$.

guest graph parameter	LIHom	LBHom	LSHom
$\lvert V(G)\rvert$	XP, W[1]-hard [26]	FPT	FPT
vertex cover number	XP (Theorem 5), W[1]-hard	FPT	FPT
c-deletion set number (fixed c)	paraNP-c ($c \geq 2$) (Theorem 6)	FPT	FPT
fracture number	paraNP-c	FPT (Theorem 2)	FPT (Theorem 2)
tree-depth	paraNP-c	paraNP-c (Theorem 7)	paraNP-c (Theorem 7)
path-width	paraNP-c [19]	paraNP-c [19]	paraNP-c [19]
treewidth	paraNP-c	paraNP-c	paraNP-c
maximum degree	paraNP-c [34]	paraNP-c [56]	paraNP-c [60]
treewidth plus maximum degree	XP, W[1]-hard	XP [19]	XP [19]
feedback vertex set number	paraNP-c	paraNP-c (Theorem 8)	paraNP-c (Theorem 8)

Table 1. The results in purple are our new results and are annotated with the corresponding theorem numbers. The results in black are either known results, some of which are now also implied by our new results, or follow immediately from other results in the table. In particular, LIHom is W[1]-hard with parameter $\lvert V(G)\rvert$, as Clique is W[1]-hard when parameterized by the clique number [26], so we can let G be the complete graph in this case.

1.4 Our Results

We prove a number of new parameterized complexity results for LBHom, LSHom and LIHom by considering some property of the guest graph G as the parameter. In particular, we consider the graph parameters introduced in Section 1.3.

Our two main results, proven in Section 4, show that LBHom and LSHom are fixed-parameter tractable parameterized by $c + \text{ds}_c$, or equivalently, the fracture number $\text{fr}(G)$, of the guest graph G. In the same section, we also prove that Role Assignment is FPT when parameterized by the fracture number. Recall that td $\triangleright$ fr. However, assuming P $\neq$ NP, the FPT results for LBHom and LSHom involving the fracture number cannot be strengthened to the tree-depth of the guest graph. Namely, we prove in Section 6 that LBHom and LSHom are paraNP-complete when parameterized by the tree-depth of the guest graph. Recall also that pw $\triangleright$ td. Hence, these paraNP-completeness results imply the known paraNP-completeness results for path-width of the guest graph [19].

In Section 6 we also prove that LBHom and LSHom are paraNP-complete when parameterized by the feedback vertex set number of the guest graph. In fact, this result and the paraNP-hardness for tree-depth motivated us to consider the fracture number as a natural graph parameter for obtaining an FPT-algorithm.

For a fixed integer $k \geq 1$, the k-FoldCover problem is the restriction of LBHom to input pairs (G, H) where $\lvert V(G)\rvert = k\lvert V(H)\rvert$. This problem was introduced as the k-Graph

Covering problem by Bodlaender [8]. The aforementioned paraNP-completeness result of [19] on LBHom for pairs (G, H) where G has path-width 5 holds even for 3-FoldCover. In fact, this was the first proof for showing that k-FoldCover is NP-complete for some fixed integer $k \geq 1$.[8] Similarly, we will show in Section 6 that the two paraNP-completeness results for LBHom that we described above also hold even for the 3-FoldCover problem.

In Section 5 we prove that LIHom is in XP and W[1]-hard when parameterized by the vertex cover number of the guest graph, or equivalently, the c-deletion set number for $c = 1$. However, we show that the XP result for LIHom cannot be generalised to hold for $c \geq 2$. In fact, in Section 5, we will give a full complexity dichotomy. Namely, we determine the computational complexity of LIHom on graphs with c-deletion set number at most k for every fixed pair of integers c and k.

Table 1 summarizes the new and known results for LBHom, LSHom and LIHom.

1.5 Algorithmic Framework

Our FPT results for LBHom, LSHom and Role Assignment and our XP result for LIHom are proven via a new algorithmic framework (described in detail in Section 3). This framework involves a reduction to an integer linear program (ILP). We emphasize that in our framework the host graph H is *not* fixed, but part of the input. This is in contrast to other frameworks that include the locally constrained homomorphism problems (and that consequently work for more powerful graph parameters), such as the framework of locally checkable vertex partitioning problems [13,75] or the framework of Gerber and Kobler [46] based on (feasible) interval degree constraint matrices.

The main ideas behind our algorithmic ILP framework are as follows. Let G and H be the guest and host graphs, respectively. First, we observe that if G has a c-deletion of size at most k and there is a locally surjective homomorphism from G to H, then H must also have a c-deletion set of size at most k. However, it does not suffice to compute c-deletion sets D_G and D_H for G and H, guess a partial homomorphism h from D_G to D_H, and use the structural properties of c-deletion sets to decide whether h can be extended to a desired homomorphism from G to H. This is because a homomorphism from G to H does not necessarily map D_G to D_H. Moreover, even if it did, vertices in $G \setminus D_G$ can still be mapped to vertices in D_H. Consequently, components of $G \setminus D_G$ can still be mapped to more than one component of $H \setminus D_H$.

The above makes it difficult to decompose the homomorphism from G to H into small independent parts. To overcome this challenge, we prove that there are small sets D_G and D_H of vertices in G and H, respectively, such that every locally surjective homomorphism from G to H satisfies:

1. the pre-image of D_H is a subset of D_G,
2. D_H is a c'-deletion set for H for some c' bounded in terms of only $c + k$, and
3. all but at most k components of $G \setminus D_G$ have at most c vertices, whilst the treewidth of the remaining (and possibly large) components is bounded in terms of $c + k$.

[8] The proof can easily be adapted to hold for every $k \geq 3$, as observed by Klavík [55], while the cases $k = 1$ (which is equivalent to Graph Isomorphism) and $k = 2$ are still open.

As D_G and D_H are small, we can enumerate all possible homomorphisms from some subset of D_G to D_H. Condition 2 allows us to show that any locally surjective homomorphism from G to H can be decomposed into locally surjective homomorphisms from a small set of components of $G \setminus D_G$ (plus D_G) to one component of $H \setminus D_H$ (plus D_H). This enables us to formulate the question of whether a homomorphism from a subset of D_G to D_H can be extended to a desired homomorphism from G to H in terms of an ILP. Finally, Condition 3 allows us to efficiently compute the possible parts of the decomposition, that is, which (small) sets of components of $G \setminus D_G$ can be mapped to which components of $H \setminus D_H$.

2 Preliminaries

We use standard notation from graph theory, as can be found in e.g. [24]. Let G be a graph. We denote the vertex set and edge set of G by $V(G)$ and $E(G)$, respectively. Let $X \subseteq V(G)$ be a set of vertices of G. The *subgraph of G induced by X*, denoted $G[X]$, is the graph with vertex set X and edge set $\{uv \in E(G) \mid u, v \in X\}$. When the underlying graph is clear from the context, we will sometimes refer to an induced subgraph simply by its set of vertices. We use $G \setminus X$ to denote the subgraph of G induced by $V(G) \setminus X$. Similarly, for $Y \subseteq E(G)$ we let $G \setminus Y$ be the subgraph of G obtained by deleting all edges in Y from G.

For a graph G and a vertex $u \in V(G)$, we let $N_G(u) = \{v \mid uv \in E(G)\}$ and $N_G[v] = N_G(v) \cup \{v\}$ denote the open and closed neighbourhood of v in G, respectively. We let $\Delta(G)$ be the maximum degree of G. Recall that we assume that the guest graph G does not contain self-loops, while the host graph H is permitted to have self-loops. In this case, by definition, $u \in N_H(u)$ if $uu \in E(H)$.

2.1 Parameterized Complexity

In parameterized complexity [22,27,40], the complexity of a problem is studied not only with respect to the input size, but also with respect to some problem parameter(s). A *parameterized problem* Q is a subset of $\Omega^* \times \mathbb{N}$, where Ω is a fixed alphabet. Each instance of Q is a pair (I, κ), where $\kappa \in \mathbb{N}$ is called the *parameter*. A parameterized problem Q is *fixed-parameter tractable* (FPT) if there is an algorithm, called an FPT-*algorithm*, that decides whether an input (I, κ) is a member of Q in time $f(\kappa) \cdot |I|^{\mathcal{O}(1)}$, where f is a computable function and $|I|$ is the size of the input instance. The class FPT denotes the class of all fixed-parameter tractable parameterized problems.

A parameterized problem Q is FPT-*reducible* to a parameterized problem Q' if there is an algorithm, called an FPT-*reduction*, that transforms each instance (I, κ) of Q into an instance (I', κ') of Q' in time $f(\kappa) \cdot |I|^{\mathcal{O}(1)}$, such that $\kappa' \leq g(\kappa)$ and $(I, \kappa) \in Q$ if and only if $(I', \kappa') \in Q'$, where f and g are computable functions. By FPT-*time*, we denote time of the form $f(\kappa) \cdot |I|^{\mathcal{O}(1)}$, where f is a computable function. A problem Π is in $\mathsf{W}[1]$ if it is FPT-reducible to Independent Set parameterized by the solution size and $\mathsf{W}[1]$-hard if the latter problem is FPT-reducible to Π; it is easy to verify that $\mathsf{FPT} \subseteq \mathsf{W}[1]$. As an analogue to the conjecture that $\mathsf{P} \neq \mathsf{NP}$, it is widely believed that $\mathsf{FPT} \neq \mathsf{W}[1]$.

The class XP contains all parameterized problems that can be solved in $\mathcal{O}(|I|^{f(\kappa)})$ time, where f is a computable function. The class paraNP is the class of parameterized problems that can be solved by non-deterministic algorithms in $f(\kappa) \cdot |I|^{\mathcal{O}(1)}$ time, where f is a computable function.

2.2 Graph Parameters

A *tree-decomposition* $\mathcal{T}$ of a graph G is a pair (T, χ), where T is a tree and χ is a function that assigns each tree node t a set $\chi(t) \subseteq V(G)$ of vertices such that the following conditions hold:

(i) For every edge $uv \in E(G)$, there is a tree node t such that $u, v \in \chi(t)$.
(ii) For every vertex $v \in V(G)$, the set of tree nodes t with $v \in \chi(t)$ induces a non-empty subtree of T.

The sets $\chi(t)$ are called *bags* of the decomposition $\mathcal{T}$ and $\chi(t)$ is the bag associated with the tree node t. The *width* of a tree-decomposition (T, χ) is the size of a largest bag minus 1. The *treewidth* $\mathrm{tw}(G)$ of G is the minimum width over all tree-decompositions of G. If T is a path, then we obtain the notions of path-decomposition and *path-width.*

For a rooted forest F, the *closure* $C(F)$ is the graph with vertex set $V(F)$ such that two vertices u and v are adjacent in $C(F)$ if and only if u is an ancestor of v in F. We say that F is a *tree-depth decomposition* of a graph G if G is a subgraph of $C(F)$. The *depth* of F is equal to the height of F plus 1. The *tree-depth* $\mathrm{td}(G)$ of G is the minimum depth over all tree-depth decompositions of G.

We also need the following (well-known) fact on c-deletion sets.

Proposition 1 ([61]). *Let G be a graph and let k and c be natural numbers. Then, deciding whether G has a c-deletion set of size at most k is fixed-parameter tractable parameterized by $k + c$.*

Definition 1. *A* (k, c)-extended deletion set *for a graph G is a set $D \subseteq V(G)$ such that:*

- *every component of $G \setminus D$ either has at most c vertices or has a c-deletion set of size at most k and*
- *at most k components of $G \setminus D$ have more than c vertices.*

The following proposition summaries some known relationships between the parameters we consider.

Proposition 2 ([67]). *Let G be a graph and let k and c be natural numbers. Then:*

- *if G has a c-deletion set of size at most k, then $\mathrm{td}(G) \leq k + c$.*
- *if G has a (k', c)-extended deletion set of size at most k, then $\mathrm{td}(G) \leq k' + k + c$.*
- $\mathrm{tw}(G) \leq \mathrm{td}(G)$.

2.3 Locally Constrained Homomorphisms

Recall that we allow self-loops for the host graph H, but not for the guest graph G (see also Section 1). Here we show some basic properties of locally constrained homomorphisms.

Lemma 1. *Let G and H be non-empty connected graphs and let ϕ be a locally surjective homomorphism from G to H. Then ϕ is surjective.*

Proof. Suppose not, and let C be the set of vertices in $V(H) \setminus \phi(V(G))$. Note that $C \neq \emptyset$ (because otherwise ϕ is surjective) and $\phi(V(G)) \neq \emptyset$ (because G is non-empty). Because H is connected, there is an edge $uv \in E(H)$ such that $u \in V(C)$ and $v \in \phi(V(G))$. But then, the mapping $\phi_x : N_G(x) \to N_H(v)$ is not surjective for any vertex $x \in \phi^{-1}(v)$. □

Lemma 2. *Let G and H be non-empty connected graphs with a homomorphism ϕ from G to H and let $I \subseteq \phi(V(G))$. Let $P = \phi^{-1}(I)$ and $\phi_R = \phi|_P$. If ϕ is a locally injective, surjective or bijective homomorphism, then ϕ_R is a locally injective, surjective or bijective homomorphism, respectively, from $G[P]$ to $H[I]$.*

Proof. Clearly, ϕ_R is a homomorphism from $G[P]$ to $H[I]$ and since ϕ_R is a restriction of ϕ, it follows that if ϕ is locally injective, then so is ϕ_R. It remains to show that if ϕ is locally surjective, then so is ϕ_R. Suppose, for contradiction, that ϕ is locally surjective, but ϕ_R is not. Then there is a vertex $v \in P$ such that $\phi_R(N_G(v) \cap P) \subsetneq N_H(\phi_R(v)) \cap I$. However, since ϕ does not map any vertex in $V(G) \setminus P$ to a vertex of I, it follows that $\phi(N_G(v)) \cap I \subsetneq N_H(\phi(v)) \cap I$, so $\phi(N_G(v)) \neq N_H(\phi(v))$. Thus ϕ is not surjective, a contradiction. □

Lemma 3. *Let G and H be graphs, let $D \subseteq V(G)$, and let ϕ be a homomorphism from G to H. Then, for every component C_G of $G \setminus D$ such that $\phi(C_G) \cap \phi(D) = \emptyset$, there is a component C_H of $H \setminus \phi(D)$ such that $\phi(C_G) \subseteq C_H$. Moreover, if ϕ is locally injective/surjective/bijective, then $\phi|_{D\cup C_G}$ is a homomorphism from $G' = G[D \cup C_G]$ to $H' = H[\phi(D) \cup C_H]$ that is locally injective/surjective/bijective for every $v \in V(C_G)$.*

Proof. Suppose for a contradiction that this is not the case. Then, there is a component C_G of $G \setminus D$ and an edge $uv \in E(C_G)$ such that $\phi(u)$ and $\phi(v)$ are in different components of $H \setminus \phi(D)$. Therefore, $\phi(u)\phi(v) \notin E(H)$, contradicting our assumption that ϕ is a homomorphism.

Towards showing the second statement, first note that $\phi_R := \phi|_{D\cup C_G}$ is a homomorphism from G' to H'. Moreover, $N_G[v] = N_{G'}[v]$ for every vertex $v \in V(C_G)$, so if ϕ is locally injective/surjective/bijective for a vertex $v \in V(C_G)$, then so is ϕ_R. □

The following lemma is a basic but crucial observation showing that if $G \xrightarrow{S} H$ and G has a small c-deletion set, then so does H.

Lemma 4. *Let G and H be non-empty connected graphs, let $D \subseteq V(G)$ be a c-deletion set for G, and let ϕ be a locally surjective homomorphism from G to H. Then $\phi(D)$ is a c-deletion set for H.*

Proof. Suppose not, then there is a component C_H of $H \setminus \phi(D)$ such that $|C_H| > c$. By Lemma 1, it follows that ϕ is surjective and therefore $\phi^{-1}(C_H)$ is non-empty. Let $v \in \phi^{-1}(C_H)$. Then $v \notin D$ and therefore v is in some component C_G of $G \setminus D$. Lemma 3 implies that $\phi_R = \phi|_{D \cup C_G}$ is a homomorphism from $G[D \cup C_G]$ to $H[\phi(D) \cup C_H]$ that is locally surjective for every $v \in V(C_G)$.

Now $|V(C_G)| < |V(C_H)|$, so there must be a vertex in $V(C_H) \setminus \phi_R(C_G)$. Because C_H is connected, there is an edge $xy \in E(C_H)$ such that $x \in V(C_H) \setminus \phi_R(C_G)$ and $y \in \phi_R(V(C_G))$. But then, the mapping $\phi_z : N_G(z) \to N_H(y)$ is not surjective for any vertex $z \in \phi_R^{-1}(y)$. □

2.4 Integer Linear Programming

Given a set $\mathcal{X}$ of variables and a set $\mathcal{C}$ of linear constraints (i.e. inequalities) over the variables in $\mathcal{X}$ with integer coefficients, the task in the feasibility variant of *integer linear programming* (ILP) is to decide whether there is an assignment $\alpha : \mathcal{X} \to \mathbb{Z}$ of the variables satisfying all constraints in $\mathcal{C}$. We will use the following well-known result by Lenstra [62].

Proposition 3 ([32,41,53,62]). ILP *is* FPT *parameterized by the number of variables.*

3 Our Algorithmic Framework

In this section we present our main algorithmic framework that will allow us to show that LSHom, LBHom and Role Assignment are FPT, parameterized by $k + c$ when the guest graph has c-deletion set number at most k.

To illustrate the main ideas behind our framework, let us first explain these ideas for the examples of LSHom and LBHom. In this case we are given G and H and we know that G has a c-deletion set of size at most k. Because of Lemma 4, it then follows that if (G, H) is a yes-instance of LSHom or LBHom, then H also has a c-deletion set of size at most k. Informally, our next step, which is given in Section 3.1, is to compute a small (i.e. with size bounded by a function of $k + c$) set Φ of partial locally surjective homomorphisms such that

(1) every locally surjective homomorphism from G to H augments some $\phi_P \in \Phi$, and
(2) for every $\phi_P \in \Phi$, the domain of ϕ_P is a (k, c)-extended deletion set of G and the co-domain of ϕ_P is a c'-deletion set of H, where c' is bounded by a function of $k + c$.

Here and in what follows, we say that a function $\phi : V(G) \to V(H)$ *augments* (or is an *augmentation* of) a partial function $\phi_P : W_G \to W_H$, where $W_G \subseteq V(G)$ and $W_H \subseteq V(H)$ if $v \in W_G \Leftrightarrow \phi(v) \in W_H$ and $\phi|_{W_G} = \phi_P$. This allows us to reduce our problems to (boundedly many) subproblems of the following form: Given a (k, c)-extended deletion set D_G for G, a c'-deletion set D_H for H, and a locally surjective (respectively, bijective) homomorphism ϕ_P from D_G to D_H, find a locally surjective homomorphism ϕ from G to H that augments ϕ_P.

In Section 3.2 we will then show how to formulate this subproblem as an integer linear program and in Section 3.3 we will show that we can efficiently construct and solve the ILP for this subproblem. Importantly, our ILP formulation will allow us to solve a much more general problem, where the host graph H is not explicitly given, but defined in terms of a set of linear constraints. We will then exploit this in Section 4 to solve not only LSHom and LBHom, but also the Role Assignment problem.

3.1 Partial Homomorphisms for the Deletion Set

For a graph G and $m \in \mathbb{N}$ we let $\Delta_G^m := \{v \in V(G) \mid \deg_G(v) \geq m\}$. Our aim is to show in Lemma 7 that there is a small set Φ of partial homomorphisms such that every locally surjective (respectively, bijective) homomorphism from G to H augments some $\phi_P \in \Phi$ and, for every $\phi_P \in \Phi$, the domain of ϕ_P is a (k, c)-extended deletion set for G of size at most k and its co-domain is a c'-deletion set of size at most k for H. The main idea behind finding this set Φ is to consider the set of high degree vertices in G and H, that is, the sets Δ_G^{k+c} and Δ_H^{k+c}.

We start with the following lemma.

Lemma 5. *Let G be a graph. If G has a c-deletion set of size at most k, then the set Δ_G^{k+c} is a $kc(k+c)$-deletion set of size at most k. Furthermore, every subset of Δ_G^{k+c} is a $(k-|D|, c)$-extended deletion set of G.*

Proof. Let D be a c-deletion set of G of size at most k. Then every vertex $v \in V(G) \setminus D$ has degree at most $k+c-1$, as each of its neighbours lies either in its own component of $G \setminus D$ or in D. Hence $\Delta_G^{k+c} \subseteq D$ and therefore $|\Delta_G^{k+c}| \leq k$. Let $C_1, \ldots, C_m$ be the components of $G \setminus D$ that contain a vertex adjacent to a vertex in $D \setminus \Delta_G^{k+c}$. Since $|D \setminus \Delta_G^{k+c}| \leq k$ and every vertex in $D \setminus \Delta_G^{k+c}$ has degree at most $k+c-1$, we find that $m \leq k(k+c-1)$ and $|C_1 \cup \cdots \cup C_m \cup (D \setminus \Delta_G^{k+c})| \leq kc(k+c-1)+k \leq kc(k+c)$. Since every component in $G \setminus \Delta_G^{k+c}$ is either contained in a component of $G \setminus D$ or contained in $C_1 \cup \cdots \cup C_m \cup (D \setminus \Delta_G^{k+c})$, we find that Δ_G^{k+c} is a $kc(k+c)$-deletion set.

Let $D' \subseteq \Delta_G^{k+c} \subseteq D$. We will show that D' is a $(k-|D'|, c)$-extended deletion set of G. The components of $G \setminus D'$ that contain no vertices from $D \setminus D'$ are components of $G \setminus D$ and thus have size at most c. Consider a component C of $G \setminus D'$ that contains at least one vertex from $D \setminus D'$. Let $D_C = V(C) \cap (D \setminus D')$. Every component of $C \setminus D_C$ is a component of $G \setminus D$ and thus has size at most c. Moreover, D_C has size at most $|D \setminus D'| \leq k - |D'|$.

We conclude that every component of $G \setminus D'$ either has size at most c or has a c-deletion set of size at most $k - |D'|$. Furthermore, since there are at most $k - |D'|$ vertices in $\Delta_G^{k+c} \setminus D'$, and every component of $G \setminus D'$ that has size larger than c must contain a vertex of Δ_G^{k+c}, it follows that there are at most $k - |D'|$ components of $G \setminus D'$ that have size larger than c. This completes the proof. ◻

We use Lemma 5 in the proof of our next lemma, which shows that every locally surjective (respectively, bijective) homomorphism from G to H has to augment a locally surjective (respectively, bijective) homomorphism from some induced subgraph of $G[\Delta_G^{k+c}]$ to $H[\Delta_H^{k+c}]$. Intuitively, this holds because for every locally surjective homomorphism, only vertices of high degree in G can be mapped to a vertex of high degree in H and every vertex in H must have a pre-image in G.

Lemma 6. *Let G and H be non-empty connected graphs such that G has a c-deletion set of size at most k. If there is a locally surjective homomorphism ϕ from G to H, then there is a set $D \subseteq \Delta_G^{k+c}$ and a locally surjective homomorphism ϕ_P from $G[D]$ to $H[\Delta_H^{k+c}]$ such that ϕ augments ϕ_P. If ϕ is locally bijective, then $D = \Delta_G^{k+c}$ and ϕ_P is a locally bijective homomorphism.*

Proof. Observe that for a locally surjective homomorphism ϕ from G to H, the inequality $\deg_G(v) \geq \deg_H(\phi(v))$ holds for every $v \in V(G)$; moreover the equality $\deg_G(v) = \deg_H(\phi(v))$ holds in the locally bijective case. Since ϕ is surjective by Lemma 1, this implies that $\phi(\Delta_G^{k+c}) \supseteq \Delta_H^{k+c}$ (and if ϕ is locally bijective, then $\phi(\Delta_G^{k+c}) = \Delta_H^{k+c}$). Let $D = \phi^{-1}(\Delta_H^{k+c})$, so $D \subseteq \Delta_G^{k+c}$ (note that $D = \Delta_G^{k+c}$ if ϕ is locally bijective). Now $\phi|_D$ is a surjective map from D to Δ_H^{k+c}. Furthermore, $\phi(\Delta_G^{k+c} \setminus D) \cap \phi(D) = \phi(\Delta_G^{k+c} \setminus D) \cap \Delta_H^{k+c} = \emptyset$. Moreover, for every $v \in V(G) \setminus \Delta_G^{k+c}$, $\phi(v) \notin \Delta_H^{k+c} = \phi|_D(D)$, since $\deg_G(v) \geq \deg_H(\phi(v))$. Furthermore, $\phi|_D$ is a homomorphism from $G[D]$ to $H[\Delta_H^{k+c}]$ because ϕ is a homomorphism. Additionally, $\phi|_D$ is locally surjective (respectively, bijective) by Lemma 2. □

We are now ready to show that we can easily compute all possible pre-images of Δ_H^{k+c} in any locally surjective (respectively bijective) homomorphism from G to H.

Lemma 7. *Let G and H be non-empty connected graphs, and let k and c be two non-negative integers. For any $D \subseteq \Delta_G^{k+c}$, we can compute the set Φ_D of all locally surjective (respectively, bijective) homomorphisms ϕ_P from $G[D]$ to $H[\Delta_H^{k+c}]$ in $\mathcal{O}(|D|^{|D|+2})$ time. Furthermore, $|\Phi_D| \leq |D|^{|D|}$.*

Proof. Let $D \subseteq \Delta_G^{k+c}$ and suppose there is a surjective map $\phi_P : D \to \Delta_H^{k+c}$. Then for every vertex $v \in \Delta_H^{k+c}$, there must be a vertex $x \in D$ such that $\phi_P(x) = v$. Therefore $|\Delta_H^{k+c}| \leq |D|$, so if this condition fails, then we can immediately return that $\Phi_D = \emptyset$.

Otherwise, for each vertex of D, there are $|\Delta_H^{k+c}| \leq |D|$ possible choices for where a map $\phi_P : D \to \Delta_H^{k+c}$ could map this vertex. We can list all of the at most $|D|^{|D|}$ resulting maps in $\mathcal{O}(|D|^{|D|})$ time, and for each such map, we can check whether it is a locally surjective (respectively, bijective) homomorphism in $\mathcal{O}(|D|^2)$ time. □

3.2 ILP Formulation

In this section, we will show how to formulate the subproblem obtained in the previous subsection in terms of an ILP instance. More specifically, we will show that the following problem can be formulated in terms of an ILP: given a partial locally surjective (respectively, bijective) homomorphism ϕ_P from some induced subgraph D_G of G to some induced subgraph D_H of H, can this be augmented to a locally surjective (respectively, bijective) homomorphism from G to H? See Figure 2 for an illustration of the subproblem for the simpler case when D_G is a vertex cover of G and we are looking for a locally bijective homomorphism. Moreover, we will actually show that for this to work, the host graph H does not need to be given explicitly, but can instead be defined by a certain system of linear constraints.

We now sketch the main ideas behind our translation to ILP. We postpone the formal definitions of notions introduced below until later. Suppose that there is a locally surjective (respectively, bijective) homomorphism ϕ from G to H that augments ϕ_P. Because ϕ augments ϕ_P, Lemma 3 implies that ϕ maps every component C_G of $G \setminus V(D_G)$ entirely to some component C_H of $H \setminus V(D_H)$, moreover, $\phi|_{V(D_G) \cup V(C_G)}$ is already locally surjective (respectively, bijective) for every vertex $v \in V(C_G)$.

Our aim now is to describe ϕ in terms of its parts consisting of locally surjective (respectively, bijective) homomorphisms from extensions of D_G in G, that is, sets of components of $G \setminus D_G$ plus D_G, to simple extensions of D_H in H, that is, single components of $H \setminus D_H$ plus D_H.

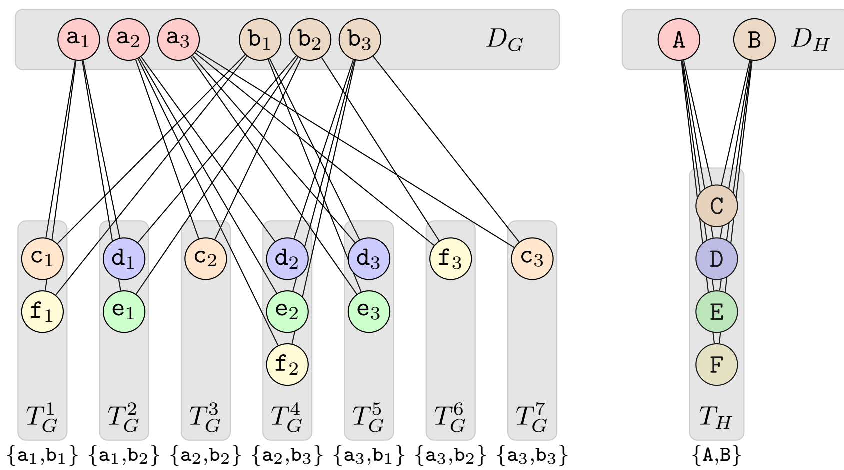

$\mathrm{Ext}_G^1 : T_G^1 + T_G^3 + T_G^7$ $(x_{\mathrm{Ext}_G^1 T_H} = 1) \rightarrow$ C
$\mathrm{Ext}_G^2 : T_G^2 + T_G^4 + T_G^5$ $(x_{\mathrm{Ext}_G^2 T_H} = 2) \rightarrow$ D, E
$\mathrm{Ext}_G^3 : T_G^1 + T_G^4 + T_G^6$ $(x_{\mathrm{Ext}_G^3 T_H} = 1) \rightarrow$ F

Fig. 2. A locally bijective homomorphism from a graph G (left) to a graph H (right), augmenting the partial homomorphism ϕ_P mapping the vertices of the vertex cover D_G into $D_H = \{$A,B$\}$. The ith vertex of G mapped to some vertex X of H is denoted $\mathtt{x}_i$. Vertices in $G \setminus D_G$ are grouped by type (e.g. $\{c_1\}$ and $\{f_1\}$ have type T_G^1), each T_G^i is characterised by the neighbours of its vertices, recalled below each column. Vertices in $H \setminus D_H$ all have the same type T_H. Rows Ext_G^i are extensions that can be minimally ϕ_P-B-mapped to T_H. This, in particular, means that each $\mathtt{a}_i$ and $\mathtt{b}_i$ must have a neighbour in Ext_G^i. This is because otherwise $\mathtt{a}_i$ or $\mathtt{b}_i$ is not mapped locally bijective to its image A or B, because A and B have a neighbor in T_H but the particular $\mathtt{a}_i$ or $\mathtt{b}_i$ has no neighbor in Ext_G^i. Note that the neighbour can then later be used as a pre-image of any vertex in {C,D,E,F}. Using Ext_G^1 once (for colour C), Ext_G^2 twice (for colours D and E) and Ext_G^3 once (for colour F) yields the given locally bijective homomorphism, and it can be verified that each $\mathtt{a}_i$ and $\mathtt{b}_i$ indeed has all four colours in its neighbourhood.

The main difficulty of doing the above comes from the fact that we need to ensure that ϕ is locally surjective (respectively, bijective) for every $d \in D_G$ and not only for the vertices within the components of $G \setminus D_G$. This is why we need to describe the parts of ϕ using sets of components of $G \setminus D_G$ and not just single components. However, as we will show, it will suffice to consider only minimal extensions of D_G in G, where an extension is minimal if no subset of it allows for a locally surjective (respectively bijective) homomorphism from it to some simple extension of D_H in H. The fact that we only need to consider minimal extensions is important for showing that we can compute the set of all possible parts of ϕ efficiently (see Section 3.3). Having shown this, we can create an ILP that has one variable $x_{\mathrm{EC}_G \mathrm{ES}_H}$ for every minimal extension EC_G and every simple extension ES_H such that there is a locally surjective (respectively, bijective) homomorphism from EC_G to ES_H that augments ϕ_P. The value of the variable $x_{\mathrm{EC}_G \mathrm{ES}_H}$ now corresponds to the number of times ϕ maps a minimal extension isomorphic to EC_G to a simple extension isomorphic to ES_H that augments ϕ_P. We can then use linear constraints on these variables to ensure that:

- H contains exactly the right number of extensions isomorphic to ES_H required by the assignment for $x_{\mathrm{EC}_G \mathrm{ES}_H}$,
- G contains exactly the right number of minimal extensions isomorphic to EC_G required by the assignment for $x_{\mathrm{EC}_G \mathrm{ES}_H}$ (if ϕ is locally bijective),
- G contains at least the number of minimal extensions isomorphic to EC_G required by the assignment for $x_{\mathrm{EC}_G \mathrm{ES}_H}$ (if ϕ is locally surjective), and
- for every simple extension ES_G of G that is not yet used in any part of ϕ, there is a homomorphism from ES_G to some simple extension of D_H in H that augments ϕ_P and is locally surjective for every vertex in $\mathrm{ES}_G \setminus D_G$ (if ϕ is locally surjective).

Together, these constraints ensure that there is a locally surjective (respectively, bijective) homomorphism ϕ from G to H that augments ϕ_P. See also Figure 3 for an illustration of the main ideas.

We are now ready to formalise these ideas. To do so, we need the following additional notation. Given a graph D, an *extension* of D is a graph EC containing D as an induced subgraph. It is *simple* if $\mathrm{Ext} \setminus D$ is connected, and *complex* in general. Given two extensions $\mathrm{EC}_1, \mathrm{EC}_2$ of D, we write $\mathrm{EC}_1 \sim_D \mathrm{EC}_2$ if there is an isomorphism τ from EC_1 to EC_2 with $\tau(d) = d$ for every $d \in D$. Then $\sim_D$ is an equivalence relation. Let the *types* of D, denoted $\mathcal{T}_D$, be the set of equivalence classes of $\sim_D$ of simple extensions of D. We write $\mathcal{T}_D^c$ to denote the set of types of D of size at most $|D| + c$, so

$$|\mathcal{T}_D^c| \leq \sum_{i=0}^{c} 2^{\binom{|D|+i}{2} - \binom{|D|}{2}} \leq c 2^{\binom{|D|+c}{2}}.$$

Given a complex extension EC of D, let C be a connected component of $\mathrm{EC} \setminus D$. Then C has type $T \in \mathcal{T}_D$ if $\mathrm{EC}[D \cup C] \sim_D T$ (depending on the context, we also say that the extension $\mathrm{EC}[D \cup C]$ has type T). The *type-count* of EC is the function $\mathrm{tc}_{\mathrm{EC}} : \mathcal{T}_D \to \mathbb{N}$ such that $\mathrm{tc}_{\mathrm{EC}}(T)$ for $T \in \mathcal{T}_D$ is the number of connected components of $\mathrm{EC} \setminus D$ with type T (in particular if EC is simple, the type-count is 1 for EC and 0 for other types). Note that two extensions are equivalent under $\sim_D$ if and only if they have the same type-counts; in particular, this implies that there is an isomorphism τ between the two extensions satisfying $\tau(d) = d$ for every $d \in D$. We write $\mathrm{EC} \preceq \mathrm{EC}'$ if $\mathrm{tc}_{\mathrm{EC}}(T) \leq \mathrm{tc}_{\mathrm{EC}'}(T)$ for all types $T \in \mathcal{T}_D$. If EC is an extension of D, we write $\mathcal{T}_D(\mathrm{EC}) = \{T \in \mathcal{T}_D \mid \mathrm{tc}_{\mathrm{EC}}(T) \geq 1\}$ for

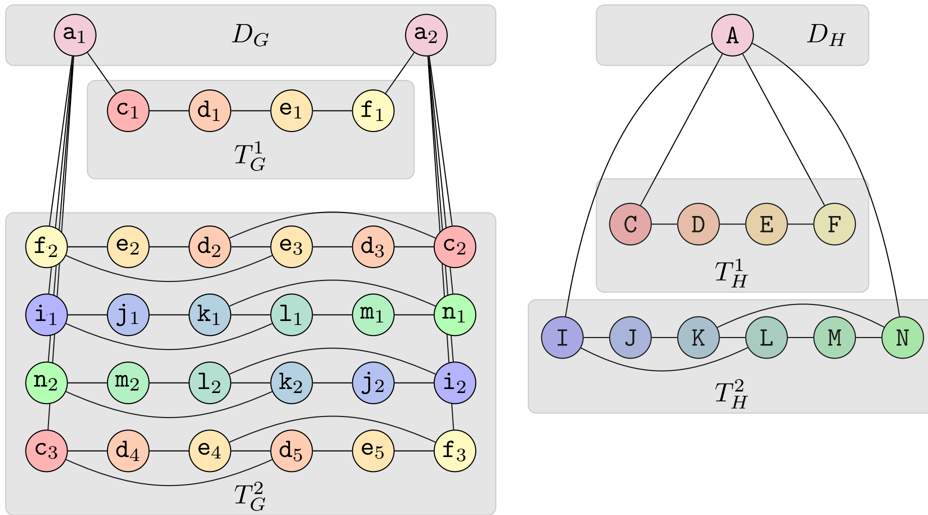

$$\mathrm{Ext}^1_G : 2 \times T^1_G \qquad \mathrm{Ext}^2_G : T^1_G + T^2_G \qquad \mathrm{Ext}^3_G : 2 \times T^2_G$$

$$x_{\mathrm{Ext}^1_G, T^1_H} = 0 \qquad x_{\mathrm{Ext}^2_G, T^1_H} = 1 \qquad x_{\mathrm{Ext}^3_G T^1_H} = 0 \qquad x_{\mathrm{Ext}^3_G T^2_H} = 1$$

Fig. 3. A locally surjective homomorphism from a graph G (left) to a graph H (right), where D_G is a 6-deletion set. The extensions $\mathrm{Ext}^1_G, \mathrm{Ext}^2_G, \mathrm{Ext}^3_G$ can be minimally ϕ_P-S-mapped to T^1_H; only Ext^3_G can also be minimally ϕ_P-S-mapped to T^2_H. Furthermore, T^1_G and T^2_G can each be weakly ϕ_P-S-mapped to some type in H (respectively, T^1_H and T^2_H). Using pair $(\mathrm{Ext}^2_G, T^1_H)$ and $(\mathrm{Ext}^3_G, T^2_H)$ once is sufficient to ensure that the mapping is locally surjective for each $\mathtt{a}_i$.

the *set of types of* EC and $\mathcal{E}_D(\mathrm{EC})$ for the set of simple extensions of D in EC. Moreover, for $T \in \mathcal{T}_D$, we write $\mathcal{E}_D(\mathrm{EC}, T)$ for the set of simple extensions of D in EC having type T.

A *target description* is a tuple (D_H, c, CH) where D_H is a graph, c is an integer and CH is a set of linear constraints over variables x_T for every $T \in \mathcal{T}^c_{D_H}$. A type-count for D_H is an integer assignment of the variables x_T. A graph H satisfies the target description (D_H, c, CH) if it is an extension of D_H, $\mathrm{tc}_H(T) = 0$ for $T \notin \mathcal{T}^c_{D_H}$, and setting $x_T = \mathrm{tc}_H(T)$ for all $T \in \mathcal{T}^c_{D_H}$ satisfies all constraints in CH. Note that target descriptions can be easily used to specify any graph H that has deletion set D_H into components of size at most c by using CH to specify the number of components of each type. However, allowing for arbitrary linear equations instead of specifying the graph H explicitly is much more flexible, and we will show how to employ this flexibility for the Role Assignment problem in Section 4.

In what follows, we assume that the following are given: the graphs D_G, D_H, an extension G of D_G, a target description $\mathcal{D} = (D_H, c, \mathrm{CH})$, and a locally surjective (respectively, bijective) homomorphism $\phi_P : D_G \to D_H$. Let EC_G be an extension of D_G with $\mathrm{EC}_G \preceq G$ and let $T_H \in \mathcal{T}^c_{D_H}$; note that we only consider $T_H \in \mathcal{T}^c_{D_H}$, because we assume that T_H is a type of a simple extension of a graph H that satisfies the target description $\mathcal{D}$. We say that EC_G can be

– *weakly ϕ_P-S-mapped* to the type T_H if there exists an augmentation $\phi : \mathrm{EC}_G \to T_H$ of ϕ_P such that ϕ is locally surjective for every $v \in \mathrm{EC}_G \setminus D_G$;
– *ϕ_P-S-mapped* (respectively, *ϕ_P-B-mapped*) to the type T_H if there exists an augmentation $\phi : \mathrm{EC}_G \to T_H$ of ϕ_P such that ϕ is locally surjective (respectively, locally bijective)
– *minimally ϕ_P-S-mapped* (respectively, *minimally ϕ_P-B-mapped*) to the type T_H if EC_G can be ϕ_P-S-mapped (respectively, ϕ_P-B-mapped) to T_H and no other extension EC'_G with $\mathrm{EC}'_G \preceq \mathrm{EC}_G$ can be ϕ_P-S-mapped (respectively, ϕ_P-B-mapped) to T_H.

Note that in the case of locally bijective homomorphisms, the two notions ϕ_P-B-mapped and minimally ϕ_P-B-mapped coincide. We also define the following sets:

– the set $\mathrm{wSM} = \mathrm{wSM}(G, \mathcal{D}, \phi_P)$ consists of all pairs (T_G, T_H) such that $T_G \in \mathcal{T}_{D_G}(G)$ can be weakly ϕ_P-S-mapped to $T_H \in \mathcal{T}^c_{D_H}$;
– the set $\mathrm{SM} = \mathrm{SM}(G, \mathcal{D}, \phi_P)$ consists of all pairs (EC_G, T_H) with $\mathrm{EC}_G \preceq G$ and $T_H \in \mathcal{T}^c_{D_H}$ such that EC_G can be minimally ϕ_P-S-mapped to T_H; and
– the set $\mathrm{BM} = \mathrm{BM}(G, \mathcal{D}, \phi_P)$ consists of all pairs (EC_G, T_H) with $\mathrm{EC}_G \preceq G$ and $T_H \in \mathcal{T}^c_{D_H}$ such that EC_G can be minimally ϕ_P-B-mapped to T_H.

See Figure 3 for an illustration of these notions.

We now build a set of linear constraints. To this end, besides variables x_T for $T \in T_H$, we introduce variables $x_{\mathrm{EC}_G T_H}$ for each $(\mathrm{EC}_G, T_H) \in \mathrm{SM}$ (respectively BM).

(S1) $\displaystyle\sum_{(\mathrm{EC}_G, T_H) \in \mathrm{SM}} \mathrm{tc}_{\mathrm{EC}_G}(T_G) * x_{\mathrm{Ext}_G T_H} \leq \mathrm{tc}_G(T_G)$ for every $T_G \in \mathcal{T}_{D_G}(G)$,

(B1) $\displaystyle\sum_{(\mathrm{EC}_G, T_H) \in \mathrm{BM}} \mathrm{tc}_{\mathrm{EC}_G}(T_G) * x_{\mathrm{EC}_G T_H} = \mathrm{tc}_G(T_G)$ for every $T_G \in \mathcal{T}_{D_G}(G)$,

(S2) $\displaystyle\sum_{\mathrm{EC}_G : (\mathrm{EC}_G, T_H) \in \mathrm{SM}} x_{\mathrm{EC}_G T_H} = x_{T_H}$ for every $T_H \in \mathcal{T}_{D_H}$,

(B2) $\displaystyle\sum_{\mathrm{EC}_G : (\mathrm{EC}_G, T_H) \in \mathrm{BM}} x_{\mathrm{EC}_G T_H} = x_{T_H}$ for every $T_H \in \mathcal{T}_{D_H}$, and

(S3) $\displaystyle\sum_{(T_G, T_H) \in \mathrm{wSM}} x_{T_H} \geq 1$ for every $T_G \in \mathcal{T}_{D_G}(G)$.

We refer to Figure 3 for an illustration.

Lemma 8. *Let D_G and D_H be graphs, let G be an extension of D_G and let $\mathcal{D} = (D_H, c, \mathrm{CH})$ be a target description. Moreover, let $\phi_P : V(D_G) \to V(D_H)$ be a locally surjective (respectively, bijective) homomorphism from D_G to D_H. There exists a graph H satisfying $\mathcal{D}$ and a locally surjective (respectively, bijective) homomorphism ϕ from G to H augmenting ϕ_P if and only if the equation system* (CH, S1, S2, S3) *(respectively,* (CH, B1, B2)) *admits a solution.*

Proof. Towards showing the forward direction of the claim, let H be a graph satisfying $\mathcal{D} = (D_H, c, \mathrm{CH})$ and let ϕ be a locally surjective (respectively, bijective) homomorphism that augments ϕ_P.

Consider $T_H \in \mathcal{T}_{D_H}(H)$ and let $\mathrm{ES}_H \in \mathcal{E}_{D_H}(H, T_H)$. Let $W_\phi(\mathrm{ES}_H) = G[\phi^{-1}(\mathrm{ES}_H)]$; note that $D_G \subseteq V(W_\phi(\mathrm{ES}_H))$ and therefore $W_\phi(\mathrm{ES}_H)$ is a (possibly) complex extension of

D_G. Then because of Lemma 2, we obtain that $\phi_R = \phi|_{W_\phi(\mathrm{ES}_H)}$ is a locally surjective (respectively, bijective) homomorphism from $W_\phi(\mathrm{ES}_H)$ to ES_H that augments ϕ_P. Moreover, because of Lemma 3, it follows that $W_\phi(\mathrm{ES}_H) \setminus D_G$ is the union of a set of components of $G - D_G$. Therefore, $W_\phi(\mathrm{ES}_H)$ can be ϕ_P-S-mapped (respectively, ϕ_P-B-mapped) to T_H. Moreover, if $W_\phi(\mathrm{ES}_H)$ can be ϕ_P-S-mapped to T_H, then $W_\phi(\mathrm{ES}_H)$ also contains a subgraph $W_\phi^{\min}(\mathrm{ES}_H)$ induced by D_G and a subset of components of $W_\phi(\mathrm{ES}_H) \setminus D_G$ that can be minimally ϕ_P-S-mapped to T_H. Note that $W_\phi^{\min}(\mathrm{ES}_H)$ is not uniquely defined. However, the concrete choice for $W_\phi^{\min}(\mathrm{ES}_H)$ does not matter. Informally, this is because $W_\phi^{\min}(\mathrm{ES}_H)$ ensures that some extension is ϕ-S-mapped to ES_H while the remaining components in $W_\phi(\mathrm{ES}_H)$ can still be weakly ϕ-S-mapped to ES_H.

Let $X_T = \{\, x_{T_H} : T_H \in \mathcal{T}_{D_H}(H) \,\}$, $X_M = \{\, x_{\mathrm{EC}_G T_H} : (\mathrm{EC}_G, T_H) \in \mathrm{SM} \,\}$ (respectively, $X_M = \{\, x_{\mathrm{EC}_G T_H} : (\mathrm{EC}_G, T_H) \in \mathrm{BM} \,\}$), and $X = X_T \cup X_M$. Let $\alpha : X \to \mathbb{N}$ be defined by setting:

- $\alpha(x_{T_H}) = \mathrm{tc}_H(T_H)$ and
- $\alpha(x_{\mathrm{EC}_G T_H}) = |\{\, \mathrm{Ext}_H \in \mathcal{E}_{D_H}(H, T_H) : \mathrm{tc}_{\mathrm{EC}_G} = \mathrm{tc}_{W_\phi^{\min}(\mathrm{ES}_H)} \,\}|$ (in the locally surjective case)
- $\alpha(x_{\mathrm{EC}_G T_H}) = |\{\, \mathrm{ES}_H \in \mathcal{E}_{D_H}(H, T_H) : \mathrm{tc}_{\mathrm{EC}_G} = \mathrm{tc}_{W_\phi(\mathrm{ES}_H)} \,\}|$ (in the locally bijective case)

We claim that the assignment α satisfies the equation system (CH, S1, S2, S3) (respectively, the equation system (CH, B1, B2)). Because H satisfies $\mathcal{D}$, it follows that α satisfies CH.

We start by showing the claim for the locally surjective case. Towards showing that (S3) is satisfied, consider a type $T_G \in T_{D_G}(G)$ and let $\mathrm{ES}_G \in \mathcal{E}_{D_G}(G, T_G)$. Then, because of Lemma 3, the mapping $\phi|_{\mathrm{ES}_G}$ maps ES_G to some type $T_H \in \mathcal{T}_{D_H}(H)$ and shows that ES_G can be weakly ϕ_P-S-mapped to T_H.

Towards showing (S1), let $T_G \in \mathcal{T}_{D_G}(G)$. Because of Lemma 3, every simple extension $\mathrm{ES}_G \in \mathcal{E}_{D_G}(G, T_G)$ satisfies $\phi(\mathrm{ES}_G) \subseteq \mathrm{ES}_H$ for some simple extension ES_H of D_H. In other words ES_G is contained in the pre-image of exactly one simple extension ES_H, showing that every simple extension $\mathrm{ES}_G \in \mathcal{E}_{D_G}(G, T_G)$ is counted at most once on the left side of the inequality in (S1) and therefore the left side is at most $\mathrm{tc}_G(T_G)$.

Towards showing (S2), let $T_H \in \mathcal{T}_{D_H}$. Then, because $V(W_\phi^{\min}(\mathrm{ES}_H)) \cap V(W_\phi^{\min}(\mathrm{ES}'_H)) = D_G$, i.e., $W_\phi^{\min}(\mathrm{ES}_H)$ and $W_\phi^{\min}(\mathrm{ES}'_H)$ share no components, for every two distinct $\mathrm{ES}_H, \mathrm{ES}'_H \in \mathcal{E}_{D_H}(H, T_H)$, we obtain:

$$\begin{aligned}
&\sum_{\mathrm{EC}_G : (\mathrm{EC}_G, T_H) \in \mathrm{SM}} \alpha(x_{\mathrm{EC}_G T_H}) \\
&= \sum_{\mathrm{EC}_G : (\mathrm{EC}_G, T_H) \in \mathrm{SM}} |\{\, \mathrm{ES}_H \in \mathcal{E}_{D_H}(H, T_H) : \mathrm{tc}_{\mathrm{EC}_G} = \mathrm{tc}_{W_\phi^{\min}(\mathrm{ES}_H)} \,\}| \\
&= \mathrm{tc}_H(T_H) \\
&= \alpha(x_{T_H}),
\end{aligned}$$

as required.

Finally, if ϕ is locally bijective, we only have to show (B1) and (B2), which can be shown similarly to (S1) and (S2), respectively. That is, (B1) can be shown similarly to (S1) by using the additional observation that due to the definition of α in terms of $W_\phi(\text{ES}_H)$ instead of $W_\phi^{\min}(\text{ES}_H)$, every simple extension ES_G also occurs in the pre-image of at least one simple extension ES_H. Moreover, (B2) can be shown in the same manner as (S2), since $V(W_\phi^{\min}(\text{ES}_H)) \cap V(W_\phi^{\min}(\text{ES}'_H)) = D_G$ also holds for every two distinct $\text{ES}_H, \text{ES}'_H \in \mathcal{E}_{D_H}(H, T_H)$.

Towards showing the reverse direction, let $\alpha : X \to \mathbb{N}$ be an assignment satisfying the equation system (CH, S1, S2, S3) (respectively, the equation system (CH,B1,B2)). Let H be the unique graph consisting of D_H and $\alpha(x_{T_H})$ extensions of D_H of type T_H for every $T_H \in \mathcal{T}_{D_H}$. Then H satisfies (D_H, c, CH) and $\text{tc}_H(T_H) = \alpha(x_{T_H})$.

We now define a function $\phi : V(G) \to V(H)$, which will be a locally surjective (respectively, locally bijective) homomorphism that augments ϕ_P as follows.

Let $\mathcal{A}$ be the multiset containing each pair $(\text{EC}_G, T_H) \in \text{SM}$ (respectively, BM) exactly $\alpha(x_{\text{EC}_G T_H})$ times. Because of (S2) (respectively, (B2)), there is a bijection γ_{T_H} between $\mathcal{A}_{T_H} = \{\, \text{EC}_G : (\text{EC}_G, T_H) \in \mathcal{A} \,\}$ and $\mathcal{E}_{D_H}(H, T_H)$ for every $T_H \in \mathcal{T}_{D_H}(H)$. Let γ be the bijection between $\mathcal{A}$ and the extensions $\mathcal{E}_{D_H}(H)$ given by $\gamma((\text{EC}_G, T_H)) = \gamma_{T_H}(\text{EC}_G)$.

Since the proof now diverges quite significantly for the locally surjective and locally bijective cases, we start by giving the proof for the former case and then show how to adapt the proof in latter (easier) case.

Because of (S1), there is a function β from $\mathcal{A}$ to the complex extensions of G such that:

- $\text{tc}_{\beta((\text{EC}_G, T_H))} = \text{tc}_{\text{EC}_G}$ for every $(\text{EC}_G, T_H) \in \mathcal{A}$,
- $\beta(A) \cap \beta(A') = D_G$ for every two distinct A and A' in $\mathcal{A}$.

Let $A = (\text{EC}_G, T_H) \in \mathcal{A}$. Because $A \in \text{SM}$, there is a locally surjective homomorphism ϕ_A from $\beta(A)$ to $\gamma(A)$ that augments ϕ_P. Let $\mathcal{E}_\mathcal{A}$ be the set of simple extensions ES_G in $\mathcal{E}_{D_G}(G)$ for which there is an $A \in \mathcal{A}$ such that ES_G is an induced subgraph of $\beta(A)$. Moreover, let $\bar{\mathcal{E}}_\mathcal{A}$ be the set of all remaining simple extensions in $\mathcal{E}_{D_G}(G)$, i.e. the set of all simple extensions ES_G in $\mathcal{E}_{D_G}(G) \setminus \mathcal{E}_\mathcal{A}$. Consider a simple extension ES_G in $\bar{\mathcal{E}}_\mathcal{A}$. Then, because of (S3), there is a $T_H \in \mathcal{T}_{D_H}(H)$ and a corresponding simple extension $\text{ES}_H \in \mathcal{E}_{D_H}(H, T_H)$ such that there is a homomorphism ϕ_{ES_G} from ES_G to ES_H that augments ϕ_P, which is locally surjective for every $v \in V(\text{ES}_G - D_G)$. We are now ready to define $\phi : V(G) \to V(H)$. That is, we set $\phi(v)$ to be equal to:

- $\phi_P(v)$ if $v \in D_G$,
- $\phi_{\text{Ext}_G}(v)$ if $v \in V(\text{ES}_G)$ for some simple extension $\text{ES}_G \in \bar{\mathcal{E}}_\mathcal{A}$, and
- $\phi_A(v)$ if $v \in V(\text{ES}_G - D_G)$ for some $\text{ES}_G = \beta(A)$ and $A \in \mathcal{A}$.

It remains to show that ϕ is a locally surjective homomorphism from G to H that augments ϕ_P. Clearly, ϕ augments ϕ_P by definition and because ϕ_{ES_G} does so too for every simple extension ES_G in $\bar{\mathcal{E}}_\mathcal{A}$, as does ϕ_A for every $A \in \mathcal{A}$.

Moreover, ϕ is a homomorphism, because every edge $\{u, v\} \in E(G)$ is contained in $G[\text{ES}_G]$ for some simple extension ES_G in $\mathcal{E}_{D_G}(G)$ and ϕ maps ES_G according to some homomorphism ϕ_{Ext_G} (if $\text{ES}_G \in \bar{\mathcal{E}}_\mathcal{A}$) or some homomorphism ϕ_A (otherwise). For basically the same reason, namely because every ϕ_{ES_G} and every ϕ_A is locally surjective for every vertex in $V(G) \setminus D_G$, we have that ϕ is locally surjective for every vertex $v \in V(G) \setminus D_G$.

Towards showing that ϕ is also locally surjective for every $d \in D_G$, let n_H be any neighbour of $\phi(d)$ in H. If $n_H \in \phi_P(D_G)$, then there is a neighbour n_G of d in D_G with $\phi(n_G) = n_H$, because ϕ_P is locally surjective. If on the other hand $n_H \in V(\text{ES}_H - D_H)$ for some $\text{ES}_H \in \mathcal{E}_{D_H}(H, T_H)$ with $T_H \in \mathcal{T}_{D_H}(H)$, then there is a neighbour n_G of d in $\beta(\gamma^{-1}(T_H))$ with $\phi(n_G) = n_H$, because ϕ (restricted to $\beta(\gamma^{-1}(T_H))$) is a locally surjective homomorphism from $\beta(\gamma^{-1}(T_H))$ to ES_H.

This completes the proof for the locally surjective case. We now complete the proof for the locally bijective case. First note that because of (B1), the function β from $\mathcal{A}$ to the complex extensions of G is bijective. Moreover, if $A = (\text{EC}_G, T_H) \in \mathcal{A}$, then because $A \in \text{BM}$, there is a locally bijective homomorphism ϕ_A from $\beta(A)$ to $\gamma(A)$ that augments ϕ_P. This now allows us to directly define $\phi : V(G) \to V(H)$. That is, we set $\phi(v)$ to be equal to:

- $\phi_P(v)$ if $v \in D_G$ and
- $\phi_A(v)$ if $v \in V(\text{EC}_G - D_G)$ for some $\text{EC}_G = \beta(A)$ and $A \in \mathcal{A}$.

It remains to show that ϕ is a locally bijective homomorphism from G to H that augments ϕ_P. Note that we can assume that ϕ is already a locally surjective homomorphism that augments ϕ_P, using the same arguments as for the locally surjective case. Thus it only remains to show that ϕ is also locally injective for every $d \in D_G$. Suppose not, then there are two distinct neighbours n_G and n'_G that are mapped to the same neighbour n_H of $\phi(d)$ in H. This is clearly not possible if both n_G and n'_G are in D_G because ϕ_P is locally bijective on D_G. Moreover, this can also not be the case if exactly one of n_G and n'_G is in D_G, because then $n_H \in V(D_H)$, but because ϕ augments ϕ_P, the other cannot be mapped to D_H. Therefore, we can assume that n_G and n'_G are outside of D_G. Let $\text{ES}_H \in \mathcal{E}_{D_H}(H)$ be the simple extension containing n_H. Then, n_G and n'_G must by mapped by $\phi_{\gamma^{-1}(\text{ES}_H)}$, but this is not possible because $\phi_{\gamma^{-1}(\text{ES}_H)}$ is locally bijective. $\square$

3.3 Constructing and Solving the ILP

The main aim of this section is to show the following theorem; see Definition 1 for a formal description of a (k, c)-extended deletion.

Theorem 1. *Let G be a graph, let D_G be a (k, c)-extended deletion set (respectively, a c-deletion set) of size at most k for G, let $\mathcal{D} = (D_H, c', \text{CH})$ be a target description and let $\phi_P : D_G \to D_H$ be a locally surjective (respectively, bijective) homomorphism from D_G to D_H. Deciding whether there is a locally surjective (respectively bijective) homomorphism that augments ϕ_P from G to any graph satisfying* CH *is* FPT *parameterized by $k + c + c'$.*

In order to prove Theorem 1, we need to show that we can construct and solve the ILP instance given in the previous section. The main ingredient for the proof of Theorem 1 is Lemma 12, which shows that we can efficiently compute the sets wSM, SM, and BM. We start by showing that the set $\mathcal{T}_{D_G}(G)$ can be computed efficiently and has small size.

Lemma 9. *For a graph G and a (k, c)-extended deletion set D_G of size at most k for G, it hold that $\mathcal{T}_{D_G}(G)$ has size at most $k + c2^{\binom{|D_G|+c}{2}}$. Moreover, the problem of computing $\mathcal{T}_{D_G}(G)$ and tc_G for a given graph G with (k, c)-extended deletion set D_G is* FPT *parameterized by $|D_G| + k + c$.*

Proof. Because $|\mathcal{T}_G(G) \setminus \mathcal{T}_G^c| \leq k$ and $|\mathcal{T}_G^c| \leq c2^{\binom{|D_G|+c}{2}}$, we obtain that $|\mathcal{T}_{D_G}(G)| \leq k + c2^{\binom{|D_G|+c}{2}}$. Moreover, we can compute $\mathcal{T}_{D_G}(G)$ starting from the empty set and adding a simple extension $G[D_G \cup C]$ for some component C of $G \setminus D_G$ if $G[D_G \cup C]$ is not equivalent with respect to $\sim_{D_G}$ to any element already added to $\mathcal{T}_{D_G}(G)$. Note that checking whether $G[D_G \cup C] \sim_D G[D_G \cup C']$ for two components C and C' of $G \setminus D_G$ is FPT parameterized by $|D_G| + k + c$, because $G[D_G \cup C]$ has treewidth at most $|D_G| + k + c$ for every component C of $G \setminus D_G$ (because of Proposition 2) and graph isomorphism is FPT parameterized by treewidth [64]. The same procedure can now also be used to compute all the non-zero entries of the function tc_G (i.e. the entries where $\text{tc}_G(T) \neq 0$), which provides us with a compact representation of tc_G. □

The following lemma is crucial for computing the sets SM and BM that are required to construct the ILP instance. Informally, we will show that if $(\text{EC}_G, \text{ES}_H) \in \text{SM}$ (or $(\text{EC}_G, \text{ES}_H) \in \text{BM}$), then EC_G consists of only boundedly many (in terms of some function of the parameters) components, which will allow us to enumerate all possibilities for EC_G in FPT-time.

Lemma 10. *Let D_G and D_H be graphs and let ϕ_P be a locally surjective (respectively, locally bijective) homomorphism from D_G to D_H. Moreover, let EC_G be an extension of D_G that can be minimally ϕ_P-S-mapped (respectively, minimally ϕ_P-B-mapped) to an extension ES_H of D_H. Then, $\text{EC}_G \setminus D_G$ consists of at most $|D_G||\text{ES}_H \setminus D_H|$ components.*

Proof. We first show the statement of the lemma for the case when ϕ_P is locally surjective and therefore EC_G can be minimally ϕ_P-S-mapped to ES_H. Let $\phi : V(\text{EC}_G) \to V(\text{ES}_H)$ be a locally surjective homomorphism that augments ϕ_P and that exists because EC_G can be ϕ_P-S-mapped to ES_H. Let EC'_G be an extension of EC_G with $\text{EC}'_G \preceq \text{EC}_G$. Then, because of Lemma 3, it follows that $\phi|_{\text{EC}'_G}$ is a homomorphism from EC'_G to ES_H that is locally surjective for every $v \in \text{EC}'_G \setminus D_G$. Therefore, $\phi|_{\text{EC}'_G}$ is a locally surjective homomorphism from EC'_G to ES_H if and only if EC'_G is such that $\phi|_{\text{EC}'_G}$ is locally surjective for every $d \in D_G$. That is, for every $d \in D_G$ and every neighbour n_H of $\phi(d)$ in ES'_H, there has to exist a neighbour n_G of d in EC'_G such that $\phi(n_G) = n_H$. Since this holds if $n_H \in D_H$ (because ϕ_P is a locally surjective homomorphism from D_G to D_H), we can assume that the above only has to hold for every $d \in D_G$ and $n_H \in \text{ES}_H \setminus D_H$.

Because ϕ is a locally surjective homomorphism from EC_G to ES_H, it follows that for every $d \in D_G$ and every neighbour n_H of $\phi(d)$ in ES_H, there is a component, say C_{d,n_H}, containing a neighbour n_G of d in EC_G such that $\phi(n_G) = \phi(n_H)$; note that because ϕ augments ϕ_P, it follows that $n_G \notin D_G$ because $n_H \notin D_H$.

Let EC'_G be the extension of D_G consisting of D_G and all components C_{d,n_H} for every $d \in D$ and $n_H \in \text{ES}_H \setminus D_H$ as above. Then, $\phi|_{\text{EC}'_G}$ is a locally surjective homomorphism from EC'_G to ES_H and since EC_G is minimally ϕ_P-S-mapped to ES_H and $\text{EC}'_G \preceq \text{EC}_G$, it follows that $\text{EC}'_G = \text{EC}_G$. However, $\text{EC}'_G \setminus D_G$ consists of at most one component for every $d \in D_G$ and every $n_H \in \text{ES}_H \setminus D_H$ and therefore it consists of at most $|D_G||\text{ES}_H \setminus D_H|$ components. This concludes the proof for the case when ϕ_P is locally surjective.

It remains to show the statement of the lemma for the case when ϕ_P is locally bijective and EC_G is minimally ϕ_P-B-mapped to ES_H. Let $\phi : V(\text{EC}_G) \to V(\text{ES}_H)$ be a locally bijective homomorphism that augments ϕ_P and that exists because EC_G can be ϕ_P-B-mapped to ES_H. Because ϕ is locally bijective, it is also locally surjective and therefore we can obtain the components C_{d,n_H} of $\mathrm{EC}_G \setminus D_G$ for $d \in D_H$ and $n_H \in \mathrm{ES}_H \setminus D_H$ using the same arguments as in the case when ϕ was locally surjective. As before, let EC'_G be the extension of D_G containing all components C_{d,n_H}. Then, as we showed above, $\phi|_{\mathrm{EC}'_G}$ is a locally surjective homomorphism from EC'_G to ES_H. Moreover, $\phi|_{\mathrm{EC}'_G}$ is also locally injective, because so is ϕ. Therefore, $\phi|_{\mathrm{EC}'_G}$ is a locally bijective homomorphism from EC'_G to ES_H. The latter implies, as EC_G can be minimally ϕ_P-B-mapped to ES_H, that $\mathrm{EC}_G = \mathrm{EC}'_G$. This concludes the proof of the lemma, because EC'_G consists of at most $|D_G||\mathrm{ES}_H \setminus D_H|$ components. □

The following proposition is a slight generalisation of [19, Theorem 4] and will allow us to efficiently decide whether an extension EC_G can be (weakly) S-mapped (respectively, B-mapped) to some extension ES_H.

Lemma 11 ([19, Theorem 4]). *Let G and H be graphs and let $\phi_P : D_G \to D_H$ be a locally surjective (respectively, bijective) homomorphism from D_G to D_H for some subgraphs D_G of G and D_H of H. Then deciding whether there is a locally surjective (respectively, bijective) homomorphism from G to H that augments ϕ_P can be achieved in $\mathcal{O}(|V(G)|((|V(H)2^{\Delta(H)})^{\mathrm{tw}(G)})^2\mathrm{tw}(G)\Delta(H))$ time and is therefore* FPT *parameterized by $\mathrm{tw}(G) + |V(H)|$.*

Proof. In [19, Theorem 4], the authors provided an algorithm that, given a graph G and a graph H, decides in $\mathcal{O}(|V(G)|((|V(H)2^{\Delta(H)})^{\mathrm{tw}(G)})^2\mathrm{tw}(G)\Delta(H))$ time whether there is a locally surjective homomorphism from G to H. The algorithm uses a standard dynamic programming approach on a tree decomposition of G, and it is straightforward to verify that the algorithm can be adapted with only minor modifications to an algorithm using the same run-time that decides whether there is a locally bijective homomorphism from G to H. Similarly, it is straightforward to adapt their algorithm to the case that one is additionally given a locally surjective (respectively, bijective) homomorphism ϕ_P from some induced subgraph D_G of G to some induced subgraph D_H of H and one only looks for a locally surjective (respectively, bijective) homomorphism from G to H that augments ϕ_P. □

The following corollary now follows directly from Lemma 11 and the definition of (weakly) S-mapped (respectively, B-mapped).

Corollary 1. *Let D_G and D_H be graphs and let ϕ_P be a locally surjective (respectively, bijective) homomorphism from D_G to D_H. Let EC_G be an extension of D_G having treewidth at most ω and let ES_H be an extension of D_H. Then, testing whether EC_G can be weakly ϕ_P-S-mapped, ϕ_P-S-mapped, or ϕ_P-B-mapped to ES_H is* FPT *parameterized by $\omega + |\mathrm{ES}_H|$.*

We are now ready to show that we can efficiently compute the sets wSM, SM, and BM, which is the last crucial step towards constructing the ILP instance.

Lemma 12. *Let G be a graph, let D_G be a (k, c)-extended deletion set (respectively, a c-deletion set) of size at most k for G, let $\mathcal{D} = (D_H, c', \mathrm{CH})$ be a target description and let ϕ_P be a locally surjective (respectively, bijective) homomorphism from D_G to D_H. Then, the sets $\mathrm{wSM} = \mathrm{wSM}(G, \mathcal{D}, \phi_P)$ and $\mathrm{SM} = \mathrm{SM}(G, \mathcal{D}, \phi_P)$ (respectively, the set $\mathrm{BM} = \mathrm{BM}(G, \mathcal{D}, \phi_P)$) can be computed in* FPT*-time parameterized by $k + c + c'$ and $|\mathrm{SM}|$*

(respectively, $|\mathrm{BM}|$) is bounded by a function depending only on $k+c+c'$. Moreover, the number of variables in the equation system (CH, S1, S2, S3) (respectively, (CH, B1, B2)) is bounded by a function depending only on $k+c+c'$.

Proof. We only show the lemma for the set SM, since the proof for the set wSM can be seen as a special case and the proof for the set BM is identical. Let $(\mathrm{EC}_G, T_H) \in \mathrm{SM}$. Then, EC_G is an extension of D_G with $\mathrm{EC}_G \preceq G$, $T_H \in \mathcal{T}_{D_H}^{c'}$, and EC_G can be minimally ϕ_P-S-mapped to T_H. Because EC_G can be minimally ϕ_P-S-mapped to ES_H, Lemma 10 implies that $\mathrm{EC}_G \setminus D_G$ consists of at most $\ell = |D_G||\mathrm{ES}_H \setminus D_H|$ components and, because $\mathrm{EC}_G \preceq G$, these are also components of $G \setminus D_G$. Therefore, there are at most $(|\mathcal{T}_{D_G}(G)|)^\ell$ non-isomorphic possibilities for EC_G, which together with Lemma 9 and the facts that $\ell \leq kc'$ and $|\mathcal{T}_{D_H}^{c'}| \leq c'2^{\binom{k+c'}{2}}$ shows that $|\mathrm{SM}| \leq (|\mathcal{T}_{D_G}(G)|)^\ell |\mathcal{T}_{D_H}^{c'}| \leq (k + c2^{\binom{k+c}{2}})^\ell (c'2^{\binom{k+c'}{2}})$. Therefore, $|\mathrm{SM}|$ is bounded by a function depending only on $k+c+c'$. Towards showing that we can compute SM is FPT-time parameterized by $k+c+c'$, first note that the set $\mathcal{T}_{D_G}(G)$ can be computed in FPT-time parameterized by $k+c$ using Lemma 9. Similarly, the set $\mathcal{T}_{D_H}^{c'}$ can be computed in FPT-time parameterized by $k+c'$ using the same idea as in Lemma 9. This now allows us to compute the set $\mathcal{A}$ containing all non-isomorphic possibilities for EC_G, i.e. the set of all extensions EC_G of D_G with $\mathrm{EC}_G \preceq G$ and $\sum_{T_G \in \mathcal{T}_{D_G}(G)} \mathrm{tc}_{\mathrm{EC}_G}(T_G) \leq \ell$ in FPT-time parameterized by $k+c+c'$, i.e. in time at most $(|\mathcal{T}_{D_G}(G)|)^\ell$. But then, SM is equal to the set of all pairs $(\mathrm{EC}_G, \mathrm{ES}_H) \in \mathcal{A} \times \mathcal{T}_{D_H}^{c'}$ such that EC_G can be minimally ϕ_P-S-mapped to ES_H. Moreover, for every such pair $(\mathrm{EC}_G, \mathrm{ES}_H)$ we can test in FPT-time parameterized by $k+c+c'$ whether EC_G can be ϕ_P-S-mapped to ES_H using Corollary 1, because the treewidth of EC_G is at most $k+c$ (Proposition 2). Therefore, we can compute SM by enumerating all pairs $(\mathrm{EC}_G, \mathrm{ES}_H) \in \mathcal{A} \times \mathcal{T}_{D_H}^{c'}$, testing for each of them whether EC_G can be ϕ_P-S-mapped to ES_H using Corollary 1, and keeping only those pairs $(\mathrm{EC}_G, \mathrm{ES}_H)$ such that EC_G can be ϕ_P-S-mapped to ES_H and EC_G is inclusion-wise minimal among all pairs $(\mathrm{EC}'_G, \mathrm{ES}_H)$. ◻

We are now ready to prove the main result of this subsection.

Theorem 1 (restated). *Let G be a graph, let D_G be a (k,c)-extended deletion set (respectively, a c-deletion set) of size at most k for G, let $\mathcal{D} = (D_H, c', \mathrm{CH})$ be a target description and let $\phi_P : D_G \to D_H$ be a locally surjective (respectively, bijective) homomorphism from D_G to D_H. Then, deciding whether there is a locally surjective (respectively bijective) homomorphism that augments ϕ_P from G to any graph satisfying* CH *is* FPT *parameterized by $k+c+c'$.*

Proof. We first compute the sets wSM and SM (respectively, the set BM), which because of Lemma 12 can be achieved in FPT-time parameterized by $k+c+c'$. This now allows us to construct the ILP instance $\mathcal{I}$ given by the equation system (CH,S1,S2,S3) (respectively, the equation system (CH,B1,B2)) in FPT-time parameterized by $k+c+c'$. Moreover, because the number of variables in $\mathcal{I}$ is bounded by a function of $k+c+c'$ and we can employ Proposition 3 to solve $\mathcal{I}$ in FPT-time parameterized by $k+c+c'$. Finally, because of Lemma 8, it follows that $\mathcal{I}$ has a solution if and only if there is a locally surjective (respectively bijective) homomorphism that augments ϕ_P from G to any graph satisfying CH, which completes the proof of the theorem. ◻

4 Applications of Our Algorithmic Framework

We are ready to show the main results of our paper, which can be obtained as an application of our framework given in the previous section. Our first result implies that LSHOM and LBHOM are FPT parameterized by the fracture number of the guest graph.

Theorem 2. *LSHOM and LBHOM are FPT parameterized by $k+c$, where k and c are such that the guest graph G has a c-deletion set of size at most k.*

Proof. Let G and H be non-empty connected graphs such that G has a c-deletion set of size at most k. Let $D_H = H[\Delta_H^{k+c}]$. We first verify whether H has a c-deletion set of size at most k using Proposition 1. If this is not the case, then we can return that there is no locally surjective (and therefore also no bijective) homomorphism from G to H because of Lemma 4. Therefore, we can assume in what follows that H also has a c-deletion set of size at most k, which together with Lemma 5 implies that $V(D_H)$ is a $kc(k+c)$-deletion set of size at most k for H. Therefore, using Lemma 9, we can compute tc_H in FPT-time parameterized by $k+c$. This now allows us to obtain a target description $\mathcal{D} = (D_H, c', \mathrm{CH})$ with $c' = kc(k+c)$ for H, i.e. $\mathcal{D}$ is satisfied only by the graph H, by adding the constraint $x_T = \mathrm{tc}_H(T_H)$ to CH for every simple extension type $T_H \in \mathcal{T}_{D_H}^{c'}$; note that $\mathcal{T}_{D_H}^{c'}$ can be computed in FPT-time parameterized by $k+c$ by Lemma 9.

Because of Lemma 6, we obtain that there is a locally surjective (respectively, bijective) homomorphism ϕ from G to H if and only if there is a set $D \subseteq \Delta_G^{k+c}$ and a locally surjective (respectively, bijective) homomorphism ϕ_P from $D_G = G[D]$ to D_H such that ϕ augments ϕ_P. Therefore, we can solve LSHOM by checking, for every $D \subseteq \Delta_G^{k+c}$ and every locally surjective homomorphism ϕ_P from $D_G = G[D]$ to D_H, whether there is a locally surjective homomorphism from G to H that augments ϕ_P. Note that there are at most 2^k subsets D and because of Lemma 7, we can compute the set Φ_D for every such subset in $\mathcal{O}(k^{k+2})$ time. Furthermore, due to Lemma 5, D is a $(k-|D|, c)$-extended deletion set of size at most k for G. Therefore, for every $D \subseteq \Delta_G^{k+c}$ and $\phi_p \in \Phi_D$, we can use Theorem 1 to decide in FPT-time, parameterized by $k+c$ (because $c' = kc(k+c)$), if there is a locally surjective (respectively, bijective) homomorphism from G to a graph satisfying $\mathcal{D}$ that augments ϕ_P. As H is the only graph satisfying $\mathcal{D}$, we proved the theorem. □

The proof of our next theorem is similar to the proof of Theorem 2. The major difference is that H is not given. Instead, we use Theorem 1 for a selected set of target descriptions. Each of these target descriptions enforces that graphs satisfying it have to be connected and have precisely h vertices, where h is part of the input for the ROLE ASSIGNMENT problem. Furthermore, we ensure that every graph H satisfying the requirements of ROLE ASSIGNMENT satisfies at least one of the selected target descriptions. The size of the set of considered target descriptions depends only on c and k, as it suffices to consider any small graph D_H and types of small simple extensions of D_H.

Theorem 3. *ROLE ASSIGNMENT is FPT parameterized by $k+c$, where k and c are such that G has a c-deletion set of size at most k.*

Proof. Let G be a non-empty connected graph such that G has a c-deletion set of size at most k and let $h \geq 1$ be an integer.

In order to use Theorem 1 in this case, we need to ensure that the target descriptions used enforce that H is connected and has h vertices. Therefore, for a fixed graph D on at most k vertices, we let CON_D be the set of all minimal sets $S \subseteq \mathcal{T}_D^{k+c}$ such that any extension H of D, that contains exactly the types in S, is connected.

Since $|\mathcal{T}_D^{k+c}|$ is bounded by $(2k+c)2^{\binom{2k+c}{2}}$, we can compute CON_D by considering every $S \subseteq \mathcal{T}_D^{k+c}$ and checking whether an extension $T \in \mathcal{T}_D$ of D containing precisely the types in S is connected. Since $|V(T)| \leq k + (k+c) \cdot |S|$ and checking connectivity takes linear time (using BFS or DFS) we can compute CON_D in time depending only on k and c. For $S \in \mathrm{CON}_D$, we set CH_S to be the set of equations containing $x_T \geq 1$ for every $T \in S$ and $|V(D_H)| + \sum_{T \in \mathcal{T}^c_{D_H}} (|V(T)| - |V(D_H)|) * x_T = h$. Note that for D and $S \in \mathrm{CON}_D$, any graph H satisfying the target description $(D, c+k, \mathrm{CH}_S)$ is connected and has h vertices.

If there is a connected graph H on h vertices and a locally surjective homomorphism ϕ from G to H, then by Lemma 6 there is a set $D \subseteq \Delta_G^{k+c}$ and a locally surjective homomorphism ϕ_P from $D_G = G[D]$ to $D_H = H[\Delta_H^{k+c}]$ such that ϕ augments ϕ_P. Note that by Lemmas 4 and 5, D_H is a $(k+c)$-deletion set of size at most k. This implies firstly that D_H is a graph on at most k vertices. Secondly, H is an extension of D_H such that $\mathrm{tc}_H(T) = 0$ for $T \notin T_{D_H}^{c+k}$ and, since H is also connected and has h vertices, H satisfies the target description $(D_H, c+k, \mathrm{CH}_S)$ for at least one $S \in \mathrm{CON}_{D_H}$.

Therefore, we can solve the Role Assignment problem by checking for every $D \subseteq \Delta_G^{k+c}$, every graph D_H on no more than k vertices, every $S \in \mathrm{CON}_{D_H}$ and every locally surjective homomorphism ϕ_P from $D_G = G[D]$ to D_H, whether there is a graph H satisfying the target description $(D_H, k+c, \mathrm{CH}_S)$ and a locally surjective homomorphism from G to H that augments ϕ_P. Note that there are at most 2^k subsets D. Furthermore, there are at most $k2^{\binom{k}{2}}$ graphs on at most k vertices and for each we can compute CON_D in time depending only on k and c. For each such graph D_H, there are at most $|\mathrm{CON}_D| \leq 2^{(2k+c)2^{\binom{2k+c}{2}}}$ subsets S to consider. Lastly, because of Lemma 7, for every $D \subseteq \Delta_G^{k+c}$ and any graph D_H on no more than k vertices, we can compute the set of locally surjective homomorphisms ϕ_P from $G[D]$ to D_H in time $\mathcal{O}(k^{k+2})$ time and there are at most $|D|^{|D|}$ ϕ_P to consider.

By Lemma 5, D is a $(k-|D|, c)$-extended deletion set of size at most k for G. Therefore, for every $D \subseteq \Delta_G^{k+c}$, every graph D_H on no more than k vertices, every $S \in \mathrm{CON}_{D_H}$ and every locally surjective homomorphism ϕ_P from $D_G = G[D]$ to D_H, we can employ Theorem 1 to decide in FPT-time parameterized by $k+c$, whether there is a graph H satisfying $(D_H, c+k, \mathrm{CH}_S)$ and a locally surjective homomorphism from G to H that augments ϕ_P. This completes the proof. ◻

5 Locally Injective Homomorphisms

The following result is well known. We include a proof for completeness.

Theorem 4 (Folklore). LIHom *is* W[1]*-hard parameterized by* $|V(G)|$*. In particular, it is* W[1]*-hard for all structural parameters of* G*.*

Proof. Let G be a complete graph on k vertices, and let H be an arbitrary graph. There exists a locally injective homomorphism ϕ from G to H if and only if H contains a clique K on k vertices. Indeed, for the forward direction, pick K to be the image of $V(G)$ under ϕ. Then $|K| = |V(G)| = k$ by the local injectivity of ϕ, and K is a clique. For the reverse direction, let ϕ be any bijection between $V(G)$ and K. The result follows from the fact that Clique is W[1]-hard. ◻

The locally injective case is more difficult in our setting since, in general, surjectivity helps to transfer structural parameters on G to similar structures on H (for example, in LSHom and LBHom the image of a deletion set is also a deletion set by Lemma 4). In LIHom however, and even in the restricted case of graphs with bounded vertex cover number, no such property can be used to help find the image of a vertex cover, and exponential-time enumerations appear to be necessary. On the positive side, once a partial homomorphism from a vertex cover of G to H has been found, our ILP framework can still be applied to map the remaining vertices in FPT-time. This leads to an XP-algorithm for vertex cover number (Theorem 5). Interestingly, this result does not extend to c-deletion set number for $c > 1$: even if the mapping of the deletion set can be guessed, the fact that the non-trivial remaining components must be mapped to distinct subgraphs of H makes the problem difficult (see Theorem 6).

Theorem 5. LIHom *is in* XP *parameterized by the vertex cover number of* G*.*

Proof. As for the surjective and bijective cases, we employ a two-step algorithm that first computes a suitable vertex cover, which can be done in time $1.2738^{\text{vc}(G)}$ [21], and guesses the image of the vertex cover and a partial homomorphism. Second, the algorithm finds a solution of an ILP which defines how to map the remaining vertices. The ILP only requires FPT-time, however the first step needs an exhaustive enumeration of subsets of H (in the injective case, the image of a vertex cover does not have to be a vertex cover), hence the XP running time.

We use the definitions of *types* and *extensions* from Section 3.2. Note that for any connected graph G with vertex cover D_G the connected components of $G \setminus D_G$ are single vertices. We can thus define the type of a vertex v of $G \setminus D_G$ to be the type of the simple extension $G[\{v\} \cup D_G]$. Note that two vertices u, v from $G \setminus D_G$ have the same type if they have the same neighbours in D_G. Hence, there are at most $2^{|D_G|}$ types in G.

In the following, let G be the given connected guest graph, and let H be the host graph. The main property of locally injective homomorphisms we use in this proof is that the size of the pre-image of any vertex of H can be bounded by twice the vertex cover number of G. This follows from the following claims.

Claim 1. *For any locally injective homomorphism $\phi : G \to H$ and every $h \in V(H)$ no two vertices in the pre-image $\phi^{-1}(h)$ share a neighbour.*

Proof of Claim. Assume that $\phi : G \to H$ is a locally injective homomorphism and there are vertices $h \in V(H)$, $v, v' \in \phi^{-1}(h)$ and $u \in V(G)$ such that u is adjacent to both v and v'. But then the restriction of ϕ to $N_G(u)$ is not injective contradicting the assumption. ⋄

Since for every vertex cover D_G and every vertex $u \in V(G)$, either $u \in D_G$ or every neighbour of u is in D_G we observe the following.

Claim 2. *If $R \subseteq V(G)$ is a subset with the property that no two vertices $u, v \in R$ share a neighbour then $|R| \leq 2 \cdot \mathrm{vc}(G)$.*

We use the property that pre-images have bounded size of locally injective homomorphisms to guess a partial homomorphism we aim to augment as well as determining which types of complex extensions our ILP needs to consider. In the following we show that picking suitable pre-images for the vertices of the host graph H yields a locally injective homomorphism.

Assume that D_G is a vertex cover of G and $D_H \subseteq V(H)$. We further fix $\phi_P : D_G \to D_H$ to be a partial homomorphism such that for every $h \in D_H$ no two vertices $u, v \in \phi_P^{-1}(h)$ share a neighbour in G. We say that a (possibly empty) subset R of $V(G \setminus D_G)$ is a *candidate pre-image* of a vertex $h \in V(H \setminus D_H)$ if the following two conditions hold:

1. $\phi_P(N_G(R)) \subseteq N_H(h) \cap D_H$ and
2. any two vertices in R do not share a neighbour.

Claim 3. *There is a locally injective homomorphism from G to H augmenting ϕ_P if and only if there is a partition $\{R_h \mid h \in V(H)\}$ of $V(G)$ such that $R_h = \phi_P^{-1}(h)$ for every $h \in D_H$, and R_h is a candidate pre-image for every $h \in V(H) \setminus D_H$.*

Proof of Claim. First assume that $\phi : V(G) \to V(H)$ is a locally injective homomorphism from G to H augmenting ϕ_P. Define $R_h = \phi^{-1}(h)$ for every $h \in V(H)$. Since ϕ is a mapping, we get that $\{R_h \mid h \in V(H)\}$ is a partition of $V(G)$. Since ϕ augments ϕ_P, we know that $\phi_P^{-1}(h) = \phi^{-1}(h)$. The latter implies that $\phi_P^{-1}(h) = R_h$ for every $h \in D_h$. Finally, let $h \in V(H) \setminus D_H$. As ϕ augments ϕ_P the set R_h must be disjoint from D_G. Hence, $N_G(R_h) \subseteq D_G$ and $\phi_P(N_G(R_h)) \subseteq D_H$. Additionally, since ϕ is a homomorphism, $\phi(N_G(R_h) = \phi_P(N_G(R_h)) \subseteq N_H(h)$ implying 1. Hence, by Claim 1, we find that R_h is a candidate pre-image.

On the other hand, assume that $\{R_h \mid h \in V(H)\}$ is a partition of $V(G)$ such that $R_h = \phi_P^{-1}(h)$ for every $h \in D_H$ and R_h is a candidate pre-image for every $h \in V(H) \setminus D_H$. Define a mapping $\phi : V(G) \to V(H)$ where $\phi(v) = h$ if $v \in R_h$. Note that ϕ is well defined as $\{R_h \mid h \in V(H)\}$ is a partition of $V(G)$. We now argue that ϕ is a homomorphism, i.e., we argue that for any two vertices $u, v \in V(G)$ it holds that if $uv \in E(G)$ then $\phi(u)\phi(v) \in E(H)$. For $u, v \in D_G$ this directly follows from ϕ_P being a homomorphism. For $u \in D_G$ and $v \notin D_G$ it follows from Condition 1 while $uv \notin E(G)$ for any $u, v \notin D_G$. We are left to prove that ϕ is locally injective.

Let $v \in V(G)$ and $u, u' \in N_G(v)$ be any vertices. We prove that $\phi(u) \neq \phi(u')$. If $u, u' \in D_G$, then u, u' share a neighbour which implies $\phi_P(u) \neq \phi_P(u')$ by choice of ϕ_P. If $u \in D_G$ and $u' \notin D_G$, then $\phi(u) \in D_H$ and $\phi(u') \notin D_H$ as ϕ is an augmentation of ϕ_P. If $u, u' \notin D_G$, then u, u' share a neighbour, and they cannot be in the same candidate pre-image by Condition 2. ◇

Let H^{D_H} be the graph obtained from H by deleting all edges that are not incident to D_H. In particular, D_H is a vertex cover of H^{D_H}.

Claim 4. *There is a locally injective homomorphism from G to H augmenting ϕ_P if and only if there is a locally injective homomorphism from G to H^{D_H} augmenting ϕ_P.*

Proof of Claim. By Claim 3 there is a locally injective homomorphism from G to H augmenting ϕ_P if and only if there is a partition $\{R_h \mid h \in V(H)\}$ of $V(G)$ such that $R_h = \phi_P^{-1}(h)$ for every $h \in D_H$ and R_h is a candidate pre-image for every $h \in V(H) \setminus D_H$. Observe that R_h being a canidate pre-image is independent on whether we consider the graph H or the graph H^{D_H}. This directly yields the claim using Claim 3. ⋄

We are now ready to describe our algorithm. Since in the first step we want to guess all possible partial homomorphisms we first observe the following. Assume that $\phi : V(G) \to V(H)$ is a locally injective homomorphism from G to H, that D'_G is a vertex cover of G of size $\mathrm{vc}(G)$, and that D_H is the image $\phi(D'_G)$ of D'_G. It is not necessarily true that the pre-image $\phi^{-1}(D_H)$ is D'_G. Hence our goal is to guess the image D_H of D'_G as well as the pre-image D_G of D_H. Note that guessing D_G could be done in FPT time by essentially only considering the type of pre-image of every vertex $h \in D_H$. However, we simplify the analysis here and guess D_G with XP many guesses.

Given two graphs G and H as input the algorithm proceeds as follows. Since we can consider any connected component of G individually, we assume that G is connected. First we compute any vertex cover D'_G of G of size $\mathrm{vc}(G)$. We proceed by considering every set $D_H \subseteq V(H)$ of size at most $\mathrm{vc}(G)$ and any set $\{R_h \mid h \in D_H\}$ of pairwise disjoint subsets of $D(G)$, for which the following conditions hold:

- $D'_G \subseteq D_G$ where we define D_G to be $\bigcup_{h \in D_H} R_h$,
- the partial mapping $\phi_P : D_G \to D_H$ defined by setting $\phi_P(v) = h$ whenever $v \in R_h$ is a homomorphism from D_G to D_H and
- for every $h \in D_H$ no two vertices $u, v \in R_h$ share a neighbour in G.

In total we consider $(|V(H)| + 1)^{\mathrm{vc}(G)}$ subsets D_H of $V(H)$. Furthermore, we only need to consider sets $\{R_h \mid h \in D_H\}$ for which $|R_h|$ is at most $2 \times \mathrm{vc}(G)$ by the condition that no two vertices $u, v \in R_h$ share a neighbour and Claim 2. Hence, we consider $(|V(G)| + 1)^{2 \cdot \mathrm{vc}(G)^2}$ such sets and check for each whether the conditions are met. In the next step of the algorithm we decide for every partial homomorphism ϕ_P considered in the first step, whether there is a locally injective homomorphism $\phi : V(G) \to V(H)$ from G to H augmenting ϕ_P. Note that by Claim 4 there is such a homomorphism if and only if there is a locally injective homomorphism $\phi : V(G) \to V(H^{D_H})$ from G to H^{D_H} augmenting ϕ_P. We encode the existence of such a homomorphism into an ILP.

Observe that D_G is a vertex cover of G of size at most $2 \cot \mathrm{vc}(G)^2$ (the set D_G consists of a set R_h of size at most $2 \cdot \mathrm{vc}(G)$ for each of the at most $\mathrm{vc}(G)$ vertices $h \in D_H$) and D_H is a vertex cover of H^{D_H}. Types considered in the following are the types of G and H^{D_H} with regards to vertex covers D_G and D_H. To construct the ILP we first compute tc_G and $\mathrm{tc}_{H^{D_H}}$ in polynomial time. Let IM be the set of pairs (EC_G, T_H), where $\mathrm{EC}_G \in \mathcal{E}_{D_G}(G)$ is a complex extension of D_G and $T_H \in \mathcal{T}_{D_H}(H^{D_H})$ is a simple extension of D_H, such that $V(\mathrm{EC}_G \setminus D_G)$ is a candidate pre-image of $T_H \setminus D_H$.

Observe that the size of every candidate pre-image is at most $2 \cdot \mathrm{vc}(G)$ by Condition 2 and Claim 2. Therefore, we get that the size of IM is at most $|\mathcal{T}_{D_G}(G)|^{2 \cdot \mathrm{vc}(G)} \cdot |\mathcal{T}_{D_H}(H^{D_H})| = 2^{4 \cdot \mathrm{vc}(G)^3} \cdot 2^{\mathrm{vc}(G)}$. We introduce a variable $x_{\mathrm{EC}_G T_H}$ for each pair $(\mathrm{EC}_G, T_H) \in \mathrm{IM}$ and let

$$X = \{x_{\mathrm{EC}_G T_H} : (\mathrm{EC}_G, T_H) \in \mathrm{IM}\}.$$

The variable $x_{\mathrm{EC}_G T_H}$ represents the number of vertices h in $V(H) \setminus D_H$ with type T_H whose pre-image is $V(\mathrm{EC}_G) \setminus D_G$. We introduce two types of constraints (see below). Constraint (I1) enforces that the pre-images R_h form a partition of $V(G \setminus D_G)$ by counting vertices of each type in each EC_G and checking that the sum corresponds to the count in G. Constraint (I2) corresponds to the fact that each vertex in H needs to be assigned a (possibly empty) pre-image (the number of pairs involving a type T_H must correspond to the type-count of T_H in H).

(I1) $$\sum_{(\mathrm{EC}_G, T_H) \in \mathrm{IM}} \mathrm{tc}_{\mathrm{EC}_G}(T_G) * x_{\mathrm{EC}_G T_H} = \mathrm{tc}_G(T_G) \text{ for every } T_G \in \mathcal{T}_{D_G}(G),$$

(I2) $$\sum_{\mathrm{EC}_G : (\mathrm{EC}_G, T_H) \in \mathrm{IM}} x_{\mathrm{EC}_G T_H} = \mathrm{tc}_H(T_H) \text{ for every } T_H \in \mathcal{T}_{D_H}(H^{D_H}).$$

We accept the input if any of the considered ILP's admits a solution, which can be checked in FPT time.

It remains to show that the algorithm is correct. If there is a locally injective homomorphism from G to H, then it must augment one of the partial homomorphisms $\phi_P : D_G \to D_H$ considered in the first step. By Claim 3 there is a partition $\{R_h \mid h \in V(H)\}$, such that R_h is a candidate pre-image for every $h \in V(H) \setminus D_H$, and hence $(R_h \cup D_G, D_H \cup \{h\}) \in \mathrm{IM}$. Hence, we obtain a feasible solution of the constraints (I1) and (I2) considering the following assignment α. The assignment $\alpha : X \to \mathbb{N}$ sets variable $x_{\mathrm{EC}_G T_H}$ to the number of vertices $h \in V(H) \setminus D_H$ of type T_H for whom the extension $R_h \cup D_G$ has type EC_G.

Conversely, let $\alpha : X \to \mathbb{N}$ be an assignment which is a solution to the constraints (I1) and (I2) for some partial homomorphism $\phi_P : D_G \to D_H$. For each $T_H \in \mathcal{T}_{D_H}(H^{D_H})$, we first partition the vertices of $V(H^{D_H}) \setminus D_H$ into sets $S_{\mathrm{EC}_G T_H}$ for every $(\mathrm{EC}_G, T_H) \in \mathrm{IM}$ such that $|S_{\mathrm{EC}_G T_H}| = \alpha(x_{\mathrm{EC}_G T_H})$. This is possible due to constraint (I2). We then choose a partition $\{R_h \mid h \in V(H^{D_H})\}$ of $V(G)$ as follows. For each $h \in D_H$, we set R_h to be $\phi_P^{-1}(h)$. For $h \in S_{\mathrm{EC}_G T_H}$, we choose $R_h \subseteq V(G) \setminus D_G$ such that $R_h \cup D_H$ has type EC_G. This is possible because of constraint (I1). Since $\{R_h \mid h \in V(H^{D_H})\}$ is a partition satisfying that $R_h = \phi_P^{-1}(h)$ for every $h \in D_H$ and R_h is a candidate pre-image for every $h \in V(H) \setminus D_H$, we conclude that there is a locally injective homomorphism augmenting ϕ_P by Claim 3. □

Corollary 2. *For any constant k,* LIHom *is polynomial-time solvable for graphs G where $\mathrm{ds}_1(G)$ is at most k.*

We actually obtain the following dichotomy for the complexity of LIHom, where the $c = 1$, $k \geq 1$ case is already given by Corollary 2.

Theorem 6. *Let $c, k \geq 1$. Then* LIHom *is polynomial-time solvable on guest graphs G where $\mathrm{ds}_c(G)$ is at most k if either $c = 1$ and $k \geq 1$ or $c = 2$ and $k = 1$; otherwise, it is* NP*-complete.*

Theorem 6 follows from Corollary 2 and the following three lemmas.

Lemma 13. LIHom *is polynomial-time solvable for graphs G where $\mathrm{ds}_2(G)$ is at most 1.*

Proof. Let G and H be connected graphs such that G has 2-deletion set number at most 1. If G has a 2-deletion set containing no vertices, then G contains at most two vertices, in which case we can solve LIHom in polynomial time. Otherwise, we can find a 2-deletion set $\{v\}$ in polynomial time by trying all possibilities for v. Let p be the number of edges in $G[N_G(v)]$ and let w be a vertex of H. We claim that there is a locally injective homomorphism ϕ from G to H such that $\phi(v) = w$ if and only if $H[N_H(w)]$ has a matching on at least p edges and $d_G(v) \leq d_H(w)$.

Indeed, if such a locally injective homomorphism ϕ exists, then $d_G(v) \leq d_H(w)$ because ϕ is locally injective. Furthermore, for every edge xy in $G[N_G(v)]$, the homomorphism ϕ maps the vertices x and y to adjacent vertices of $H[N_H(w)]$, and since ϕ is locally injective, it cannot map two vertices of $N_G(v)$ to the same vertex in $N_H(w)$. Therefore $H[N_H(w)]$ must have a matching on at least p edges.

Now suppose that $H[N_H(w)]$ has a matching M on at least p edges and $d_G(v) \leq d_H(w)$. For each edge xy in $G[N_G(v)]$, let $\phi(x)$ and $\phi(y)$ be the endpoints of an edge in M (choosing a different edge of M for each edge xy). For the remaining vertices $x \in N_G(v)$, assign the remaining vertices of $N_H(w)$ arbitrarily, such that no two vertices of $N_G(v)$ are assigned the same value (this can be done since $d_G(v) \leq d_H(w)$). Let $\phi(x) = w$ for all remaining vertices of G (i.e. the vertex v and all vertices non-adjacent to v that have a common neighbour with v). By construction, ϕ is a locally injective homomorphism from G to H.

The size of a maximum matching in a graph can be found in polynomial time [30]. Thus, by branching over the possible vertices $w \in V(H)$, we obtain a polynomial-time algorithm for LIHom. ◻

For the NP-hardness part of Theorem 6 we use a reduction from H'-PARTITION, defined in the following, with $H' = P_3$ in Lemma 14 and $H' = K_3$ in Lemma 15. Let H' be a fixed graph on h vertices. The H'-Partition problem takes as input a graph G' on hn vertices and the task is to decide whether the vertex set of G' can be partitioned into sets $V_1, \ldots, V_n$, each of size h, such that $G'[V_i]$ contains H' as a subgraph for all $i \in \{1, \ldots, n\}$. This problem is known to be NP-complete if $H' \in \{K_3, P_3\}$ [45,54].

Lemma 14. *For $c \geq 2$ and $k \geq 2$, LIHom is NP-hard on graphs G where $\mathrm{ds}_c(G)$ is k.*

Proof. We first consider the case when $k = 2$. Consider an instance G' of the P_3-Partition problem on $3n$ vertices, where $n \geq c$. We construct a graph G as follows. For $i \in \{1, \ldots, n\}$, add vertices a_i, b_i, c_i and d_i and edges a_ib_i and c_id_i. Then add vertices u and v and make u adjacent to a_i, b_i and d_i and v adjacent to a_i, c_i and d_i for all $i \in \{1, \ldots, n\}$. Finally, add the edge uv. Note that $\{u, v\}$ is a minimum-size c-deletion set for G since $\deg_G(u) = \deg_G(v) > c$. Now let H be the graph obtained from G' by adding two vertices u' and v' that are adjacent to all the vertices in $V(G')$ and to each other. For an illustration of the construction see Figure 4. We claim that there is a locally injective homomorphism ϕ from G to H if and only if G' is a yes-instance of the P_3-Partition problem.

Suppose that G' is a yes-instance of the P_3-Partition problem and, for $i \in \{1, \ldots, n\}$, let v_i^1, v_i^2, v_i^3 be the three vertices in V_i, such that v_i^2 is adjacent to v_i^1 and v_i^3 (v_i^1 may or may not be adjacent to v_i^3). Let $\phi : V(G) \to V(H)$ be the function such that $\phi(u) = u'$, $\phi(v) = v'$, and for $i \in \{1, \ldots, n\}$, $\phi(a_i) = v_i^1$, $\phi(b_i) = \phi(c_i) = v_i^2$ and $\phi(d_i) = v_i^3$. Note that ϕ is injective on $N_G(a_i)$ for every $i \in [n]$ as $\phi(u), \phi(v)$ and $\phi(b_i)$ are pairwise different. For

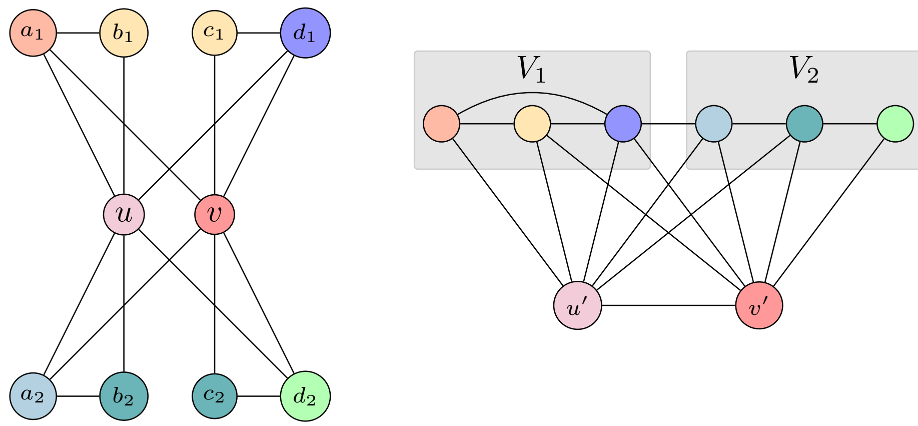

Fig. 4. An instance of LIHom consisting of the graph G (left) and the graph H (right) corresponding to the instance $G' = (V, E)$ of P_3-Partition, where $V = [6]$ and $E = \{12, 13, 23, 34, 45, 56\}$. As G' is a yes-instance of the P_3-Partition problem (partition is indicated by gray boxes), there is a locally injective homomorphism from G to H which is indicated by colours.

the same reason, ϕ is injective on $N_G(b_i), N_G(c_i)$ and $N_G(d_i)$ for every $i \in [n]$. To see that ϕ is injective on $N_G(u)$ observe that ϕ restricted to $V(G) \setminus \{c_i : i \in [n]\}$ is a bijection. Since u is not adjacent to any vertex in $\{c_i : i \in [n]\}$, this proves that ϕ is injective on $N_G(u)$. For v we conclude with a symmetric argument. Hence, ϕ is a locally injective homomorphism from G to H.

Now suppose that ϕ is a locally injective homomorphism from G to H. Now $deg_G(u) = deg_G(v) = 3n + 1$. Since H has $3n + 2$ vertices and ϕ is a locally injective homomorphism, it follows that $\phi(u)$ and $\phi(v)$ must be universal vertices in H. By symmetry, we may therefore assume that $\phi(u) = u'$ and $\phi(v) = v'$. Now u is adjacent to v and the vertices a_i, b_i and d_i for all $i \in \{1, \ldots, n\}$. Similarly, v is adjacent to u and the vertices a_i, c_i and d_i for all $i \in \{1, \ldots, n\}$. Since $deg_G(u) = 3n + 1$, and ϕ is locally injective, it follows that $\phi(\{a_i, b_i, d_i \mid i \in \{1, \ldots, n\}\}) = V(G')$. Similarly, since $deg_G(v) = 3n + 1$, it follows that $\phi(\{a_i, c_i, d_i \mid i \in \{1, \ldots, n\}\}) = V(G')$. Therefore $\phi(\{b_i \mid i \in \{1, \ldots, n\}\}) = \phi(\{c_i \mid i \in \{1, \ldots, n\}\})$. Renumbering the indices of the c_i and d_i vertices if necessary, we may therefore assume by symmetry that $\phi(b_i) = \phi(c_i)$ for all $i \in \{1, \ldots, n\}$. Now, for all $i \in \{1, \ldots, n\}$, the vertices a_i and b_i are adjacent in G, so $\phi(a_i)$ and $\phi(b_i)$ are adjacent in H. Furthermore the vertices c_i and d_i are adjacent in G, so $\phi(c_i) = \phi(b_i)$ and $\phi(d_i)$ are adjacent in H. We now set $V_i = \{\phi(a_i), \phi(b_i), \phi(d_i)\}$ and note that the V_i sets partition $V(G')$, and that $G'[V_i]$ contains a P_3 subgraph for all $i \in \{1, \ldots, n\}$. This completes the proof of the case when $k = 2$.

To extend the proof to graphs with c-deletion number $k > 2$, we add $(k - 1)$ universal vertices to H and replace u with a k-clique K each of whose vertices is adjacent to v and a_i, b_i and d_i for all $i \in \{1, \ldots, n\}$. □

Lemma 15. *For $c \geq 3$ and $k \geq 1$,* LIHom *is* NP*-hard on graphs G where $\mathrm{ds}_c(G)$ is k.*

Proof. We first consider the case when $k = 1$. Consider an instance G' of the K_3-PARTITION problem on $3n$ vertices, where $n \geq c$. Let H be the graph obtained from G' by adding a universal vertex w. Let G be the graph obtained by taking the disjoint union of n copies of K_3 and adding a universal vertex v. Note that $\{v\}$ forms a minimum-size c-deletion set for G since $\deg_G(v) > c$. We claim that there is a locally injective homomorphism ϕ from G to H if and only if G' is a yes-instance of the K_3-PARTITION problem.

Indeed, suppose there is such a ϕ. Since ϕ is locally injective and the graphs G and H each have $3n$ vertices, the universal vertex v must be mapped to a universal vertex of H; without loss of generality, we may therefore assume that $\phi(v) = w$. Since v and w are universal vertices of the same degree, it follows that ϕ is a bijection from $V(G)$ to $V(H)$. Every K_3 in the disjoint union part of G must therefore be mapped to a K_3 in $H \setminus \{w\} = G'$. Therefore G' is a yes-instance of the K_3-PARTITION problem.

Now suppose that G' is a yes-instance of the K_3-PARTITION problem. We let $\phi(v) = w$, and map the vertices of each K_3 in the disjoint union part of G to some V_i from the K_3-partition of H, mapping each K_3 to a different set V_i. Clearly this is a locally injective homomorphism. This completes the proof of the case when $k = 1$. To extend the proof to graphs with c-deletion number $k > 1$, we add $(k - 1)$ universal vertices to G and H. $\square$

6 Bounded Tree-depth and Feedback Vertex Set Number

By Theorem 6, we already obtained paraNP-hardness for LIHOM parameterized by tree-depth or feedback vertex set number. In this section we show that our tractability results for LSHOM and LBHOM cannot be significantly extended, since both problems become paraNP-hard parameterized by tree-depth. Furthermore, the reduction we give here also provides paraNP-hardness for both LSHOM and LBHOM parameterized by the feedback vertex set number. We show this by replacing cycles with stars in the reduction provided in [19] for path-width. This strengthens their result from path-width to tree-depth and feedback vertex set number.

Theorem 7. *LBHOM, or more specifically, 3-FOLDCOVER, and LSHOM are NP-complete on input pairs (G, H) where G has tree-depth at most 6 and H has tree-depth at most 4.*

Proof. First note that LBHOM, 3-FOLDCOVER and LSHOM are in NP. To prove NP-hardness for 3-FOLDCOVER and LSHOM we use a reduction from the 3-PARTITION problem. This problem takes as input a multiset A of $3m$ integers, denoted in what follows by $\{a_1, a_2, \dots, a_{3m}\}$, and a positive integer $b > 2$, such that $\frac{b}{4} < a_i < \frac{b}{2}$ for all $i \in \{1, \dots, 3m\}$ and $\sum_{1 \leq i \leq 3m} a_i = mb$. The task is to determine whether A can be partitioned into m disjoint sets $A_1, \dots, A_m$ such that $\sum_{a \in A_i} a = b$ for all $i \in \{1, \dots, m\}$. Note that the restrictions on the size of each element in A implies that each set A_i in the desired partition must contain exactly three elements, which is why such a partition $A_1, \dots, A_m$ is called a 3-*partition* of A. The 3-PARTITION problem is strongly NP-complete [45], that is, it remains NP-complete even if the problem is encoded in unary. If $b = 3$, then 3-PARTITION can be solved in polynomial time (since all a_i's are 1). Hence, we assume that $b > 3$.

We first prove NP-hardness for 3-FOLDCOVER, which implies NP-hardness for LBHOM. Given an instance (A, b) of 3-PARTITION, we construct an instance of 3-FOLDCOVER consisting of connected graphs G and H with $|V(G)| = 3|V(H)|$ as follows. To construct G we do as follows:

- Take $3m$ disjoint copies $S_1, \ldots, S_{3m}$ of $K_{1,b}$ (stars), one for each element of A. For each $i \in \{1, \ldots, 3m\}$, the vertices of S_i are labelled $c^i, u_1^i, \ldots, u_b^i$, where c_i is the vertex of degree b in S_i (the centre of the star).
- Add two new vertices p_j^i and q_j^i for each $i \in \{1, \ldots, 3m\}$, $j \in \{1, \ldots, b\}$, as well as two new edges $u_j^i p_j^i$ and $u_j^i q_j^i$.
- Add three new vertices x, y and z.
- Make x adjacent to the vertices $p_1^i, p_2^i \ldots, p_{a_i}^i$ and $q_1^i, q_2^i \ldots, q_{a_i}^i$ for every $i \in \{1, \ldots, 3m\}$.
- Make y adjacent to every vertex p_j^i that is not adjacent to x.
- Make z adjacent to every vertex q_j^i that is not adjacent to x.

See Figure 5 for an example.

To construct H, we take m disjoint copies $\tilde{S}_1, \ldots, \tilde{S}_m$ of $K_{1,b}$, where the vertices of each star $\tilde{S}_i$ are labelled $\tilde{c}^i, \tilde{u}_1^i, \ldots, \tilde{u}_b^i$. For each $i \in \{1, \ldots, m\}$ and $j \in \{1, \ldots, b\}$, we add two vertices $\tilde{p}_j^i$ and $\tilde{q}_j^i$ and make both of them adjacent to $\tilde{u}_j^i$. Finally, we add a vertex $\tilde{x}$ and make it adjacent to each of the vertices $\tilde{p}_j^i$ and $\tilde{q}_j^i$. This finishes the construction of H. Note that $|V(G)| = 3|V(H)|$. See again Figure 5 for an illustration.

We now show that there exists a locally bijective homomorphism from G to H if and only if (A, b) is a yes-instance of 3-Partition.

Let us first assume that there exists a locally bijective homomorphism ϕ from G to H. Since ϕ is a degree-preserving mapping, we must have $\phi(x) = \tilde{x}$. Moreover, since ϕ is locally bijective, the restriction of ϕ to $N_G(x)$ is a bijection from $N_G(x)$ to $N_H(\tilde{x})$. Again using the definition of a locally bijective mapping, this time considering the neighbourhoods of the vertices in $N_H(\tilde{x})$, we deduce that there is a bijection from the set $N_G^2(x) := \{u_j^i \mid 1 \leq i \leq 3m, 1 \leq j \leq a_i\}$, i.e. from the set of vertices in G at distance 2 from x, to the set $N_H^2(\tilde{x}) := \{\tilde{u}_j^k \mid 1 \leq k \leq m, 1 \leq j \leq b\}$ of vertices that are at distance 2 from $\tilde{x}$ in H.

For every $k \in \{1, \ldots, m\}$, we define a set $A_k \subseteq A$ such that A_k contains element $a_i \in A$ if and only if $\phi(u_1^i) \in \{\tilde{u}_1^k, \ldots, \tilde{u}_b^k\}$. Since ϕ is a bijection from $N_G^2(x)$ to $N_H^2(\tilde{x})$, the sets $A_1, \ldots, A_m$ are disjoint; moreover each element $a_i \in A$ is contained in exactly one of them. Since ϕ is degree preserving, each c^i has to be mapped onto a $\tilde{c}^j$ (note that we use the assumption $b > 3$ here). Additionally, since ϕ is locally bijective, for every $i \in \{1, \ldots, 3m\}$ there is a bijection from $N_G(c^i) = \{u_1^i, \ldots, u_b^i\}$ to $N_H(\tilde{c}^j) = \{\tilde{u}_1^j, \ldots, \tilde{u}_b^j\}$ for the $j \in \{1, \ldots, m\}$ for which $\phi(c^i) = \tilde{c}^j$. Combining this and the previous argument implies that $\sum_{a \in A_i} a = b$ for all $i \in \{1, \ldots, m\}$. Hence $A_1, \ldots, A_m$ is a 3-partition of A.

For the reverse direction, suppose that there exists a 3-partition $A_1, \ldots, A_m$ of A. We define a mapping ϕ as follows. We first set $\phi(x) = \phi(y) = \phi(z) = \tilde{x}$. Let $A_i = \{a_r, a_s, a_t\}$ be any set of the 3-partition. We map the vertices of S_r, S_s, S_t to the vertices of $\tilde{S}_i$ in the following way:

- $\phi(c_r) = \phi(c_s) = \phi(c_t) = \tilde{c}_i$;
- $\phi(u_j^r) = \tilde{u}_j^i$ for each $j \in \{1, \ldots, b\}$;
- $\phi(u_j^s) = \tilde{u}_{a_r+j}^i$ for each $j \in \{1, \ldots, a_s + a_t\}$;
- $\phi(u_j^s) = \tilde{u}_{a_r+j-b}^i$ for $j \in \{a_s + a_t + 1, \ldots, b\}$;
- $\phi(u_j^t) = \tilde{u}_{a_r+a_s+j}^i$ for each $j \in \{1, \ldots, a_t\}$; and
- $\phi(u_j^s) = \tilde{u}_{a_r+j-b}^i$ for $j \in \{a_t + 1, \ldots, b\}$.

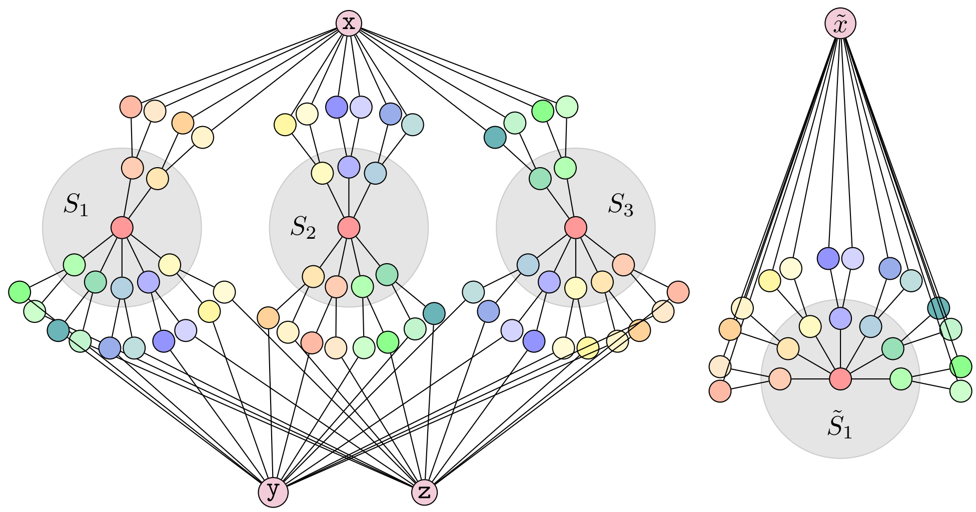

Fig. 5. An instance of LBHom consisting of the graph G (left) and the graph H (right) corresponding to the instance (A, b) of 3-Partition, where $A = \{2, 3, 2\}$ and $b = 7$. As (A, b) is a yes-instance of the 3-Partition problem, there is a locally bijective homomorphism from G to H which is indicated by colours.

It remains to map the vertices p^i_j and q^i_j for each $i \in \{1, \ldots, 3m\}$ and $j \in \{1, \ldots, b\}$. Let p^i_j, q^i_j be a pair of vertices in G that are adjacent to x, and let u^i_j be the second common neighbour of p^i_j and q^i_j. Suppose that $\tilde{u}^k_\ell$ is the image of u^i_j, i.e. suppose that $\phi(u^i_j) = \tilde{u}^k_\ell$. Then we map p^i_j and q^i_j to $\tilde{p}^k_\ell$ and $\tilde{q}^k_\ell$, respectively. We now consider the neighbours of y and z in G. By construction, the neighbourhood of y consists of the $2mb$ vertices in the set $\{p^i_j \mid a_{i+1} \leq j \leq b\}$, while $N_G(z) = \{q^i_j \mid a_{i+1} \leq j \leq b\}$.

Observe that $\tilde{x}$, the image of y and z, is adjacent to two sets of mb vertices: one of the form $\tilde{p}^k_\ell$, the other of the form $\tilde{q}^k_\ell$. Hence, we need to map half the neighbours of y to vertices of the form $\tilde{p}^k_\ell$ and half the neighbours of y to vertices of the form $\tilde{q}^k_\ell$ in order to make ϕ a locally bijective homomorphism. The same must be done with the neighbours of z. For every vertex $\tilde{u}^k_\ell$ in H, we do as follows. By construction, exactly three vertices of G are mapped to $\tilde{u}^k_\ell$, and exactly two of these vertices, say u^i_j and u^g_h, are at distance 2 from y in G. We set $\phi(p^i_j) = \tilde{p}^k_\ell$ and $\phi(p^g_h) = \tilde{q}^k_\ell$. We also set $\phi(q^i_j) = \tilde{q}^k_\ell$ and $\phi(q^g_h) = \tilde{p}^k_\ell$. This completes the definition of the mapping ϕ. For an illustration of the map ϕ, see Figure 5.

Since the mapping ϕ preserves adjacencies, it is a homomorphism. In order to show that ϕ is locally bijective, we first observe that the degree of every vertex in G is equal to the degree of its image in H, in particular,

$$d_G(x) = d_G(y) = d_G(z) = d_H(\tilde{x}) = 2mb.$$

From the above description of ϕ we get a bijection between the vertices of $N_H(\tilde{x})$ and the vertices of $N_G(v)$ for each $v \in \{x, y, z\}$. For every vertex p^i_j that is adjacent to x and u^i_j in G, its image $\tilde{p}^k_\ell$ is adjacent to the images $\tilde{x}$ of x and $\tilde{u}^k_\ell$ of u^i_j. For every vertex p^i_j

that is adjacent to y (respectively, z) and u_j^i in G, its image $\tilde{p}_\ell^k$ or $\tilde{q}_\ell^k$ is adjacent to $\tilde{x}$ of y (respectively, z) and $\tilde{u}_\ell^k$ of u_j^i. Hence the restriction of ϕ to $N_G(p_j^i)$ is bijective for every $i \in \{1, \ldots, 3m\}$ and $j \in \{1, \ldots, b\}$, and the same clearly holds for the restriction of ϕ to $N_G(q_j^i)$.

The vertices of each star S_i are mapped to the vertices of some star $\tilde{S}_k$ in such a way that the centres are mapped to centres. This, together with the fact that the image $\tilde{u}_\ell^k$ of every vertex u_j^i is adjacent to the images $\tilde{p}_\ell^k$ and $\tilde{q}_\ell^k$ of the neighbours p_j^i and q_j^i of u_j^i, shows that the restriction of ϕ to $N_G(u_j^i)$ is bijective for every $i \in \{1, \ldots, 3m\}$ and $j \in \{1, \ldots, b\}$. Finally, the neighbourhood of c^i is clearly mapped to the neighbourhood of $\phi(c^i)$ for every $i \in \{1, \ldots, 3m\}$. We conclude that ϕ is a locally bijective homomorphism from G to H.

In order to show that the tree-depth of G is at most 6, we observe that removing x, y and z yields a forest of depth 2. Similarly, H has tree-depth 4 since removing $\tilde{x}$ leaves a tree of depth 2. This completes the proof for 3-FoldCover and therefore LBHom.

In order to prove NP-hardness for LSHom we can use the same reduction as for LBHom. For this we can argue that there is a locally bijective homomorphism from G to H, for the graphs G and H constructed above, if and only if there is a locally surjective homomorphism from G to H. While the one direction is clear, if $G \xrightarrow{B} H$ then $G \xrightarrow{S} H$, for the converse direction we can make use of the following statement due to Kristiansen and Telle [60]:

(*) If $G \xrightarrow{S} H$ and $\mathrm{drm}(G) = \mathrm{drm}(H)$, then $G \xrightarrow{B} H$.

Here $\mathrm{drm}(G)$, $\mathrm{drm}(H)$ refers to the degree refinement matrix of G or H respectively, which is defined as follows. An *equitable partition* of a connected graph G is a partition of its vertex set into blocks $B_1, \ldots, B_k$ such that every vertex in B_i has the same number $m_{i,j}$ of neighbours in B_j. Then $\mathrm{drm}(G) = (m_{i,j})$ for $m_{i,j}$ corresponding to the coarsest equitable partition of G. We can easily observe that

$$\mathrm{drm}(G) = \mathrm{drm}(H) = \begin{pmatrix} 0 & 0 & 2mb & 0 \\ 0 & 0 & 2 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & b & 0 & 0 \end{pmatrix},$$

corresponding to the equitable partitions $B_1 = \{x, y, z\}$, $B_2 = \{u_j^i \mid i \in \{1, \ldots, 3m\}, j \in \{1, \ldots, b\}\}$, $B_3 = \{p_j^i, q_j^i \mid i \in \{1, \ldots, 3m\}, j \in \{1, \ldots, b\}\}$ and $B_4 = \{c^i \mid i \in \{1, \ldots, 3m\}\}$ in G and a similar equitable partition in H. Hence by (*) we find that $G \xrightarrow{B} H$ if and only if $G \xrightarrow{S} H$, completing the proof for LSHom. ◻

Theorem 8. LBHom*, or more specifically,* 3-FoldCover*, and* LSHom *are* NP*-complete on input pairs* (G, H) *where* G *and* H *have feedback vertex set number at most* 3 *and* 1*, respectively.*

Proof. To prove the statement we use the same reductions as in the proof of Theorem 7. This is sufficient, as the set $\{x, y, z\}$ is a feedback vertex set of G and the set $\{\tilde{x}\}$ is a feedback vertex set of H for graphs G and H defined in the proof of Theorem 7. ◻

7 Conclusions

We presented a fairly comprehensive picture concerning the parameterized complexity of three locally constrained graph homomorphism problems, namely LSHom, LBHom, and LIHom, when parameterized by some property of the guest graph. Our hardness results showed that the fracture number is the most suitable graph parameter of the guest graph for obtaining (parameterized) algorithms for these problems. We developed our algorithms through a general ILP-based framework. Besides the three locally constrained graph homomorphism problems, we also illustrated the applicability of our framework for the Role Assignment problem. This yielded three FPT results and one XP result in total.

As future resarch, we aim to extend our ILP-based framework. If successful, this will then also enable us to address the parameterized complexity of other graph homomorphism variants such as quasi-covers [37] and pseudo-covers [15,17,18]. We also recall the open problem from [19]: are LBHom and LSHom in FPT when parameterized by the treewidth of the guest graph plus the maximum degree of the guest graph?

References

1. Abello, J., Fellows, M.R., Stillwell, J.: On the complexity and combinatorics of covering finite complexes. Australasian Journal of Combinatorics **4**, 103–112 (1991)
2. Angluin, D.: Local and global properties in networks of processors. Proc. STOC 1980 pp. 82–93 (1980)
3. Angluin, D., Gardiner, A.: Finite common coverings of pairs of regular graphs. Journal of Combinatorial Theory, Series B **30**, 184–187 (1981)
4. Barefoot, C.A., Entringer, R., Swart, H.: Vulnerability in graphs – a comparative survey. Journal of Combinatorial Mathematics and Combinatorial Computing **1**, 13–22 (1987)
5. Biggs, N.J.: Algebraic Graph Theory. Cambridge University Press (1974)
6. Biggs, N.J.: Constructing 5-arc transitive cubic graphs. Journal of the London Mathematical Society II **26**, 193–200 (1982)
7. Bílka, O., Lidický, B., Tesař, M.: Locally injective homomorphism to the simple weight graphs. Proc. TAMC 2011, LNCS **6648**, 471–482 (2011)
8. Bodlaender, H.L.: The classification of coverings of processor networks. Journal of Parallel and Distributed Computing **6**, 166–182 (1989)
9. Bodlaender, H.L., Gilbert, J.R., Hafsteinsson, H., Kloks, T.: Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. Journal of Algorithms **18**, 238–255 (1995)
10. Bok, J., Fiala, J., Hliněný, P., Jedličková, N., Kratochvíl, J.: Computational complexity of covering multigraphs with semi-edges: Small cases. Proc. MFCS 2021, LIPIcs **202**, 21:1–21:15 (2021)
11. Bok, J., Fiala, J., Jedličková, N., Kratochvíl, J., Rzążewski, P.: List covering of regular multigraphs. Proc. IWOCA 2022, LNCS **13270**, 228–242 (2022)
12. Bok, J., Fiala, J., Jedličková, N., Kratochvíl, J., Seifrtová, M.: Computational complexity of covering disconnected multigraphs. Proc. FCT 2021, LNCS **12867**, 85–99 (2021)
13. Bui-Xuan, B., Telle, J.A., Vatshelle, M.: Fast dynamic programming for locally checkable vertex subset and vertex partitioning problems. Theoretical Computer Science **511**, 66–76 (2013)
14. Bulteau, L., Dabrowski, K.K., Köhler, N., Ordyniak, S., Paulusma, D.: An algorithmic framework for locally constrained homomorphisms. Proc. WG 2022, LNCS **13453**, 114–128 (2022)

15. Chalopin, J.: Local computations on closed unlabelled edges: The election problem and the naming problem. Proc. SOFSEM 2005, LNCS **3381**, 82–91 (2005)
16. Chalopin, J., Métivier, Y., Zielonka, W.: Local computations in graphs: The case of cellular edge local computations. Fundamenta Informaticae **74**, 85–114 (2006)
17. Chalopin, J., Paulusma, D.: Graph labelings derived from models in distributed computing: A complete complexity classification. Networks **58**, 207–231 (2011)
18. Chalopin, J., Paulusma, D.: Packing bipartite graphs with covers of complete bipartite graphs. Discrete Applied Mathematics **168**, 40–50 (2014)
19. Chaplick, S., Fiala, J., van 't Hof, P., Paulusma, D., Tesař, M.: Locally constrained homomorphisms on graphs of bounded treewidth and bounded degree. Theoretical Computer Science **590**, 86–95 (2015)
20. Chekuri, C., Rajaraman, A.: Conjunctive query containment revisited. Theoretical Computer Science **239**, 211–229 (2000)
21. Chen, J., Kanj, I.A., Xia, G.: Improved upper bounds for vertex cover. Theoretical Computer Science **411**, 3736–3756 (2010)
22. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer (2015)
23. Dalmau, V., Kolaitis, P.G., Vardi, M.Y.: Constraint satisfaction, bounded treewidth, and finite-variable logics. Proc. CP 2002, LNCS **2470**, 310–326 (2002)
24. Diestel, R.: Graph Theory, Graduate Texts in Mathematics, vol. 173. Springer Verlag, New York, 2nd edn. (2000)
25. Dourado, M.C.: Computing role assignments of split graphs. Theoretical Computer Science **635**, 74–84 (2016)
26. Downey, R.G., Fellows, M.R.: Fixed-parameter tractability and completeness II: on completeness for W[1]. Theoretical Computer Science **141**, 109–131 (1995)
27. Downey, R.G., Fellows, M.R.: Fundamentals of Parameterized Complexity. Texts in Computer Science, Springer (2013)
28. Dvořák, P., Eiben, E., Ganian, R., Knop, D., Ordyniak, S.: Solving integer linear programs with a small number of global variables and constraints. Proc. IJCAI 2017 pp. 607–613 (2017)
29. Dvořák, P., Krawczyk, M., Masařík, T., Novotná, J., Rzążewski, P., Zuk, A.: List locally surjective homomorphisms in hereditary graph classes. Proc. ISAAC 2022, LIPIcs **248**, 30:1–30:15 (2022)
30. Edmonds, J.: Paths, trees, and flowers. Canadian Journal of Mathematics **17**, 449–467 (1965)
31. Everett, M.G., Borgatti, S.P.: Role colouring a graph. Mathematical Social Sciences **21**, 183–188 (1991)
32. Fellows, M.R., Lokshtanov, D., Misra, N., Rosamond, F.A., Saurabh, S.: Graph layout problems parameterized by vertex cover. Proc. ISAAC 2008, LNCS **5369**, 294–305 (2008)
33. Fiala, J., Kloks, T., Kratochvíl, J.: Fixed-parameter complexity of lambda-labelings. Discrete Applied Mathematics **113**, 59–72 (2001)
34. Fiala, J., Kratochvíl, J.: Partial covers of graphs. Discussiones Mathematicae Graph Theory **22**, 89–99 (2002)
35. Fiala, J., Kratochvíl, J.: Locally constrained graph homomorphisms - structure, complexity, and applications. Computer Science Review **2**, 97–111 (2008)
36. Fiala, J., Kratochvíl, J., Pór, A.: On the computational complexity of partial covers of theta graphs. Discrete Applied Mathematics **156**, 1143–1149 (2008)
37. Fiala, J., Marek: Dichotomy of the H-Quasi-Cover problem. Proc. CSR 2013, LNCS **7913**, 310–321 (2013)
38. Fiala, J., Paulusma, D.: A complete complexity classification of the role assignment problem. Theoretical Computer Science **349**, 67–81 (2005)
39. Fiala, J., Paulusma, D.: Comparing universal covers in polynomial time. Theory of Computing Systems **46**, 620–635 (2010)

40. Flum, J., Grohe, M.: Parameterized Complexity Theory, Texts in Theoretical Computer Science. An EATCS Series, vol. XIV. Springer Verlag, Berlin (2006)
41. Frank, A., Tardos, É.: An application of simultaneous diophantine approximation in combinatorial optimization. Combinatorica **7**, 49–65 (1987)
42. Freuder, E.C.: Complexity of k-tree structured constraint satisfaction problems. Proc. AAAI 1990 pp. 4–9 (1990)
43. Fujita, S., Furuya, M.: Safe number and integrity of graphs. Discrete Applied Mathematics **247**, 398–406 (2018)
44. Fujita, S., MacGillivray, G., Sakuma, T.: Safe set problem on graphs. Discrete Applied Mathematics **215**, 106–111 (2016)
45. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., New York, NY, USA (1979)
46. Gerber, M.U., Kobler, D.: Algorithms for vertex-partitioning problems on graphs with fixed clique-width. Theoretical Computer Science **299**, 719–734 (2003)
47. Grohe, M.: The complexity of homomorphism and constraint satisfaction problems seen from the other side. Journal of the ACM **54**, 1:1–1:24 (2007)
48. Heggernes, P., van 't Hof, P., Paulusma, D.: Computing role assignments of proper interval graphs in polynomial time. Journal of Discrete Algorithms **14**, 173–188 (2012)
49. Hell, P., Nešetřil, J.: On the complexity of H-coloring. Journal of Combinatorial Theory, Series B **48**, 92–110 (1990)
50. Hell, P., Nešetřil, J.: The core of a graph. Discrete Mathematics **109**, 117–126 (1992)
51. Hell, P., Nešetřil, J.: Graphs and Homomorphisms. Oxford University Press (2004)
52. van 't Hof, P., Paulusma, D., van Rooij, J.M.M.: Computing role assignments of chordal graphs. Theoretical Computer Science **411**, 3601–3613 (2010)
53. Kannan, R.: Minkowski's convex body theorem and integer programming. Mathematics of Operations Research **12**, 415–440 (1987)
54. Kirkpatrick, D.G., Hell, P.: On the completeness of a generalized matching problem. Proc. STOC 1978 p. 240–245 (1978)
55. Klavík, P.: Extension Properties of Graphs and Structures. Ph.D. thesis, Charles University, Prague (2017)
56. Kratochvíl, J.: Regular codes in regular graphs are difficult. Discrete Mathematics **133**, 191–205 (1994)
57. Kratochvíl, J., Proskurowski, A., Telle, J.A.: Covering regular graphs. Journal of Combinatorial Theory, Series B **71**, 1–16 (1997)
58. Kratochvíl, J., Proskurowski, A., Telle, J.A.: On the complexity of graph covering problems. Nordic Journal of Computing **5**, 173–195 (1998)
59. Kratochvíl, J., Telle, J.A., Tesař, M.: Computational complexity of covering three-vertex multigraphs. Theoretical Computer Science **609**, 104–117 (2016)
60. Kristiansen, P., Telle, J.A.: Generalized H-coloring of graphs. Proc. ISAAC 2000, LNCS **1969**, 456–466 (2000)
61. Kronegger, M., Ordyniak, S., Pfandler, A.: Backdoors to planning. Artificial Intelligence **269**, 49–75 (2019)
62. Lenstra Jr., H.W.: Integer programming with a fixed number of variables. Mathematics of Operations Research **8**, 538–548 (1983)
63. Lidický, B., Tesař, M.: Complexity of locally injective homomorphism to the theta graphs. Proc. IWOCA 2010, LNCS **6460**, 326–336 (2010)
64. Lokshtanov, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Fixed-parameter tractable canonization and isomorphism test for graphs of bounded treewidth. SIAM Journal on Computing **46**, 161–189 (2017)
65. Massey, W.S.: Algebraic Topology: An Introduction. Harcourt, Brace and World (1967)

66. Nešetřil, J.: Homomorphisms of derivative graphs. Discrete Mathematics **1**, 257–268 (1971)
67. Nešetřil, J., Ossona de Mendez, P.: Sparsity: Graphs, Structures, and Algorithms, Algorithms and Combinatorics, vol. 28. Springer (2012)
68. Okrasa, K., Rzążewski, P.: Subexponential algorithms for variants of the homomorphism problem in string graphs. Journal of Computer and System Sciences **109**, 126–144 (2020)
69. Pandey, S., Raman, V., Sahlot, V.: Parameterizing role coloring on forests. Proc. SOFSEM 2021, LNCS **12607**, 308–321 (2021)
70. Pandey, S., Sahlot, V.: Role coloring bipartite graphs. Discrete Applied Mathematics **322**, 276–285 (2022)
71. Pekeč, A., Roberts, F.S.: The role assignment model nearly fits most social networks. Mathematical Social Sciences **41**, 275–293 (2001)
72. Purcell, C., Rombach, M.P.: On the complexity of role colouring planar graphs, trees and cographs. Journal of Discrete Algorithms **35**, 1–8 (2015)
73. Purcell, C., Rombach, M.P.: Role colouring graphs in hereditary classes. Theoretical Computer Science **876**, 12–24 (2021)
74. Roberts, F.S., Sheng, L.: How hard is it to determine if a graph has a 2-role assignment? Networks **37**, 67–73 (2001)
75. Telle, J.A., Proskurowski, A.: Algorithms for vertex partitioning problems on partial k-trees. SIAM Journal on Discrete Mathematics **10**(4), 529–550 (1997)
76. White, D.R., Reitz, K.P.: Graph and semigroup homomorphisms on networks of relations. Social Networks **5**, 193–235 (1983)