Algorithmica (2024) 86:944–1005
https://doi.org/10.1007/s00453-023-01178-0

Parameterised and Fine-Grained Subgraph Counting, Modulo 2

Leslie Ann Goldberg[1] · **Marc Roth**[1]

Received: 24 April 2023 / Accepted: 9 October 2023 / Published online: 2 November 2023
© The Author(s) 2023

Abstract

Given a class of graphs $\mathcal{H}$, the problem $\oplus\textsc{Sub}(\mathcal{H})$ is defined as follows. The input is a graph $H \in \mathcal{H}$ together with an arbitrary graph G. The problem is to compute, modulo 2, the number of subgraphs of G that are isomorphic to H. The goal of this research is to determine for which classes $\mathcal{H}$ the problem $\oplus\textsc{Sub}(\mathcal{H})$ is fixed-parameter tractable (FPT), i.e., solvable in time $f(|H|) \cdot |G|^{O(1)}$. Curticapean, Dell, and Husfeldt (ESA 2021) conjectured that $\oplus\textsc{Sub}(\mathcal{H})$ is FPT if and only if the class of allowed patterns $\mathcal{H}$ is *matching splittable*, which means that for some fixed B, every $H \in \mathcal{H}$ can be turned into a matching (a graph in which every vertex has degree at most 1) by removing at most B vertices. Assuming the randomised Exponential Time Hypothesis, we prove their conjecture for (I) all hereditary pattern classes $\mathcal{H}$, and (II) all tree pattern classes, i.e., all classes $\mathcal{H}$ such that every $H \in \mathcal{H}$ is a tree. We also establish almost tight fine-grained upper and lower bounds for the case of hereditary patterns (I).

Keywords Modular counting · Parameterised complexity · Fine-grained complexity · Subgraph counting

Mathematics Subject Classification Theory of computation · Problems, reductions and completeness Mathematics of computing · Discrete mathematics

1 Introduction

The last two decades have seen remarkable progress in the classification of subgraph counting problems: Given a small *pattern* graph H and a large *host* graph G, how

---

For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. All data is provided in full in the results section of this paper.

✉ Marc Roth
marc.roth.cs@gmail.com

1 Department of Computer Science, University of Oxford, Oxford, UK

often does H occur as a subgraph if G? Since it was discovered that subgraph counts from small patterns reveal global properties of complex networks [26, 27], subgraph counting has also found several applications in fields such as biology [2, 33] genetics [35], phylogeny [25], and data mining [36]. Moreover, the theoretical study of subgraph counting and related problems has led to many deep structural insights, establishing both new algorithmic techniques and tight lower bounds under the lenses of fine-grained and parameterised complexity theory [4, 6, 10, 13, 14, 16, 19].

Without any additional restrictions, the subgraph counting problem is infeasible. The complexity class #W[1] is the parameterised complexity class analogous to NP (see Sect. 2 for more detail). Under standard assumptions, problems that are #W[1]-hard are not *fixed-parameter tractable* (FPT). The canonical complete problem for #W[1], the problem of counting k-cliques, corresponds to the special case of the subgraph counting problem where H is a clique of size k. This problem cannot be solved in time $f(k) \cdot n^{o(k)}$ for any function f unless the Exponential Time Hypothesis (ETH) fails [8, 9]. Due to this hardness result, the research focus in this area shifted to the question: Under which restrictions on the patterns H and the hosts G is algorithmic progress possible? More precisely, under which restrictions can the problem be solved in time $f(|H|) \cdot |G|^{O(1)}$, for some computable function f? Instances that can be solved within such a run time bound are called *fixed-parameter tractable* (FPT); allowing a potential super-polynomial overhead in the size of the pattern $|H|$ formalises the assumption that H is assumed to be (significantly) smaller than G.

If only the patterns are restricted, then the situation is fully understood. Formally, given a class $\mathcal{H}$ of patterns, the problem #SUB($\mathcal{H}$) asks, given as input a graph $H \in \mathcal{H}$ and an arbitrary graph G, to compute the number of subgraphs of G that are isomorphic to H. Following initial work by Flum and Grohe [19] and by Curticapean [11], Curticapean and Marx [14] proved that, under standard assumptions, #SUB($\mathcal{H}$) is FPT if and only if $\mathcal{H}$ has bounded matching number, that is, if there is a positive integer B such that the size of any matching in any graph in $\mathcal{H}$ is at most B. They also proved that all FPT cases are polynomial-time solvable.

In stark contrast, almost nothing is known for the decision version SUB($\mathcal{H}$). Here, the task is to correctly decide whether there is a copy of $H \in \mathcal{H}$ in G, rather than to count the copies. It is known that SUB($\mathcal{H}$) is FPT whenever $\mathcal{H}$ has bounded treewidth (see e.g. [20, Chapter 13]), and it is conjectured that those are all FPT cases. However, resolving this conjecture belongs to the "most infamous" open problems in parameterised complexity theory [18, Chapter 33.1].

1.1 Counting Modulo 2

To interpolate between the fully understood realm of (exact) counting and the barely understood realm of decision, Curticapean, Dell and Husfeldt proposed the study of counting subgraphs, modulo 2 [12]. Formally, they introduced the problem $\oplus$SUB($\mathcal{H}$), which expects as input a graph $H \in \mathcal{H}$ and an arbitrary graph G, and the goal is to compute *modulo 2* the number of subgraphs of G isomorphic to H.

The study of counting modulo 2 received significant attention from the viewpoint of classical, structural, and fine-grained complexity theory. For example, one way to state

Toda's Theorem [34] is PH $\subseteq$ $P^{\oplus P}$, implying that counting satisfying assignments of a CNF, modulo 2, is at least as hard as the polynomial hierarchy. Another example is the quest to classify the complexity of counting modulo 2 the homomorphisms to a fixed graph, which was very recently resolved by Bulatov and Kazeminia [7]. There has also been work by Abboud, Feller, and Weimann [1] on the fine-grained complexity of counting modulo 2 the number of triangles in a graph that satisfy certain weight constraints.

In their work [12], Curticapean, Dell and Husfeldt proved that the problem of counting k-matchings modulo 2, that is, the problem $\oplus\text{SUB}(\mathcal{H})$ where $\mathcal{H}$ is the class of all 1-regular graphs, is fixed-parameter tractable, where the parameter k is $|H|$. Since the exact counting version of this problem is #W[1]-hard [11], their result provides an example where counting modulo 2 is *strictly* easier than exact counting (subject to complexity assumptions). The complexity class $\oplus$W[1] can be defined via the complete problem of counting k-cliques modulo 2. Crucially, $\oplus$W[1]-hard problems are not fixed-parameter tractable, unless the randomised ETH (rETH) fails. Curticapean et al. [12] proved that counting k-paths modulo 2 is $\oplus$W[1]-hard. Since *finding* a k-path in a graph G is fixed-parameter tractable via colour-coding [3], this hardness result provides an example where counting modulo 2 is *strictly* harder than decision (subject to complexity assumptions). Combining those observations, it appears that counting subgraphs modulo 2 may lie strictly in between the complexity of decision and the complexity of exact counting.

A *matching* is a graph whose maximum degree is at most 1. The *matching-split number* of a graph H is the minimum size of a set $S \subseteq V(H)$ such that $H \setminus S$ is a matching. A class of graphs $\mathcal{H}$ is called *matching splittable* if there is a positive integer B such that the matching-split number of any $H \in \mathcal{H}$ is at most B. For example, the class of all matchings is matching splittable while the class of all cycles is not. Curticapean, Dell and Husfeldt extended their FTP algorithm for counting k-matchings modulo 2 to obtain an FPT algorithm for $\oplus\text{SUB}(\mathcal{H})$ for any matching-splittable class $\mathcal{H}$. On this basis, they then made the following conjecture.

Conjecture 1 ([12]) *$\oplus\text{SUB}(\mathcal{H})$ is FPT if and only if $\mathcal{H}$ is matching splittable.*

A class $\mathcal{H}$ of graphs is called *hereditary* if it is closed under vertex removal. Intriguingly, if Conjecture 1 is true, then the FPT criterion for counting subgraphs modulo 2 ($\oplus\text{SUB}(\mathcal{H})$) would coincide with the polynomial-time criterion for finding subgraphs ($\text{SUB}(\mathcal{H})$) for hereditary pattern classes $\mathcal{H}$ as established by Jansen and Marx.

Theorem 2 ([24]) *Let $\mathcal{H}$ be a hereditary class of graphs and assume* P $\neq$ NP. *Then* $\text{SUB}(\mathcal{H})$ *is solvable in polynomial time if and only if $\mathcal{H}$ is matching splittable.*

Jansen and Marx also conjecture that the condition of $\mathcal{H}$ being hereditary can be removed.

Conjecture 3 ([24]) $\text{SUB}(\mathcal{H})$ *is solvable in polynomial time if and only if $\mathcal{H}$ is matching splittable.*

Conjectures 1 and 3 have the remarkable consequence that $\oplus\text{SUB}(\mathcal{H})$ is FPT if and only if $\text{SUB}(\mathcal{H})$ is solvable in polynomial time. In the current work we establish this consequence for all hereditary pattern classes.

1.2 Our Contributions

We resolve Conjecture 1 for all hereditary classes $\mathcal{H}$, as well as for every class $\mathcal{H}$ consisting only of trees; note that the upper bounds were shown in [12] and that the lower bounds are the novel part.

Theorem 4 *Let $\mathcal{H}$ be a hereditary class of graphs. If $\mathcal{H}$ is matching splittable, then $\oplus\textsc{Sub}(\mathcal{H})$ is fixed-parameter tractable. Otherwise, the problem is $\oplus\mathsf{W}[1]$-complete and, assuming rETH, cannot be solved in time $f(|H|) \cdot |G|^{o(|V(H)|/\log|V(H)|)}$ for any function f.*

Theorem 5 *theoremmain Let $\mathcal{T}$ be a recursively enumerable class of trees. If $\mathcal{T}$ is matching splittable, then $\oplus\textsc{Sub}(\mathcal{T})$ is fixed-parameter tractable. Otherwise $\oplus\textsc{Sub}(\mathcal{T})$ is $\oplus\mathsf{W}[1]$-complete.*

The requirement that the class of trees $\mathcal{T}$ needs to be recursively enumerable is a standard technicality - the reason for it is that the function f in the running time in the standard definition of an FPT algorithm is required to be computable. It turns out that having $\mathcal{T}$ recursively enumerable is enough for this.

In order to prove our classifications, we adapt the by-now-standard technique for analysing subgraph counting problems established by Curticapean, Dell and Marx [13]. Let $\#\mathsf{Sub}(H \to G)$ denote the number of subgraphs of a graph G that are isomorphic to a graph H and let $\#\mathsf{Hom}(F \to G)$ denotes the number of homomorphisms (edge-preserving mappings) from a graph F to a graph G. Given a graph H, there is a function a_H from graphs to rationals with finite support such that the following holds for any graph G:

$$\#\mathsf{Sub}(H \to G) = \sum_F a_H(F) \cdot \#\mathsf{Hom}(F \to G), \tag{1}$$

where the sum is over all (isomorphism types of) graphs. Since a_H has finite support, $a_H(F) = 0$ for all but finitely-many graphs F. Thus, Eq. (1) allows us to express the solution to the *exact* counting problem as a finite linear combination of homomorphism counts. In a nutshell, the framework of [13] states that computing the function $G \mapsto \#\mathsf{Sub}(H \to G)$ is hard to compute if and only if there is a graph F of high treewidth with $a_H(F) \neq 0$. This translates the complexity of (exact) subgraph counting to the purely combinatorial problem of understanding the coefficients a_H. One might hope that this strategy transfers to counting modulo 2 as well. Unfortunately, this is not possible as Eq. (1) might not be well-defined if arithmetic is done modulo 2. The reason for this is the fact that the coefficients $a_H(F)$ are of the form $\mu(F, H) \times |\mathsf{Aut}(H)|^{-1}$, where $\mu(F, H)$ is an integer, and $\mathsf{Aut}(H)$ is the automorphism group of the graph H [13]. Thus there is, a priori, no hope to extend the framework to counting modulo 2 for pattern graphs with an even number of automorphisms. In fact, according to Curticapean, Dell and Husfeldt [12], the absence of a comparable framework for counting modulo 2 is one of the main challenges for establishing the hardness part of Conjecture 1, and it is the main reason why the reductions in [12] use more classical, gadget-based reductions.

In this work, we solve the problem of patterns with an even number of automorphisms by considering a colourful intermediate problem. More concretely, we will equip each edge of the pattern H with a distinct colour and show that it will be sufficient to consider only automorphisms that preserve the colours. If H has no isolated vertices, then this is only the trivial automorphism. Formally, the coloured approach will be based on the notion of so-called *fractured* graphs introduced by Peyerimhoff et al. [30].

1.3 Organisation of the Paper

We start by introducing some basic terminology in Sect. 2. The formal definitions of our graph colourings, as well as colour-preserving homomorphisms and embeddings can be found in Sect. 2.1, and the majority of the paper will consider the coloured setting as it allows us to get rid of automorphism groups of even size. This is formalised in Sect. 2.2 using the framework of *fractured graphs* originally introduced in [30]. An introduction to parameterised and fine-grained complexity theory, including the definition of our computational problems and the statement of the randomised Exponential Time Hypothesis, can be found in Sect. 2.3; moreover, this section contains a self-contained and formal exposition of the complexity monotonicty principle for coloured graphs in the modular setting, stating intuitively that the computation, modulo 2, of a finite linear combination of homomorphism counts between coloured graphs is precisely as hard as computing, modulo 2, the hardest term with an odd coefficient. Additionally, Sect. 2.3 contains the formal statement of the reduction from the coloured setting to the uncoloured setting via the principle of inclusion and exclusion. Note that this reduction is necessary for obtaining our main results (Theorems 4 and 5), which classify the complexity of the uncoloured problem $\oplus\textsc{Sub}(\mathcal{H})$.

Having completed the set-up, we continue in Sect. 3 with the treatment of $\oplus\textsc{Sub}(\mathcal{H})$ for hereditary $\mathcal{H}$, i.e., with the proof of Theorem 4. We note that, on a technical level, understanding the hereditary case is much easier than the case of trees. However, almost all of the key techniques and ideas that become necessary to classify the case of trees are already used in Sect. 3, although in a much simpler way. For this reason, we consider Sect. 3 also as a warm-up for getting used to the framework of fractured graphs. Concretely, we can outline our treatment of hereditary classes as follows: Using a result of Jansen and Marx [24], each hereditary class of graphs $\mathcal{H}$ is either matching splittable, or it fully contains one of the following four subclasses: (I) The class of all cliques, (II) the class of all bicliques, (III) the class of all triangle packings (disjoint unions of triangles), or (IV) the class of all P_2-packings (disjoint unions of paths with two edges). For proving the classification of $\oplus\textsc{Sub}(\mathcal{H})$ for hereditary $\mathcal{H}$ (Theorem 4), it thus suffices to show that each of the four cases (I) - (IV) is hard. Since the problems of *deciding* whether a graph contains a k-clique or whether a graph contains a k-by-k-biclique are already hard, the problem of counting their respective occurences modulo 2 (cases (I) and (II)) can easily shown to be hard using a variation of the Isolation Lemma due to Williams et al. [37]. The majority of Sect. 3 is thus dedicated to establishing hardness for triangle packings (III) in Sect. 3.1 and for P_2-packings (IV) in Sect. 3.2.

The classification of $\oplus\textsc{Sub}(\mathcal{T})$ for classes of trees $\mathcal{T}$ (Theorem 5) can be found in Sect. 4. In the first step, we establish a graph-theoretical classification of classes of trees that are not matching splittable. To this end, we first introduce three structural invariants of trees (the definitions are rather technical and can be found right at the beginning of Sect. 4): The *fork number*, the *star number*, and the C-*number*. We then show that each class $\mathcal{T}$ of trees is either matching splittable, or it satisfies at least one of the following properties:

(1) $\mathcal{T}$ has unbounded C-number,
(2) $\mathcal{T}$ has unbounded star number, or
(3) $\mathcal{T}$ has unbounded fork number.

The central steps of the proof of Theorem 5 are then hardness proofs for the previous three cases: Case (1) is treated in Sect. 4.1, Case (2) is treated in Sect. 4.2, and Case (3) is treated in Sect. 4.3. Finally, we collect the intractabilty results for all cases in Sect. 4.4 to prove Theorem 5.

2 Preliminaries

Let $f : A_1 \times A_2 \to B$ be a function. For each $a_1 \in A_1$ we write $f(a_1, \star) : A_2 \to B$ for the function that maps $a_2 \in A_2$ to $f(a_1, a_2)$.

Graphs in this work are undirected and without self loops. A *homomorphism* from a graph H to a graph G is a mapping φ from the vertices $V(H)$ of H to the vertices $V(G)$ of G such that for each edge $e = \{u, v\} \in E(H)$ of H, the image $\varphi(e) = \{\varphi(u), \varphi(v)\}$ is an edge of G. A homomorphism is called an *embedding* if it is injective. We write $\mathsf{Hom}(H \to G)$ and $\mathsf{Emb}(H \to G)$ for the sets of homomorphisms and embeddings, respectively, from H to G. An embedding $\varphi \in \mathsf{Emb}(H \to G)$ is called an *isomorphism* if it is bijective and $\{u, v\} \in E(H) \Leftrightarrow \{\varphi(u), \varphi(v)\} \in E(G)$. We say that H and G are *isomorphic*, denoted by $H \cong G$, if an isomorphism from H to G exists. A *graph invariant* ι is a function from graphs to rationals such that $\iota(H) = \iota(G)$ for each pair of isomorphic graphs H and G.

A *subgraph* of G is a graph G' with $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$. We write $\mathsf{Sub}(H \to G)$ for the set of all subgraphs of G that are isomorphic to H. Given a subset of vertices $S \subseteq V(G)$ of a graph G, we write $G[S]$ for the graph induced by S, that is, $G[S]$ has vertices S and edges $\{\{u, v\} \subseteq S \mid \{u, v\} \in E(G)\}$.

We denote by $\mathsf{tw}(G)$ the *treewidth* of the graph G. Since we will rely on treewidth purely in a black-box manner, we omit the technical definition and refer the reader to [15, Chapter 7].

Given any graph invariant ι (such as treewidth) and a class of graphs $\mathcal{G}$, we say that ι is *bounded* in $\mathcal{G}$ if there is a non-negative integer B such that, for all $G \in \mathcal{G}$, $\iota(G) \leq B$. Otherwise we say that ι is *unbounded* in $\mathcal{G}$.

Given a graph $H = (V, E)$, a *splitting set* of H is a subset of vertices S such that every vertex in $H[V \setminus S]$ has degree at most 1. The *matching-split number* of H is the minimum size of a splitting set of H. A class of graphs $\mathcal{H}$ is called *matching splittable* if the matching-split number of $\mathcal{H}$ is bounded.

2.1 Colour-Preserving Homomorphisms and Embeddings

A homomorphism c from a graph G to a graph Q is sometimes called a "Q-colouring" of G. A *Q-coloured* graph is a pair consisting of a graph G and a homomorphism c from G to Q. Note that the identity function id_Q on $V(Q)$ is a Q-colouring of Q. If a homomorphism c from G to Q is vertex surjective, then we call (G, c) a *surjectively Q-coloured* graph.

Definition 6 (c_{E}) A Q-colouring c of a graph G induces a (not necessarily proper) edge-colouring $c_{\mathrm{E}}\colon E(G) \to E(Q)$ given by $c_{\mathrm{E}}(\{u, v\}) = \{c(u), c(v)\}$.

Notation: Given a Q-coloured graph (G, c) and a vertex $u \in V(Q)$, we will use the capitalised letter U to denote the subset of vertices of G that are coloured by c with u, that is, $U := c^{-1}(u) \subseteq V(G)$.

Given two Q-coloured graphs (H, c_H) and (G, c_G), we call a homomorphism φ from H to G *colour-preserving* if for each $v \in V(H)$ we have $c_G(\varphi(v)) = c_H(v)$. We note the special case in which $Q = H$ and c_H is the identity id_Q; then the condition simplifies to $c_G(\varphi(v)) = v$. A colour-preserving embedding of (H, c_H) in (G, c_G) is a vertex injective colour-preserving homomorphism from (H, c_H) to (G, c_G). We write $\mathsf{Hom}((H, c_H) \to (G, c_G))$ and $\mathsf{Emb}((H, c_H) \to (G, c_G))$ for the sets of all colour-preserving homomorphisms and embeddings, respectively, from (H, c_H) to (G, c_G).

Let k be a positive integer, let H be a graph with k edges, and let (G, γ) be a pair consisting of a graph G and a function that maps each edge of G to one of k distinct colours. We refer to γ as a "k-edge colouring" of G. For example, in most of our applications we will fix a graph Q with k edges and a Q-colouring c of G and we will take γ to be the edge-colouring c_E from Definition 6. We write $\mathsf{ColSub}(H \to (G, \gamma))$ for the set of all subgraphs of G that are isomorphic to H and that contain each of the k edge colours precisely once.

2.2 Fractures and Fractured Graphs

In this work, we will crucially rely on and extend the framework of *fractured* graphs as introduced in [30].

Definition 7 (*Fractures*) Let Q be a graph. For each vertex v of Q, let $E_Q(v)$ be the set of edges of Q that are incident to v. A *fracture* of Q is a tuple $\rho = (\rho_v)_{v \in V(Q)}$, where for each vertex v of Q, ρ_v is a partition of $E_Q(v)$.

Note that a fracture describes how to split (or how to *fracture*) each vertex of a given graph: for each vertex v, create a vertex v^B for each block B in the partition ρ_v; edges originally incident to v are made incident to v^B if and only if they are contained in B. We call the resulting graph the *fractured graph* $H \,\sharp\, \rho$; a formal definition is given in Definition 8, a visualisation is given in Fig. 1.

Fig. 1 Illustration of the construction of a fractured graph from [30]. The left picture shows a vertex v of a graph Q with incident edges $E_Q(v) = \{$●, ●, ●, ●, ●, ●$\}$. The right picture shows the splitting of v in the construction of the fractured graph $Q \sharp \sigma$ for a fracture σ satisfying that the partition σ_v contains two blocks $B_1 = \{$●, ●, ●$\}$, and $B_2 = \{$●, ●, ●$\}$ (Color figure online)

Definition 8 (*Fractured Graph* $Q \sharp \rho$) Given a graph Q, we consider the matching M_Q containing one edge for each edge of Q; formally,

$$V(M_Q) := \bigcup_{e=\{u,v\} \in E(Q)} \{u_e, v_e\} \quad \text{and} \quad E(M_Q) := \{\{u_e, v_e\} \mid e = \{u, v\} \in E(Q)\}.$$

For a fracture ρ of Q, we define the graph $Q \sharp \rho$ to be the quotient graph of M_Q under the equivalence relation on $V(M_Q)$ which identifies two vertices v_e, w_f of M_Q if and only if $v = w$ and e, f are in the same block B of the partition ρ_v of $E_Q(v)$. We write v^B for the vertex of $Q \sharp \rho$ given by the equivalence class of the vertices v_e (for which $e \in B$) of M_Q.

Definition 9 (*Canonical Q-colouring* c_ρ) Let Q be a graph and let ρ be a fracture of Q. The *canonical Q-colouring* of the fractured graph $Q \sharp \rho$ maps v^B to v for each $v \in V(Q)$ and block $B \in \rho_v$, and is denoted by c_ρ.

Observe that c_ρ is the identity in $V(Q)$ if ρ is the coarsest fracture (that is, each partition ρ_v only contains one block, in which case $Q \sharp \rho = Q$).

2.3 Parameterised and Fine-Grained Computation

A *parameterised computational problem* is a pair consisting of a function $P : \Sigma^* \to \{0, 1\}$ and a computable parameterisation $\kappa : \Sigma^* \to \mathbb{N}$. A *fixed-parameter tractable* (FPT) *algorithm* for (P, κ) is an algorithm that computes P and runs, on input $x \in \Sigma^*$, in time $f(\kappa(x)) \cdot |x|^{O(1)}$ for some computable function f. We call (P, κ) *fixed-parameter tractable* (FPT) if an FPT algorithm for (P, κ) exists.

A parameterised Turing-reduction from (P, κ) to (P', κ') is an FPT algorithm for (P, κ) that is equipped with oracle access to P' and for which there is a computable function g such that, on input x, each oracle query y satisfies $\kappa'(y) \leq g(\kappa(x))$. We write $(P, \kappa) \leq^{\text{fpt}}_{\text{T}} (P', \kappa')$ if a parameterised Turing-reduction from (P, κ) to (P', κ') exists. This guarantees that fixed-parameter tractability of (P', κ') implies fixed-parameter tractability of (P, κ). For a more comprehensive introduction, we refer the reader the standard textbooks [15] and [20].

2.3.1 Counting Modulo 2 and the rETH

The lower bounds in this work will rely on the hardness of the parameterised complexity class $\oplus\mathsf{W}[1]$, which can be considered a parameterised equivalent of $\oplus\mathsf{P}$. Following [12], we define $\oplus\mathsf{W}[1]$ via the complete problem $\oplus$CLIQUE: Given as input a graph G and a positive integer k, the goal is to compute the number of k-cliques in G modulo 2, i.e., to compute $\oplus\mathsf{Sub}(K_k \to G)$. The problem is parameterised by k. A parameterised problem (P, κ) is called $\oplus\mathsf{W}[1]$-*hard* if $\oplus\text{CLIQUE} \leq^{\text{fpt}}_{\text{T}} (P, \kappa)$, and it is called $\oplus\mathsf{W}[1]$-*complete* if, additionally, $(P, \kappa) \leq^{\text{fpt}}_{\text{T}} \oplus\text{CLIQUE}$.

Modifications of the classical Isolation Lemma (see e.g. [5] and [37]) yield a *randomised* parameterised Turing reduction from finding a k-clique to computing the parity of the number of k-cliques. In combination with existing fine-grained lower bounds for finding a k-clique [8, 9], it can then be shown that $\oplus$CLIQUE cannot be solved in time $f(k) \cdot |G|^{o(k)}$ for any function f, unless the randomised Exponential Time Hypothesis fails:

Definition 10 (*rETH*, [23]) The *randomised Exponential Time Hypothesis* (rETH) asserts that 3-SAT cannot be solved by a randomised algorithm in time $\exp o(n)$, where n is the number of variables of the input formula.

As an immediate consequence, the rETH implies that $\oplus\mathsf{W}[1]$-hard problems are not fixed-parameter tractable.

For the lower bounds in this work, we won't reduce from $\oplus$CLIQUE directly, but instead from the following, more general problem:

Definition 11 ($\oplus$CP- HOM) Let $\mathcal{H}$ be a class of graphs. The problem $\oplus$CP- HOM$(\mathcal{H})$ has as input a graph $H \in \mathcal{H}$ and a surjectively H-coloured graph (G, c). The goal is to compute $\oplus\mathsf{Hom}((H, \mathsf{id}_H) \to (G, c))$. The problem is parameterised by $|H|$.

The following lower bound was proved independently in [28, 30] and [12].

Theorem 12 *Let $\mathcal{H}$ be a recursively enumerable class of graphs. If the treewidth of $\mathcal{H}$ is unbounded then $\oplus$CP- HOM$(\mathcal{H})$ is $\oplus\mathsf{W}[1]$-hard and, assuming the rETH, it cannot be solved in time $f(|H|) \cdot |G|^{o(\mathsf{tw}(H)/\log \mathsf{tw}(H))}$ for any function f.*

Next is the central problem in this work.

Definition 13 ($\oplus$SUB) Let $\mathcal{H}$ be a class of graphs. The problem $\oplus$SUB$(\mathcal{H})$ has as input a graph $H \in \mathcal{H}$ and a graph G. The goal is to compute $\oplus\mathsf{Sub}(H \to G)$. The problem is parameterised by $|H|$.

For example, writing $\mathcal{K}$ for the set of all complete graphs, the problem $\oplus$SUB$(\mathcal{K})$ is equivalent to $\oplus$CLIQUE.

2.3.2 Complexity Monotonicity and Inclusion–Exclusion

Throughout this work, we will rely on two important tools introduced in [30]. For the sake of being self-contained, we encapsulate them below in individual lemmas.

The first tool is an adaptation of the so-called Complexity Monotonicity principle to the realm of fractured graphs and modular counting (see [30, Sections 4.1 and

6.3] for a detailed treatment and for a proof). Intuitively, the subsequent lemma states that evaluating, modulo 2, a linear combination of colour-prescribed homomorphism counts from fractured graphs, is as hard as evaluating its hardest term with an odd coefficient.

Lemma 14 ([30]) *There is a deterministic algorithm* $\mathbb{A}$ *and a computable function* f *such that the following conditions are satisfied:*

1. $\mathbb{A}$ *expects as input a graph* Q *and a* Q*-coloured graph* (G, c).
2. $\mathbb{A}$ *is equipped with oracle access to a function*

$$(G', c') \mapsto \sum_{\rho} a(\rho) \cdot \oplus\mathsf{Hom}((Q \sharp \rho, c_\rho) \to (G', c')) \mod 2,$$

   where the sum is over all fractures of Q *and* a *is a function from fractures of* Q *to integers.*
3. *Each oracle query* (G', c') *is of size at most* $f(|Q|) \cdot |G|$.
4. $\mathbb{A}$ *computes* $\oplus\mathsf{Hom}((Q \sharp \rho, c_\rho) \to (G, c))$ *for each fracture* ρ *with* $a(\rho) \neq 0 \mod 2$.
5. *The running time of* $\mathbb{A}$ *is bounded by* $f(|Q|) \cdot |G|^{O(1)}$.

The second tool is a standard application of the inclusion–exclusion principle (see e.g. [30, Sections 4.2 and 6.3]). It will be used in the final steps of our reductions to remove the colourings.

Lemma 15 ([30]) *There is a deterministic algorithm* $\mathbb{A}$ *that satisfies the following conditions:*

1. $\mathbb{A}$ *expects as input a graph* H *with* k *edges, a graph* G *and a* k*-edge colouring* γ *of* G.
2. $\mathbb{A}$ *is equipped with oracle access to the function* $\oplus\mathsf{Sub}(H \to \star)$*, and each oracle query* G' *satisfies* $|G'| \le |G|$.
3. $\mathbb{A}$ *computes* $\oplus\mathsf{ColSub}(H \to (G, \gamma))$.
4. *The running time of* $\mathbb{A}$ *is bounded by* $2^{|H|} \cdot |G|^{O(1)}$.

3 Classification for Hereditary Graph Classes

In this section, we will completely classify the complexity of $\oplus\textsc{Sub}(\mathcal{H})$ for hereditary classes. Let us start by restating the classification theorem.

Theorem 4 Let $\mathcal{H}$ be a hereditary class of graphs. If $\mathcal{H}$ is matching splittable, then $\oplus\textsc{Sub}(\mathcal{H})$ is fixed-parameter tractable. Otherwise, the problem is $\oplus$W[1]-complete and, assuming rETH, cannot be solved in time $f(|H|) \cdot |G|^{o(|V(H)|/\log |V(H)|)}$ for any function f.

The proof of Theorem 4 is split in four cases, which stem from a structural property of non matching splittable hereditary graph classes $\mathcal{H}$ due to Jansen and Marx [24]. For the statement, we need to consider the following classes:

- $\mathcal{F}_\omega$ is the class of all complete graphs.
- $\mathcal{F}_\beta$ is the class of all complete bipartite graphs.
- $\mathcal{F}_{P_2}$ is the class of all P_2-packings, that is, disjoint unions of paths with two edges.[1]
- $\mathcal{F}_{K_3}$ is the class of all triangle packings, that is, disjoint unions of the complete graph of size 3.

Theorem 16 (Theorem 3.5 in [24]) *Let $\mathcal{H}$ be a hereditary class of graphs. If $\mathcal{H}$ is not matching splittable then at least one of the following are true: (1.) $\mathcal{F}_\omega \subseteq \mathcal{H}$, (2.) $\mathcal{F}_\beta \subseteq \mathcal{H}$, (3.) $\mathcal{F}_{P_2} \subseteq \mathcal{H}$, or (4.) $\mathcal{F}_{K_3} \subseteq \mathcal{H}$.*

Thus, it suffices to consider cases 1. - 4. to prove Theorem 4. We start with the easy cases of cliques and bicliques; they follow implicitly from previous works [12, 17, 28] and we only include a proof for completeness. Note that a tight bound under rETH is known for those cases:

Lemma 17 *Let $\mathcal{H}$ be a hereditary class of graphs. If $\mathcal{F}_\omega \subseteq \mathcal{H}$ or $\mathcal{F}_\beta \subseteq \mathcal{H}$ then $\oplus$Sub($\mathcal{H}$) is $\oplus$W[1]-hard and, assuming rETH, cannot be solved in time $f(|H|) \cdot |G|^{o(|V(H)|)}$ for any function f.*

Proof If $\mathcal{F}_\omega \subseteq \mathcal{H}$ then $\oplus$W[1]-hardness follows immediately from the fact that $\oplus$Clique is the canonical $\oplus$W[1]-complete problem [12]. For the rETH lower bound, we can reduce from the problem of *deciding* the existence of a k-clique via a (randomised) reduction using a version of the Isolation Lemma due to Williams et al. [37, Lemma 2.1]. This reduction does not increase k or the size of the host graph and is thus tight with respect to the well-known lower bound for the clique problem due to Chen et al. [8, 9]: Deciding the existence of a k-clique in an n-vertex graph cannot be done in time $f(k) \cdot n^{o(k)}$ for any function f, unless ETH fails. Our lower bound under rETH follows since the reduction is randomised.

If $\mathcal{F}_\beta \subseteq \mathcal{H}$, then the claim holds by [17, Theorem 5], which established the problem of counting, modulo 2, the induced copies of a k-by-k-biclique in an n-vertex bipartite graph to be $\oplus$W[1]-hard and not solvable in time $f(k) \cdot n^{o(k)}$ for any function f, unless rETH fails. Since a copy of a biclique (with at least one edge) in a bipartite graph must always be induced, the claim follows. This concludes the proof of Lemma 17. □

The more interesting cases are $\mathcal{F}_{P_2} \subseteq \mathcal{H}$ and $\mathcal{F}_{K_3} \subseteq \mathcal{H}$. One reason for this is that, in contrast to cliques and bicliques, the decision version of those instances are fixed-parameter tractable. Hence a reduction from the decision version via e.g. an isolation lemma does not help. In other words, establishing hardness for those cases requires us to rely on the full power of counting modulo 2. More precisely, we will rely on the framework of fractures graphs (see Sect. 2). Both cases can be considered simpler applications of the machinery used in the later sections, so we will present all steps in great detail. While this might seem unnecessary given the simplicity of the constructions, we hope that it enables the reader to make themselves familiar with the general reduction strategies which will be used throughout the later sections of this work.

[1] To avoid confusion, we remark that [24] uses P_3 to denote the path of two edges (and three vertices). In the current work, it will be more convenient to use the number of edges of a path as index.

3.1 Triangle Packings

The goal of this subsection is to establish hardness of $\oplus\textsc{Sub}(\mathcal{F}_{K_3})$. To this end, let Δ be an infinite computable class of cubic bipartite expander graphs, and let $\mathcal{Q} = \{L(H) \mid H \in \Delta\}$ where $L(H)$ is constructed as follows: Each $v \in V(H)$ becomes a triangle with vertices v_x, v_y, and v_z corresponding to the three neighbours x, y, and z of v. Finally, for every edge $\{u, v\} \in E(H)$ we identify v_u and u_v. In fact, $L(H)$ is just the *line graph* of H: Every edge of H becomes a vertex in $L(H)$, and two vertices of $L(H)$ are made adjacent if and only if the corresponding edges in H are incident. Since all $H \in \Delta$ are bipartite (and thus triangle-free), we can easily observe the following.[2]

Observation 18 *The mapping $v \mapsto (v_x, v_y, v_z)$ is a bijection from vertices of H to triangles in $L(H)$.*

We also consider the fracture of $L(H)$ that splits $L(H)$ back into $|V(H)|$ triangles; consider Fig. 2 for an illustration.

Definition 19 ($\tau(H)$) Let $H \in \Delta$ and recall that each vertex w of $L(H)$ is obtained by identifying v_u and u_v for some edge $\{u, v\} \in E(H)$. Moreover, w has four incident edges e_x, e_y, e_a, e_b, to v_x, v_y, u_a, u_b, respectively, where x, y, u are the neighbours of v in H and v, a, b are the neighbours of u in H. We define $\tau(H)_w := \{\{e_x, e_y\}, \{e_a, e_b\}\}$, and we proceed similar for all vertices of $L(H)$.

Next, we use that $\mathsf{tw}(L(H)) = \Omega(\mathsf{tw}(H))$ (see e.g. [22]). Moreover, $\mathsf{tw}(L(H)) \leq |V(L(H))|$ since the treewidth of a graph is always bounded by the number of its vertices. Additionally, $|V(L(H))| = |E(H)|$ by construction. Since the graphs in Δ are cubic, we further have that $|E(H)| = \Theta(|V(H)|)$ for $H \in \Delta$. We combine those bounds with the fact that expander graphs have treewidth linear in the number of vertices (see e.g. [21]); therefore Δ and thus $\mathcal{Q}$ have unbounded treewidth. Putting these facts together, we obtain the following.

Fact 20 *$\mathcal{Q}$ has unbounded treewidth and* $\mathsf{tw}(L(H)) = \Theta(|V(L(H))|) = \Theta(|V(H)|)$ *for $H \in \Delta$.*

We are now able to establish hardness of $\oplus\textsc{Sub}(\mathcal{F}_{K_3})$. The proof will heavily rely on the transformation from edge-coloured subgraphs to homomorphisms established in [30].

Lemma 21 *The problem $\oplus\textsc{Sub}(\mathcal{F}_{K_3})$ is $\oplus\mathsf{W}[1]$-hard. Furthermore, on input kK_3 and G, the problem cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ for any function f, unless rETH fails.*

Proof We reduce from $\oplus\textsc{cp-Hom}(\mathcal{Q})$, which, by Fact 20 and Theorem 12, is $\oplus\mathsf{W}[1]$-hard and for $L(H) \in \mathcal{Q}$, it cannot be solved in time $f(|L(H)|) \cdot |G|^{o(|V(L(H))|/\log|V(L(H))|)}$, unless rETH fails.

[2] Observation 18 is also an immediate consequence of Whitney's Isomorphism Theorem implying that a triangle of a line graph corresponds to either a claw or to a triangle in its primal graph.

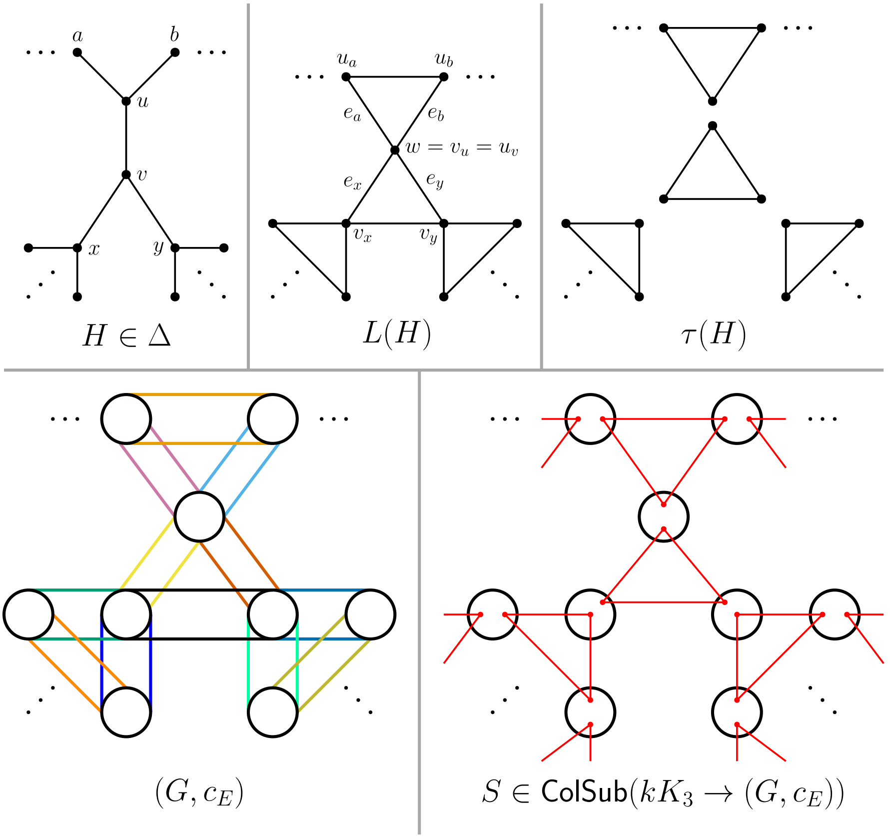

Fig. 2 (*Top*:) A cubic bipartite graph $H \in \Delta$, its line graph $L(H)$, and the fractured graph induced by $\tau(H)$. (*Below*:) An $L(H)$-coloured graph (G, c); emphasised in distinct colours is the edge-colouring c_E of G induced by the mapping $\{u, v\} \mapsto \{c(u), c(v)\}$. Additionally we depict an element $S \in \mathsf{ColSub}(kK_3 \to (G, c_E))$, that is, a subgraph of G isomorphic to kK_3 that contains each edge colour of G precisely once (Color figure online)

Let L and (G, c) be an input instance to $\oplus$CP- HOM$(\mathcal{Q})$. Recall that Δ is computable — that is, there is an algorithm that takes a graph H and determines whether it is in Δ. Thus, there is an algorithm that takes input $L \in \mathcal{Q}$ and finds a graph $H \in \Delta$ with $L = L(H)$. The run time of this algorithm depends on $|L|$ but clearly not on (G, c). Let $k = |V(H)|$ and note that $|E(L(H))| = 3k$, since, by construction, each vertex v of H becomes a triangle of $L(H)$. We consider the graph G as a $3k$-edge-coloured graph, coloured by c_E. That is, each edge $e = \{x, y\}$ of G is assigned the colour $c_E(e) = \{c(x), c(y)\}$ which is an edge of L (see Fig. 2 for an illustration).

Now, for *any* L-coloured graph (G', c') recall that $\mathsf{ColSub}(kK_3 \to (G', c'_E))$ is the set of subgraphs of G' that are isomorphic to kK_3 and that include each edge colour (each edge of L) precisely once. We will see later that $\oplus\mathsf{ColSub}(kK_3 \to (G', c'_E))$ can be computed using our oracle for $\oplus$SUB$(\mathcal{F}_{K_3})$ using the principle of inclusion and exclusion.

It was shown in [30, Lemma 4.1] that there is a unique function a such that for every L-coloured graph (G', c') we have[3]

$$\#\mathsf{ColSub}(kK_3 \to (G', c'_E)) = \sum_{\rho} a(\rho) \cdot \mathsf{Hom}(L \sharp \rho \to (G', c')). \tag{2}$$

where the sum is over all fractures of L. Additionally, it was shown in [30, Corollary 4.3] that

$$a(\top) = \sum_{\rho \in \mathsf{F}(kK_3, L)} \prod_{w \in V(L)} (-1)^{|\rho_w|-1} \cdot (|\rho_w| - 1)!, \tag{3}$$

where $\top$ is the fracture in which each partition consists only of one block (that is, $L \sharp \top = L$), and $\mathsf{F}(kK_3, L)$ is the set of all fractures ρ of L such that $L \sharp \rho \cong kK_3$. However, note that, by Observation 18, there is only way to fracture L into k disjoint triangles, and this fracture is given by $\tau(H)$. Thus, (3) simplifies to

$$a(\top) = \prod_{w \in V(L)} (-1)^{|\tau(H)_w|-1} \cdot (|\tau(H)_w| - 1)!, \tag{4}$$

which is odd since each partition of $\tau(H)$ consists of precisely two blocks (so in fact the expression in (4) is $(-1)^{|V(L)|}$).

Note that the algorithm for $\oplus$CP-HOM($\mathcal{Q}$) is supposed to compute $\oplus\mathsf{Hom}((L, \mathsf{id}_L) \to (G, c))$ which is equal to $\oplus\mathsf{Hom}(L \sharp \top \to (G, c_\top))$. Since $a(\top)$ is odd, we can invoke Lemma 14 to recover this term by evaluating the entire linear combination (2), that is, by evaluating the function $\oplus\mathsf{ColSub}(kK_3 \to \star)$. More concretely, this means that we need to compute $\oplus\mathsf{ColSub}(kK_3 \to (G', c'_E))$ for some L-coloured graphs (G', c') of size at most $f(|L|) \cdot |G|$ for some computable function f (see 3. in Lemma 14). This can easily be done using Lemma 15 since we have oracle access to the function $\oplus\mathsf{Sub}(kK_3 \to \star)$. We emphasise that, by condition 2. of Lemma 15, each oracle query $\hat{G}$ satisfies $|\hat{G}| \leq |G'|$, where (G', c') is the L-coloured graph for which we wish to compute $\oplus\mathsf{ColSub}(kK_3 \to (G', c'_E))$. Since $|(G', c')| \leq f(|L|) \cdot |G|$, we obtain that $|\hat{G}| \leq f(|L|) \cdot |G|$ as well.

Since, by Fact 20, $k = \Theta(|kK_3|) = \Theta(|V(L)|) = \Theta(\mathsf{tw}(L))$, our reduction yields $\oplus$W[1]-hardness and transfers the conditional lower bound under rETH as desired. □

3.2 P_2-Packings

Next we establish hardness for the case of P_2-packings. The strategy will be similar in spirit to the construction for triangle packings; however, rather then identifying a unique fracture for which the technique applies, we will encounter an *odd* number of possible fractures in the current section.

[3] In the language of [30], Eq. (2) is obtained by choosing Φ as the property of being isomorphic to kK_3.

Let Δ be a computable infinite class of 4-regular expander graphs, and let $\mathcal{Q}$ be the class of all subdivisions of graphs in Δ, that is $\mathcal{Q} = \{H^2 \mid H \in \Delta\}$, where H^2 is obtained from H by subdividing each edge once.

We start by establishing an easy but convenient fact on the treewidth of the graphs in $\mathcal{Q}$.

Lemma 22 *$\mathcal{Q}$ has unbounded treewidth and* $\mathsf{tw}(H^2) = \Theta(|V(H)|)$ *for* $H \in \Delta$.

Proof As in Sect. 3.1, $\mathsf{tw}(H) = \Theta(|V(H)|)$ for $H \in \Delta$, since expanders have treewidth linear in the number of vertices. Since H is a minor of H^2, and since taking minors cannot increase treewidth (see [15, Exercise 7.7]), we thus have that $\mathsf{tw}(H^2) = \Omega(|V(H)|))$. Finally, we have $\mathsf{tw}(H^2) \leq |V(H^2)|$ since the treewidth is at most the number of vertices, and $|V(H^2)| = O(|V(H)|)$ since H is 4-regular. In combination, we obtain $\mathsf{tw}(H^2) = \Theta(|V(H)|)$ for $H \in \Delta$. Note that this also implies that $\mathcal{Q}$ has unbounded treewidth (as Δ is infinite). □

For what follows, given a subdivision H^2 of a graph H, it will be convenient to assume that $V(H^2) = V(H) \cup S_E$, where $S_E = \{s_e \mid e \in E(H)\}$ is the set of the subdivision vertices.

Definition 23 (*Odd Fractures*) Let $H \in \Delta$ and let τ be a fracture of H^2. We say that τ is *odd* if the following two conditions are satisfied:

1. For each $s \in S_E$ the partition τ_s consists of two singleton blocks.
2. For each $v \in V(H)$ the partition τ_v consists of two blocks of size 2.

Consider Fig. 3 for a depiction of an odd fracture.

The following two lemmas are crucial for our construction.

Lemma 24 *Let $H \in \Delta$. The number of odd fractures of H^2 is odd.*

Proof The first condition in Definition 23 leaves only one choice for subdivision vertices. Let us thus consider a vertex $v \in V(H) = V(H^2) \setminus S_E$. Since H is 4-regular, there are 4 incident edges to v. Now note that there are precisely 3 partitions of a 4-element set with two blocks of size 2. Thus the total number of odd fractures of H^2 is $3^{|V(H)|}$, which is odd. □

Lemma 25 *Let $H \in \Delta$, let $k = 2|V(H)|$ and let τ be a fracture of H^2 such that τ_v consists of at most 2 blocks for each $v \in V(H^2)$. Then $H^2 \sharp \tau \cong kP_2$ if and only if τ is odd.*

Proof First observe that $|E(H^2)| = 2|E(H)| = 4|V(H)| = 2k$. Thus the number of edges of $H^2 \sharp \tau$ is equal to $2k$ (for each fracture τ of H^2), which is also equal to the number of edges of kP_2.

Thus, $H^2 \sharp \tau$ is isomorphic to kP_2 if and only if each connected component of $H^2 \sharp \tau$ is a path of length 2. It follows immediately by Definition 23 that τ being odd implies that $H^2 \sharp \tau$ consists only of disjoint P_2. It thus remains to show the other direction.

Assume for contradiction that there is a subdivision vertex $s \in S_E$ of H^2 such that τ_s consists of only one block (recall that s has degree 2, thus τ_s either consists of

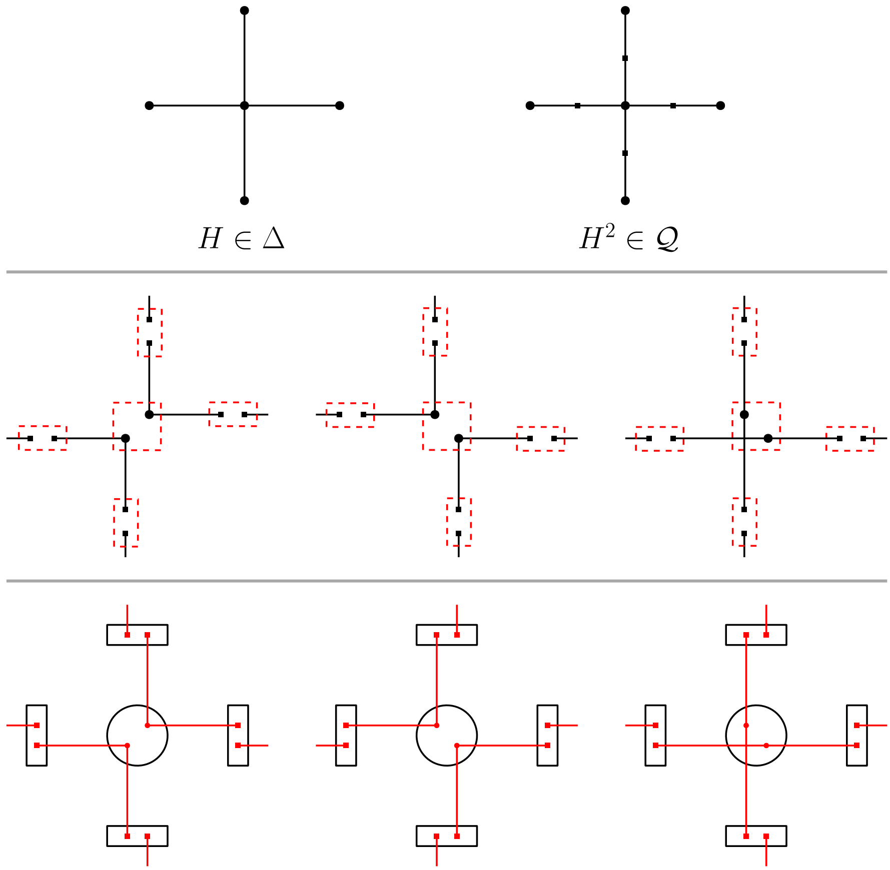

Fig. 3 (*Top*:) Subdividing a 4-regular expander in Δ depicted by the neighbourhood of an individual vertex. (*Centre*:) Illustrations of odd fractures (Definition 23). For each non-subdivision vertex, there are only three ways to satisfy 2. in Definition 23. This observation is used in Lemma 24 to show that the number of odd fractures is a power of 3. (*Bottom*:) Elements of $\mathsf{ColSub}(kP_2 \to (G, c_E))$ inducing fractures of H^2 such that each partition has at most two blocks. Lemma 25 shows that those are precisely the odd fractures of H^2

two singleton blocks, or of one block of size 2). Let $e = \{u, v\} \in E(H)$ be the edge corresponding to s, that is, s was created by subdividing e. Since $H^2 \sharp \tau$ is a union of P_2, we can infer that τ_v and τ_u contain a singleton block (otherwise we would have created a connected component which is not isomorphic to P_2). Now recall that both u and v have degree 4, since H is 4-regular. We obtain a contradiction as follows: By assumption of the lemma, we know that τ_v and τ_u can have at most two blocks. Since we have just shown that both contain a singleton block, it follows that both τ_v and τ_u contain one further block of size 3. However, a block of size 3 yields a vertex of degree 3 in the fractured graph $H^2 \sharp \tau$, contradicting the fact that $H^2 \sharp \tau$ consists only of disjoint P_2.

Thus we have established that, for each $s \in S_E$, the partition τ_s consists of two singleton blocks. Given this fact, the only way for $H^2 \sharp \tau$ being a disjoint union of P_2

is that each partition τ_v, for $v \in V(H) = V(H^2) \setminus S_E$, consists of two blocks of size 2. □

We are now able to prove our hardness result.

Lemma 26 *The problem $\oplus\text{SUB}(\mathcal{F}_{P_2})$ is $\oplus$W[1]-hard. Furthermore, on input kP_2 and G, the problem cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ for any function f, unless rETH fails.*

Proof We reduce from $\oplus$CP-HOM$(\mathcal{Q})$, which, by Lemma 22 and Theorem 12, is $\oplus$W[1]-hard and for $H' \in \mathcal{Q}$, it cannot be solved in time $f(|H'|) \cdot |G|^{o(|V(H')|/\log|V(H')|)}$, unless rETH fails.

Let H' and (G, c) be an input instance to $\oplus$CP-HOM$(\mathcal{Q})$. There is an algorithm that takes as input a graph $H' \in \mathcal{Q}$ and finds a graph $H \in \Delta$ with $H' = H^2$—this is basically 2-colouring. The run time of this algorithm depends on $|H'|$ but clearly not on (G, c). Let $k = 2|V(H)|$ and note that $|E(H^2)| = 2|E(H)| = 4|V(H)| = 2k$. We consider the graph G as a $2k$-edge-coloured graph, coloured by c_E. That is, each edge $e = \{x, y\}$ of G is assigned the colour $c_E(e) = \{c(x), c(y)\}$ which is an edge of $H' = H^2$.

Now, for *any* H^2-coloured graph (G', c') recall that $\mathsf{ColSub}(kP_2 \to (G', c'_E))$ is the set of subgraphs of G' that are isomorphic to kP_2 and that include each edge colour (each edge of H^2) precisely once. We will see later that $\oplus\mathsf{ColSub}(kP_2 \to (G', c'_E))$ can be computed using our oracle for $\oplus\text{SUB}(\mathcal{F}_{P_2})$ using the principle of inclusion and exclusion.

It was shown in [30, Lemma 4.1] that there is a unique function a such that, for every H^2-coloured graph (G', c'),

$$\#\mathsf{ColSub}(kP_2 \to (G', c'_E)) = \sum_{\rho} a(\rho) \cdot \mathsf{Hom}(H^2 \sharp \rho \to (G', c')). \tag{5}$$

where the sum is over all fractures of H^2. As in Sect. 3.1 from [30, Corollary 4.3] we know that

$$a(\top) = \sum_{\rho \in \mathsf{F}(kP_2, H^2)} \prod_{w \in V(H^2)} (-1)^{|\rho_w|-1} \cdot (|\rho_w| - 1)!, \tag{6}$$

where $\top$ is the fracture in which each partition consists only of one block and $\mathsf{F}(kP_2, H^2)$ is the set of all fractures ρ of H^2 such that $H^2 \sharp \rho \cong kP_2$.

Our next goal is to show that $a(\top) = 1 \mod 2$. First, suppose that a fracture ρ contains a partition ρ_w with at least three blocks. Then $(|\rho_w| - 1)! = 0 \mod 2$. Thus such fractures do not contribute to $a(\top)$ if arithmetic is done modulo 2. Next, note that if, for each w, the partition ρ_w contains at most 2 blocks, then

$$\prod_{w \in V(H^2)} (-1)^{|\rho_w|-1} \cdot (|\rho_w| - 1)! = 1 \mod 2.$$

Let $\mathsf{Odd}(kP_2, H^2)$ be the set of all fractures ρ of H^2 such that $H^2 \sharp \rho \cong kP_2$ and each partition of ρ consists of at most 2 blocks. Our analysis then yields $a(\top) =$

$|\mathsf{Odd}(kP_2, H^2)| \mod 2$. Finally, Lemma 25 states that $\mathsf{Odd}(kP_2, H^2)$ is precisely the set of odd fractures, and Lemma 24 thus implies that $|\mathsf{Odd}(kP_2, H^2)| = 1 \mod 2$. Consequently, $a(\top) = 1 \mod 2$ as well, and we have achieved the goal.

Next we can proceed similarly to the case of triangle packings. As in that case, the goal is to compute $\oplus\mathsf{Hom}((H^2, \mathsf{id}_{H^2}) \to (G, c)))$ which is equal to $\oplus\mathsf{Hom}((H^2 \sharp \top, c_\top) \to (G, c))$. Since $a(\top)$ is odd, we can invoke Lemma 14 to recover this term by evaluating the entire linear combination (5), that is, if we can evaluate the function $\oplus\mathsf{ColSub}(kP_2 \to \star)$. This can be done by using Lemma 15. Each call to the oracle is of the form $\oplus\mathsf{Sub}(kP_2 \to \hat{G})$ where $|\hat{G}|$ is bounded by $f(k) \cdot |G|$.

Now recall that $k \in \Theta(|V(H)|)$. By Lemma 22, we thus have $k = \Theta(\mathsf{tw}(H^2))$. Hence our reduction yields $\oplus$W[1]-hardness and transfers the conditional lower bound under rETH as desired. □

We can now conclude the treatment of hereditary pattern classes by proving Theorem 4, which we restate for convenience.

Theorem 4 Let $\mathcal{H}$ be a hereditary class of graphs. If $\mathcal{H}$ is matching splittable, then $\oplus$SUB($\mathcal{H}$) is fixed-parameter tractable. Otherwise, the problem is $\oplus$W[1]-complete and, assuming rETH, cannot be solved in time $f(|H|) \cdot |G|^{o(|V(H)|/\log|V(H)|)}$ for any function f.

Proof The fixed-parameter tractability result was shown in [12]. For the hardness result, using the fact that $\mathcal{H}$ is not matching splittable and Theorem 16 we obtain four cases.

- If $\mathcal{H}$ contains all cliques or all bicliques, then hardness follows from Lemma 17.
- If $\mathcal{H}$ contains all triangle packings, then hardness follows from Lemma 21.
- If $\mathcal{H}$ contains all P_2-packings, then hardness follows from Lemma 26.

Since the case distinction is exhaustive, the proof is concluded. □

4 Classification for Trees

Our overall goal is to prove Theorem 5, which we restate for convenience:

Theorem 5 Let $\mathcal{T}$ be a recursively enumerable class of trees. If $\mathcal{T}$ is matching splittable, then $\oplus$SUB($\mathcal{T}$) is fixed-parameter tractable. Otherwise $\oplus$SUB($\mathcal{T}$) is $\oplus$W[1]-complete.

Outline of Sect. 4

We begin our analysis by investigating the structural properties of classes of trees that are not matching splittable. In Lemma 32 we prove that for each such class $\mathcal{T}$ (at least) one of the following parameters are unbounded: The *fork number* (Definition 29), the *star number* (Definition 30), or the C-*number* (Definition 31).

The remainder of this section is then split into three, largely independent, parts: Sect. 4.1 establishes hardness of $\oplus$SUB($\mathcal{T}$) for classes of trees $\mathcal{T}$ of unbounded C-number, Sect. 4.2 shows hardness for unbounded star number, and Sect. 4.3 shows hardness for unbounded fork number.

We start by introducing some terminology for trees which will be used in the remainder of this section.

Definition 27 (*2-paths*) A 2-*path* of length a of a tree T is a path $x_0, x_1, \ldots, x_a$ such that $\deg(x_0) \neq 2$, $\deg(x_1) = \cdots = \deg(x_{a-1}) = 2$ and $\deg(x_a) \neq 2$.

Next we introduce rays, which are restricted 2-paths that will be crucial in our analysis.

Definition 28 (*Source, ray,* $\mathsf{deg}_{\mathsf{L},a}$, $\mathsf{deg}_{\mathsf{L}}$, *and* $\mathsf{deg}_{\mathsf{NL}}$) Let T be a tree. A *source* of T is any vertex with degree greater than 2. A *ray* of length a of T is a 2-path $x_0, x_1, \ldots, x_a$ such that $\deg(x_0) > 2$ and $\deg(x_a) = 1$. We call x_0 the *source of the ray*. Given a vertex s of degree at least 3, we write $\mathsf{deg}_{\mathsf{L},a}(s)$ for the number of rays of length a with source s. We set

$$\mathsf{deg}_{\mathsf{L}}(s) := \sum_a \mathsf{deg}_{\mathsf{L},a}(s).$$

Finally, we set $\mathsf{deg}_{\mathsf{NL}}(s) := \deg(s) - \mathsf{deg}_{\mathsf{L}}(s)$.

Next, we introduce parameters $\mathrm{F}_{a,b}$, S_c and C_d. Our goal is then to show that, for every non-matching-splittable class of trees, at least one of those two parameters is unbounded.

Definition 29 (*Forks and* $\mathrm{F}_{a,b}$) Let a, b be positive integers. A source s of a tree T is called an *a-b-fork* if $\mathsf{deg}_{\mathsf{NL}}(s) = 1$ and one of the following is true

- $a \neq b$ and $\mathsf{deg}_{\mathsf{L},a}(s), \mathsf{deg}_{\mathsf{L},b}(s) > 0$.
- $a = b$ and $\mathsf{deg}_{\mathsf{L},a}(s) > 1$.

The *a-b-fork number* of T, denoted by $\mathrm{F}_{a,b}(T)$ is the maximum size of an independent set containing only *a*-*b*-forks. Finally, we say that a class of trees $\mathcal{T}$ has *unbounded fork number* if for every positive integer B there are positive integers a and b and a tree $T \in \mathcal{T}$ such that $\mathrm{F}_{a,b}(T) \geq B$.

Definition 30 (*Stars and* S_c) A *star* of size $k > 1$ in a tree T is a collection of k distinct rays that have a common source s. For a positive integer $c \geq 3$, a *c-star* of size k in a tree T is a collection of k distinct rays of length c that have a common source s.

The *c-star number* of a tree T, denoted by $\mathrm{S}_c(T)$ is the maximum size of a c-star in T. Finally, we say that a class of trees $\mathcal{T}$ has *unbounded star number* if for every positive integer B there exists $c \geq 3$, and a tree $T \in \mathcal{T}$ such that $\mathrm{S}_c(T) \geq B$.

Definition 31 (C-*gadgets and* C_d) A C-gadget,[4] of *order* d and *length* k in a tree T is a path $x_0, \ldots, x_k$ such that one of the following is true for each inner vertex $x_i \in \{1, \ldots, k-1\}$:

(i) $\deg(x_i) = 2$, that is $N(x_i) = \{x_{i-1}, x_{i+1}\}$, or
(ii) x_i is a source and every neighbour $v \in N(x_i) \setminus \{x_{i-1}, x_{i+1}\}$ is contained in a ray of length at most d from x_i to a leaf.

[4] C stands for *caterpillar* the shape of which resembles the structure of a C-gadget.

The C_d-number of a tree T, denoted by $C_d(T)$ is the length of the longest C-gadget of order d. Finally, we say that a class of trees $\mathcal{T}$ has *unbounded* C-*number* if there exists $d > 0$ such that for every positive integer B, and a tree $T \in \mathcal{T}$ such that $C_d(T) \geq B$.

Note that the ordering of the quantifiers in the definition of the C_d-number is different from the ordering in the definition of the c-star-number. This is due to technical reasons which are important for the proof of Lemma 32.

Lemma 32 *Let $\mathcal{T}$ be a class of trees. If $\mathcal{T}$ is not matching splittable, then $\mathcal{T}$ has either unbounded fork number, unbounded star number, or unbounded* C-*number.*

Proof We can assume that there is an overall bound d on the length of 2-paths in trees in $\mathcal{T}$: Otherwise, $\mathcal{T}$ already has unbounded C-number (see (i) in Definition 31)). Hence the length of every ray in any tree in $\mathcal{T}$ is bounded by d as well. Thus

- $\mathcal{T}$ has unbounded fork number if and only if for every positive integer B there are $a, b \in \{1, \ldots, d\}$ and a tree $T \in \mathcal{T}$ such that $F_{a,b}(T) \geq B$.
- $\mathcal{T}$ has unbounded C-number if and only if C_d is unbounded in $\mathcal{T}$ (see Definition 31)).
- $\mathcal{T}$ has unbounded star number if and only if for every positive integer s there is a $c \in \{3, \ldots, d\}$ and a tree $T \in \mathcal{T}$ such that $S_c(T) \geq s$.

We split the proof into two cases.

Case 1 $\mathcal{T}$ has unbounded diameter.

In Case 1, we show that $\mathcal{T}$ has unbounded fork number or unbounded C-number. If C_d is unbounded in $\mathcal{T}$ then $\mathcal{T}$ has unbounded C-number and we are done so assume that there is a constant h such that $C_d(T) \leq h$ for every $T \in \mathcal{T}$.

Now let B be a positive integer. We show that there are $a, b \in \{1, \ldots, d\}$ and $T \in \mathcal{T}$ with $F_{a,b}(T) \geq B$. To this end, we use the premise that $\mathcal{T}$ has unbounded diameter. Let $k > (h + 2)(Bd^2 + 1)$ be a positive integer, and let $T \in \mathcal{T}$ be such that there is a path $P = s, p_0, \ldots, p_k, t$ in T. Observe that the deletion of all edges in P decomposes T into a family of disjoint subtrees. We write T_i for the subtree that contains p_i. Now decompose P into segments $P_1, P_2, \ldots$ of length $h+2$. Note that a segment $P_j = p_{j_0}, \ldots, p_{j_{h+2}}$ yields a C-gadget of order d and length $> h$ if and only if T_{j_i} is either a star or an isolated vertex for each $i \in \{1, \ldots, h + 1\}$.

Since no such C-gadgets exist by assumption, we obtain that each segment P_j of the path P contains a vertex p_{i_j} such that T_{i_j} is neither a star nor an isolated vertex.

Assume that T_{i_j} is rooted at p_{i_j}. Since T_{i_j} is neither an isolated vertex nor a star, there must be a (proper) descendant v_{i_j} of p_{i_j} (in T_{i_j}) such that v_{i_j} is an (a_{i_j}, b_{i_j})-fork for some $a_{i_j}, b_{i_j} \in \{1, \ldots, d\}$. Now note that there are at most d^2 pairs of integers in $\{1, \ldots, d\}$. Since we have at least one fork for every segment and since there are at least $\lfloor k/(h + 2)\rfloor > Bd^2 + 1$ segments, we thus obtain by the pigeon-hole principle that there is a pair $a, b \in \{1, \ldots, d\}$ such that, for at least B segments P_{i_j}, the node v_{i_j} is an (a, b)-fork in T_{i_j} and thus also in T. Since those forks are pairwise non-adjacent, we obtain, as desired, that the (a, b)-fork number of T is at least B, concluding Case 1.

Case 2 $\mathcal{T}$ has bounded diameter.

Let D be the assumed upper bound on the diameter of trees in $\mathcal{T}$. If $\mathcal{T}$ has unbounded star number then we are finished. Assume instead that $\mathcal{T}$ has bounded star number. Then there is a positive integer s such that for all $c \in \{3, \dots, d\}$ and every tree $T \in \mathcal{T}$, $\mathrm{S}_c(T) < s$. We will show that $\mathcal{T}$ has unbounded fork number. Consider any positive integer B. We will show that there are $a, b \in \{1, \dots, d\}$ and $T \in \mathcal{T}$ with $\mathrm{F}_{a,b}(T) \geq B$.

Let $k > (D+1)(Bd^2+1)(d^2s+1)$ be a positive integer. Since $\mathcal{T}$ is not matching splittable, there is a tree $T \in \mathcal{T}$ whose matching-split number is at least k. Note that T is not a path since every path with matching-split number at least k has length greater than $k > D$, contradicting the bound on the diameter.

Now fix any vertex r of T as the root. Given a vertex v of T, we write T_v for the subtree rooted at v (assuming that r is the overall root). We call v a *rooted fork* if T_v is a star—observe that each rooted fork must indeed be a fork. Let f be the number of rooted forks. Similar to the argument in Case 1, if $f > Bd^2 + 1$, then by the pigeon-hole principle there are $a, b \in \{1, \dots, d\}$ such that $F_{a,b}(T) \geq B$.

Hence assume for contradiction that $f \leq Bd^2 + 1$. Let $\mathcal{R}$ be the set of all rays of T and recall that each ray in $\mathcal{R}$ is, by definition, a 2-path of the form $R = x_0, \dots, x_{d'}$ for $d' \leq d$, where $\deg(x_0) > 2$ and $x_{d'}$ is a leaf. We call a ray R *long* if $d' \geq 3$. Note that the source of every ray must either be a rooted fork, or it must lie on a path from the root r to one of the rooted forks.

Let T' be the subtree of T induced by all vertices that lie on paths between r and a rooted fork (including r and all rooted forks). Since there are f rooted forks and the depth of T is bounded by D, $|V(T')| \leq (D+1)f \leq (D+1)(Bd^2+1)$.

Consider a vertex v of T'. Assume for contradiction that v is the source of $> ds$ long rays (in T). Recall that for all $c \in \{3, \dots, d\}$ we have that $\mathrm{S}_c(T) < s$. Recall further that each long ray has length d' for some $3 \leq d' \leq d$. Thus we obtain a contradiction by the pigeon-hole principle.

Now let S be the set containing all vertices of T' and all vertices of long rays. Noting that each long ray has length at most d, and that the source of each long ray must be a vertex of T' by construction, we can use the observation that each vertex of T' is the source of at most ds long rays to (generously) bound

$$|S| \leq |V(T')| + |V(T')| \cdot d \cdot ds.$$

Note further that $T[V(T) \setminus S]$ consists only of isolated edges and vertices: The only vertices in $V(T) \setminus S$ are non-source vertices of rays of length < 3, the sources of which are in T'. Thus, S is a splitting set. Finally, recalling that $|V(T')| \leq (D+1)f \leq (D+1)(Bd^2+1)$, we have

$$|S| \leq |V(T')| + |V(T')| \cdot d \cdot ds \leq (D+1)(Bd^2+1)(d^2s+1),$$

contradicting the fact that the matching-split number of T is strictly larger than $(D+1)(Bd^2+1)(d^2s+1)$. This concludes Case 2, and hence the proof. □

In the next three subsections, we will prove hardness of $\oplus$SUB($\mathcal{T}$) for non-matching-splittable $\mathcal{T}$ in each of the three cases given by Lemma 32.

4.1 Unbounded C-Number

For our hardness proof, it will be useful to find a proper sub-gadget of a C-gadget in a tree.

Definition 33 (*Strong* C-*gadgets, junctions, and closedness*) Let $C = x_0, \dots, x_L$ be a C-gadget of order d and length L in a tree T. We call C a *strong* C-gadget with k *junctions* if there are integers $0 = i_0 < i_1 < \cdots < i_k < i_{k+1} = L$ such that

(I) for all $j \in \{0, \dots, k\}$, $i_{j+1} - i_j > 2d$, and
(II) for all $j \in \{1, \dots, k\}$, x_{i_j} is the source of a ray R_j of length d that does not contain one of the neighbours $x_{i_j - 1}$ and $x_{i_j + 1}$ of x_{i_j}. The vertices $x_{i_1}, \dots, x_{i_k}$ are called the *junctions*.

Finally, a strong C-gadget is called *closed* if neither x_{i_1} nor x_{i_k} are forks.[5]

Consider the bottom part of Fig. 4 for a visualisation. We start with the following lemma which establishes the existence of a strong C-gadget with many junctions inside a long enough C-gadget.

Lemma 34 *Let T be a tree such that the longest 2-path in T has length $d \geq 1$, and let k be a positive integer. Then there exists $L > 0$ (only depending on k and d) such that the following is true: If T contains an* C-*gadget of order d and length L, then there exists $1 \leq d' \leq d$ such that T contains a strong* C-*gadget of order d' with at least k junctions.*

Proof Let $f(x) = x/(k+1) - 2d - 1$ and let L be large enough such that $f^d(L) > d$. Let $H^d = x_0, \dots, x_L$ be a C-gadget of order d and length L in T.

Let $d' = d$. Note that $H^{d'}$ is a C-gadget of order d' and length at least $L = f^{d-d'}(L)$ in T. For each graph $H^{d'}$ with $d' \geq 1$ we will either

(1) construct a strong C-gadget with k junctions with order d', or
(2) find a subsequence $H^{d'-1}$ of $H^{d'}$ that is an C-gadget of order $d' - 1$ of length at least $f^{d-(d'-1)}(L)$.

If we ever do (1) we are finished. If from $d' = 1$ we do (2) then we find a 2-path of length at least $f^d(L) > d$, which is a contradiction.

Here is how we proceed from $H^{d'} = y_0, \dots, y_\ell$. We set $i_0 = 0$. Then iteratively, for each $j \in \{1, \dots, k\}$ we will either construct $H^{d'-1}$ as in (2) or we find $i_j \in \{i_{j-1} + 2d + 1, \dots, \ell\}$ such that y_{i_j} is the source of a length-d' ray that does not contain $y_{i_j} - 1$ or $y_{i_j} + 1$. If we succeed in defining $i_1, \dots, i_k, i_{k+1}$ in this way then $y_0, \dots, y_{i_{k+1}}$ is a strong C-gadget with k junctions of order d' so (1) is satisfied.

Let us now make this argument rigorous; again, assume that $H^{d'} = y_0, \dots, y_\ell$ is a C-gadget of order d' and length $\ell \geq f^{d-d'}(L)$. Set $i_0 = 0$ and, starting with $j = 0$, proceed iteratively as follows:

1. Let S_j be the set of all indices $i \in \{i_{j-1} + 2d + 1, \dots, \ell\}$ such that y_i is the source of a length-d' ray that does not contain y_{i-1} and y_{i+1}.

[5] The condition of being closed rules out the special case in which x_0 or x_L are leaves of T. More generally it rules out the case where there is a ray from x_1 including x_0 or from x_k including x_L.

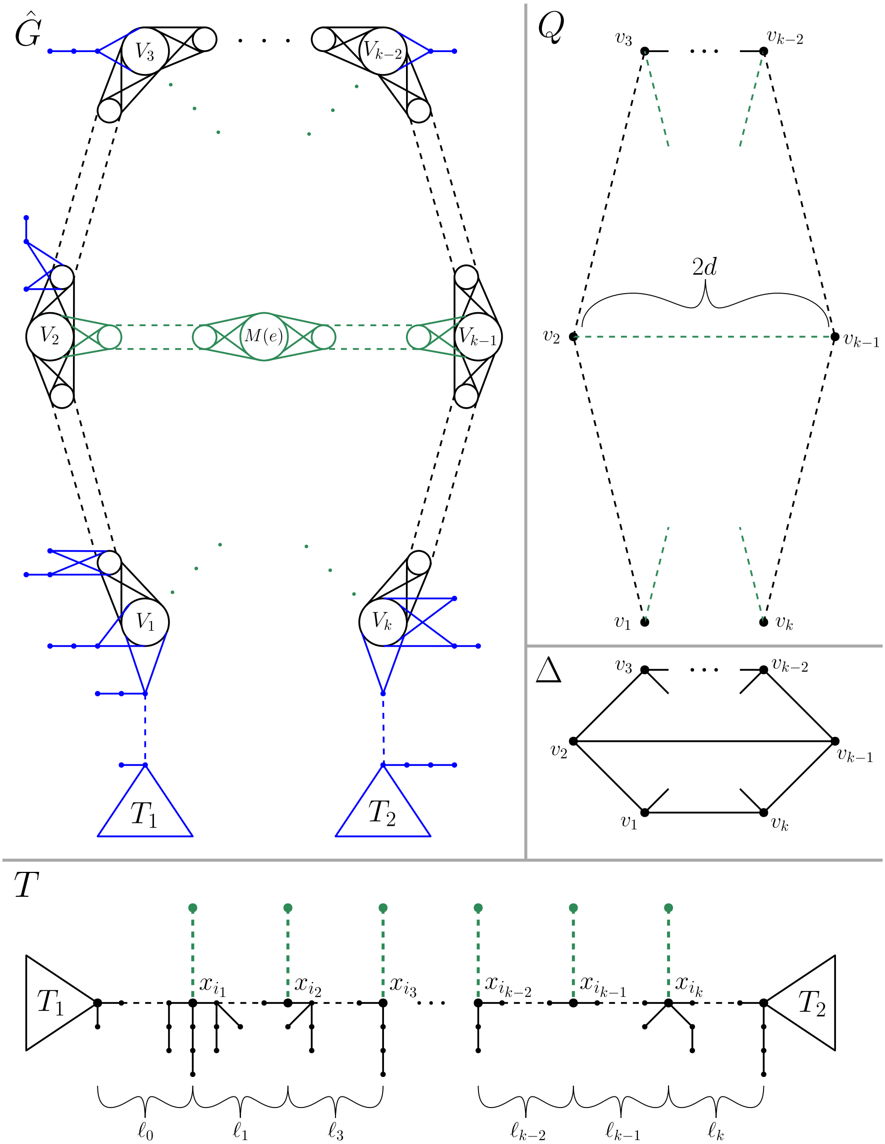

Fig. 4 (*Below*): The tree T containing a closed strong C-gadget of order d; the green dashed lines are rays of length d. (*Left*): The construction of $\hat{G} = \hat{G}(G, c, T, H)$; note that the removal of the vertices and edges coloured blue yields G (see Definition 42), and note that G is Q-coloured as depicted. (*Right*): The graphs Δ and $Q = Q(\Delta, T, H)$; we assume in the picture that $\{v_2, v_{k-1}\}$ is an edge of Δ (Color figure online)

2. If $S_j = \emptyset$ then set $\mathsf{stop} = j$ and terminate. Otherwise, set $i_j = \min S_j$ and $j \leftarrow j + 1$, and go back to 1.

We now distinguish two cases: If $\mathsf{stop} \geq k+1$, then we found indices $i_0, \dots, i_{k+1}$ such that $\hat{H}^{d'} := y_0, \dots, y_{i_{k+1}}$ is a strong hardness gadget of order d' with k junctions; hence we achieved (1) and we are done. Otherwise we have $\mathsf{stop} < k+1$. Let $I_j := \{i_j, \dots, i_{j+1} - 1\}$ for all $0 \leq j < \mathsf{stop}$, and let $I_{\mathsf{stop}} = \{i_{\mathsf{stop}}, \dots, \ell\}$. By the pigeon-hole principle, at least one of those intervals, say $I_{j'}$, has size at least $\ell/(\mathsf{stop}+1) \geq \ell/(k+1)$. Now, by construction of our iterative procedure above, we find that the sub-interval $\{i_{j'} + 2d + 1, \dots, i_{j'+1} - 1\} \subseteq I_{j'}$ contains no index i such that y_i is the source of a length-d' ray that does not contain y_{i-1} and y_{i+1}. Thus, the subsequence $H^{d'-1} := y_{i_{j'}+2d+1}, \dots, y_{i_{j'+1}-1}$ constitutes a C-gadget of order $d'-1$. Furthermore, $H^{d'-1}$ has length at least $\ell/(k+1) - 2d - 1 = f(\ell)$. Since $\ell \geq f^{d-d'}(L)$, and since f is monotonically increasing, we find that $f(\ell) \geq f^{d-(d'-1)}(L)$. Hence we achieved (2) and we can conclude this case as well. □

Now, by removing the first and the last junction, we can also ensure the existence of a closed strong C-gadget

Corollary 35 *Let T be a tree such that the longest 2-path in T has length $d \geq 1$, and let k be a positive integer. Then there exists $L > 0$ (only depending on k and d) such that the following is true: If T contains an C-gadget of order d and length L, then there exists $1 \leq d' \leq d$ such that T contains a* closed *strong C-gadget of order d' with at least k junctions.*

Proof Use Lemma 34 with $k+2$ rather than k and observe that every strong C-gadget with $k+2$ junctions also yields a closed strong C-gadget with k junctions by removing i_1 and i_{k+2} from the list of indices. Since x_{i_1} and $x_{i_{k+2}}$ must have degree at least 3 (they are inner vertices of a C-gadget and they are junctions), we obtain that neither x_{i_2} and $x_{i_{k+1}}$ can be forks of T. □

4.1.1 Constructions of Q and $\hat{G}$

For the scope of this subsection, to avoid notational clutter, we assume the following are given:

- Positive integers k and d.
- A tree T that contains a closed strong C-gadget $H = x_0, \dots, x_\ell$ of order d with k junctions $x_{i_1}, \dots, x_{i_k}$. Additionally, for each $j \in [k]$, we fix a ray $R_j = x_{i_j}, r_j^1, \dots, r_j^d$ of length d, the source of which is x_{i_j} and which does not contain one of the neighbours x_{i_j-1} and x_{i_j+1}—note that the R_j must exist as the x_{i_j} are junctions.
- A k-vertex cubic graph Δ containing a Hamiltonian cycle $v_1, \dots, v_k, v_1$.

We emphasise that the set of edges of Δ not contained in the Hamilton cycle must constitute a perfect matching, that is, a set of $k/2$ pairwise non-incident edges. This must be satisfied since Δ is cubic.

Definition 36 The *core* of H, denoted by $C(H)$, contains the subsequence $x_{i_1}, x_{i_1+1}, \ldots, x_{i_k-1}, x_{i_k}$ and the vertices of the rays R_j, that is

$$C(H) := \{x_{i_1}, x_{i_1+1}, \ldots, x_{i_k-1}, x_{i_k}\} \cup \bigcup_{j=1}^{k} V(R_j).$$

Definition 37 ($Q(\Delta, T, H)$ *and* τ_Q) Set $\ell_j := i_{j+1} - i_j$. The graph $Q = Q(\Delta, T, H)$ is obtained from Δ as follows:

1. The edge $\{v_k, v_1\}$ is deleted.
2. For each $j \in \{1, \ldots, k-1\}$ the edge $\{v_j, v_{j+1}\}$ is replaced by a path of length ℓ_j:

$$P_j = v_j, u_j^1, \ldots, u_j^{\ell_j - 1}, v_{j+1},$$

   where the u_j^t are fresh vertices.
3. Each edge $e = \{v_i, v_j\}$ not contained on the Hamilton cycle, i.e., $j \notin \{i-1, i+1\}$, is replaced by a path $P_{i,j}$ of length $2d$:

$$P_{i,j} = v_i, w_i^1, \ldots, w_i^{d-1}, m(e), w_j^{d-1}, \ldots, w_j^1, v_j,$$

   where the w_i^t and w_j^t are fresh vertices.

Finally $\tau = \tau(\Delta, T, H)$ is a fracture of Q defined as follows: For each $m(e)$, the partition $\tau_{m(e)}$ contains two singleton blocks, and for all remaining vertices v of Q the partition τ_v only contains one block.

Since Δ, T and H are fixed in this subsection, to avoid notational clutter, we just write Q and τ, rather than $Q(\Delta, T, H)$ and $\tau(\Delta, T, H)$.

It turns out that Q is isomorphic to a quotient graph of $T[C(H)]$ obtained by identifying the endpoints of the rays R_i and R_j for every $\{v_i, v_j\} \in E(\Delta)$ with $j \notin \{i-1, i+1\}$. This induces a homomorphism from $T[C(H)]$ to Q that will be useful in the construction of $\hat{G}$; hence we explicitly define this mapping below:

Definition 38 (γ) We define a function $\gamma : C(H) \to V(Q)$ as follows.

1. We map the sequence $x_{i_1}, x_{i_1+1}, \ldots, x_{i_k-1}, x_{i_k}$ in $C(H)$ to the sequence $v_1, \ldots, v_k$ in Q. More precisely, for each $j \in \{1, \ldots, k-1\}$ and $t \in \{1, \ldots, \ell_j - 1\}$, we set $\gamma(x_{i_j}) := v_j$ and $\gamma(x_{i_j+t}) := u_j^t$.
2. For each edge $e = \{v_i, v_j\}$ of Δ with $j \notin \{i-1, i+1\}$, we map $V(R_i)$ and $V(R_j)$ to the path $P_{i,j}$. More precisely, for each $t \in \{1, \ldots, d-1\}$ we set $\gamma(r_i^t) := w_i^t$ and $\gamma(r_j^t) = w_j^t$. Furthermore, we set $\gamma(r_i^d) := m(e) =: \gamma(r_j^d)$. (Note that the images of the sources of the rays R_i and R_j are already set in 1.)

Observation 39 *The function γ is an edge-bijective homomorphism from $T[C(H)]$ to Q.*

Let us provide the induced egde-bijection explicitly:

Definition 40 $(E', \gamma_{\mathrm{E}})$ Define $E' := E(T[C(H)])$, that is, $E' \subseteq E(T)$ contains all edges on the sub-path $x_{i_1}, \ldots, x_{i_k}$ of H and all edges of the rays $R_1, \ldots, R_k$. We write $\gamma_{\mathrm{E}} : E' \to E(Q)$ for the edge-bijection from E' to $E(Q)$ induced by the homomorphism γ.

Now let (G, c) be a Q-coloured graph. We state the following fact explicitly, since it will be crucial in our construction:

Observation 41 *Let (G, c) be a Q-coloured graph. The mapping $c_E \circ \gamma_E^{-1}$ is a map from $E(G)$ to E'.*

Our goal is to construct a graph $\hat{G} = \hat{G}(G, c, T, H)$ from G, and an edge-colouring $\hat{\gamma} : E(\hat{G}) \mapsto E(T)$ whose range is $E(T)$ such that

$$\oplus\mathsf{Emb}((Q\#\tau, c_\tau) \to (G, c)) = \oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma})),$$

that is, the number of colour-preserving embeddings from the fractured graph $Q\#\tau$ to (G, c) is equal, modulo 2, to the number of subgraphs of $\hat{G}$ that are isomorphic to T and that contain each edge-colour in $E(T)$ precisely once.

For what follows, let $V(\mathcal{R}) := \cup_{j=1}^{k} V(R_j)$ be the set of all vertices of the rays $R_1, \ldots, R_k$. We are now able to define $\hat{G} = \hat{G}(G, c, T, H)$; the construction is illustrated in Fig. 4. The definition uses the function c_{E} introduced in Definition 6 and the functions γ and γ_{E} introduced in Definitions 38 and 40, respectively. It also uses the mapping $c_{\mathrm{E}} \circ \gamma_{\mathrm{E}}^{-1}$ from $E(G)$ to E' (see Observation 41).

Definition 42 $(\hat{G}(G, c, T, H), \hat{\gamma}(G, c, T, H))$ Let (G, c) be a Q-coloured graph. The pair $(\hat{G}, \hat{\gamma}) = (\hat{G}(G, c, T, H), \hat{\gamma}(G, c, T, H))$ is an edge-coloured graph constructed as follows, where the co-domain of $\hat{\gamma}$ is $E(T)$:

(A) The graph $\hat{G}$ contains G as a subgraph. For each $e \in E(G)$, define $\hat{\gamma}(e) = \gamma_{\mathrm{E}}^{-1}(c_E(e))$.
(B) The vertex set of $\hat{G}$ is the union of $V(G)$ and $V(T) \setminus C(H)$.
(C) Pairs of vertices in $V(T) \setminus C(H)$ are connected by an edge in $\hat{G}$ if and only if they are adjacent in T. For each such edge e, $\hat{\gamma}(e) = e$.
(D) The remaining edges of $\hat{G}$ are defined as follows. For each edge $e \in E(T)$ that connects a vertex $z \in V(T)\setminus C(H)$ to a vertex $y \in C(H)$ there are corresponding edges in $\hat{G}$. These edges connect z to all vertices $g \in V(G)$ such that $c(g) = \gamma(y)$ For each such edge e' in $\hat{G}$, $\hat{\gamma}(e') = e$.

Observe that for each element $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ the induced subgraph

$$T_{\mathsf{col}}[G] := T_{\mathsf{col}}[V(T_{\mathsf{col}}) \cap V(G)] of T_{\mathsf{col}} is an$$

edge-colourful subgraph in G, that is, $T_{\mathsf{col}}[G]$ contains precisely one edge per edge-colour of G under the edge colouring $\hat{\gamma}$ hence it contains precisely one edge per edge-colour of G under c_{E}. As shown in Section 3 in the full version [31] of [32], $T_{\mathsf{col}}[G]$ thus induces a fracture $\rho = \rho(T_{\mathsf{col}})$ of Q: Two edges $\{v, w\}$ and $\{v, y\}$ of Q are in the same block in the partition ρ_v corresponding to vertex v of Q if and only if the edges of $T_{\mathsf{col}}[G]$ that are coloured $\gamma_{\mathrm{E}}^{-1}(\{v, w\})$ and $\gamma_{\mathrm{E}}^{-1}(\{v, y\})$ are adjacent. In what follows, we show that ρ must always be equal to $\tau(\Delta, T, H)$ (see Definition 37).

Lemma 43 *For every* $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ *we have that* $\rho(T_{\mathsf{col}}) = \tau(\Delta, T, H)$.

Proof To avoid notational clutter, we set $\rho := \rho(T_{\mathsf{col}})$ and $\tau := \tau(\Delta, T, H)$. Let T_1 and T_2 be the subtrees of T attached to the ends of the C-gadget H as shown in the bottom part of Fig. 4.

We first give an overall intuition of the proof; consider Fig. 5 for an illustration. Since T_{col} is isomorphic to T, there must be a (unique) path connecting T_1 and T_2 in $\hat{G}$ (recall that, since T_{col} is *edge-colourful* and since every edge in T_1 and T_2 has a different colour—see (C) in Definition 42 — T_{col} must contain all edges in T_1 and T_2). We claim that this path must follow the outer cycle in $\hat{G}$, in which case the designated rays in $\mathcal{R}$ of length d at the junctions must follow the inwards direction and thus induce τ. To see why the path connecting T_1 and T_2 must follow the outer cycle, first recall that V_j is the subset of $V(G)$ coloured by c with v_j. Then recall that the path between V_j and V_{j+1} along the outer cycle in $\hat{G}$ has length $\ell_j \geq 2d + 1$. Hence the designated rays in $\mathcal{R}$ cannot be used to cover all edge colours in the path between V_j and V_{j+1}.

We next provide a rigorous argument. Let

$$S := V(T_1) \cup V(T_2) \cup \{x_0, \ldots, x_{i_1-1}\} \cup \{x_{i_k+1}, \ldots, x_{k+1}\}.$$

Note that S is a subset of $V(T) \setminus V(H)$ hence it is a subset of $V(T)$ and of $V(\hat{G})$.

We first claim that every fork and every ray of length $> d$ of T must be fully contained in the subgraph of T induced by S. This claim follows from the definition of closed strong C-gadgets. In particular, the condition of being closed implies that neither x_{i_1} nor x_{i_k} is a fork.

As a consequence, every fork and every ray of length greater than d of T_{col} must be contained in the subgraph of $\hat{G}$ induced by S as well. Additionally, this implies that none of the vertices in $T_{\mathsf{col}}[G]$ can be a fork or the source of a ray of length $> d$ in T_{col}—otherwise, T_{col} would have either more forks or more rays of length $> d$ than T, contradicting the fact that T_{col} and T are isomorphic.

Recall that $V_1, \ldots, V_k$ denote the subsets of vertices of G that are coloured by c with $v_1, \ldots, v_k$. Now let P be the (unique) path P in T_{col} that connects T_1 with T_2. Then, starting with V_1 and ending with V_k, the path P must pass through a sequence of colour classes $V_1 = V_{j_1}, V_{j_2}, \ldots, V_{j_t} = V_k$ of G. The following claim formalises the idea that this sequence must correspond to the Hamilton cycle $v_1, \ldots, v_k$ in Δ.

Claim We have $t = k$ and $V_{j_i} = V_i$ for each $i \in [k]$.

Before proving the claim, we show that it implies the lemma. Since, from the claim, P must follow the outer cycle, the fracture $\rho = \rho(T_{\mathsf{col}})$ induced by T_{col} must split the inner paths of length $2d$ (otherwise T_{col} would contain a cycle). However, since there are no sources or rays of length greater than d outside of S in T_{col}, ρ must split all of the inner length-$2d$ paths at the central vertex $m(e)$. Furthermore, it cannot split additional vertices since this would disconnect T_{col}. Thus, ρ is the fracture τ, concluding the proof. ◻

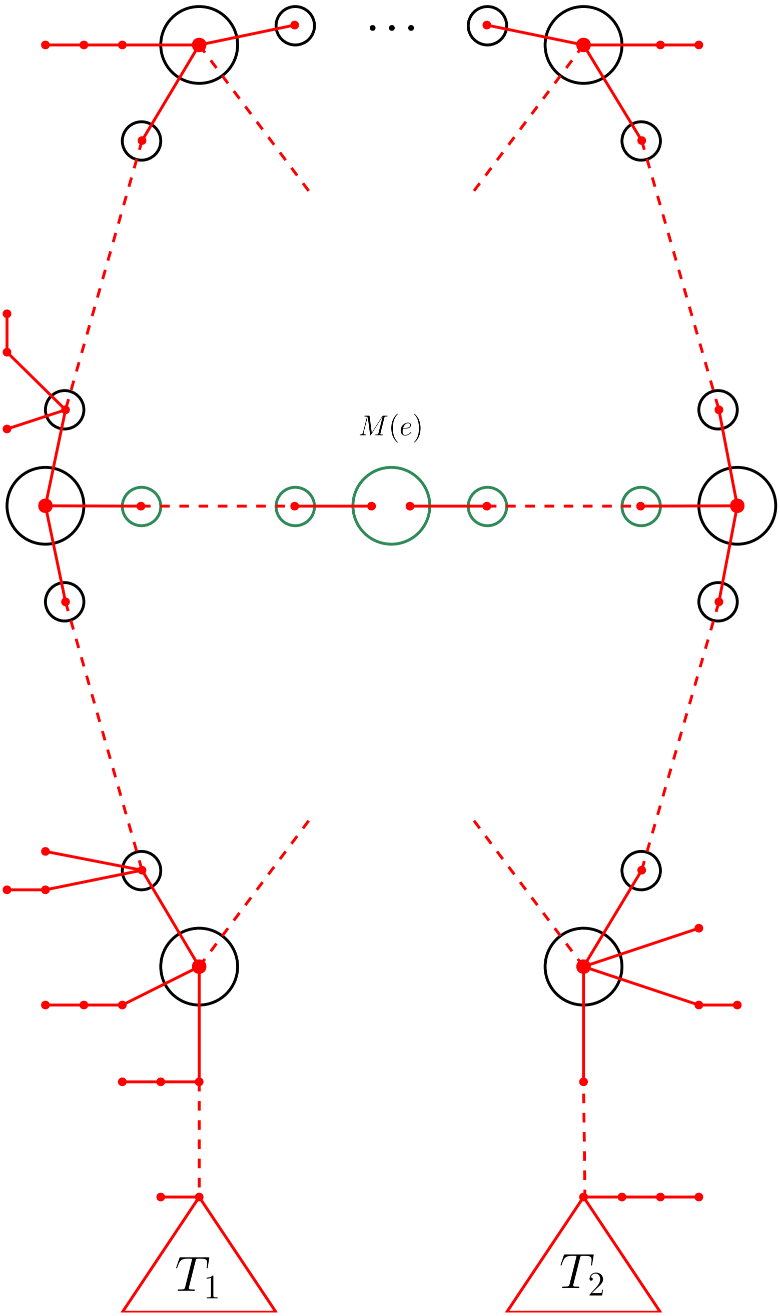

Fig. 5 Illustration of Lemma 43: The only possibility for an edge-colourful copy of T to be embedded in $\hat{G}$ is depicted in red (Color figure online)

To conclude the proof, we now prove the claim. Note first that P cannot pass through any of the colour classes V_i more than once as this would cause T_{col} to use an edge-colour multiple times. Next assume for contradiction that P misses some colour class V_a for some $a \in [2, k-1]$ (i.e., we assume that $t < k$). Since T_{col} is a connected tree containing all of the edge colours in Q there must be an index $j_i \neq a$ and a vertex $u \in V_{j_i} \cap P$ such that T_{col} contains a (unique) path P_u from u to a vertex $w \in V_a$. In order to get the contradiction, root T_{col} at u. Construct a subtree $T_{\mathsf{col}}(u)$ of T_{col} as follows: For each neighbour x of u except the ancestor of w on the path from u, we delete x and all of its descendants. Observe that the edge colours of $T_{\mathsf{col}}(u)$ are disjoint from the edge-colours of P and that $V(T_{\mathsf{col}}(u))$ is disjoint from S. Now, if $T_{\mathsf{col}}(u)$ is a path, then (using that $\ell_i > 2d$), we obtain that u is the source of a ray in T_{col} of length greater than d, contradicting the fact that every ray of length $> d$ of T_{col} is in the subgraph of $\hat{G}$ induced by S. Otherwise, $T_{\mathsf{col}}(u)$ contains a fork, contradicting the fact that all forks of T_{col} are in the subgraph of $\hat{G}$ induced by S.

Having established that $t = k$ and that no V_i is visited more than once, it remains to show that P visits the colour classes in the correct order, that is $V_{j_i} = V_i$ for each $i \in [k]$. Assume for contradiction that this is not the case, which allows us to set

$$m := \min\{i \in [k] \mid V_{j_i} \neq V_i\} - 1.$$

Note that $m \geq 1$ since $j_1 = 1$. Let $z_m \in V_m \cap P$ and $z_{m+1} \in V_{m+1} \cap P$ and recall that G contains colour classes $U_m^1, \dots, U_m^{\ell_m - 1}$ corresponding to the path

$$P_m = v_m, u_m^1, \dots, u_m^{\ell_m - 1}, v_{m+1}$$

in Q (see Definition 37). Let us now define the subtrees $T_{\mathsf{col}}(m)$ and $T_{\mathsf{col}}(m+1)$:

- For $T_{\mathsf{col}}(m)$ we root T_{col} at z_m and for each neighbour x of z_m in T_{col}, we delete x and all of its descendants unless $x \in U_m^1$.
- For $T_{\mathsf{col}}(m+1)$ we root T_{col} at z_{m+1} and for each neighbour x of z_{m+1} in T_{col}, we delete x and all of its descendants unless $x \in U_m^{\ell_m - 1}$.

Note that at least one of $T_{\mathsf{col}}(m)$ and $T_{\mathsf{col}}(m+1)$ must have depth greater than d (if rooted at z_m and z_{m+1}, respectively), since $\ell_m > 2d$ and T_{col} is edge-colourful with respect to $\hat{\gamma}$, that is, we have to make sure that we cover all of the edge colours

$$\{v_m, u_m^1\}, \{u_m^1, u_m^2\}, \dots, \{u_m^{\ell_m - 1}, v_{m+1}\}$$

Finally, regardless of which one of the two subtrees has depth greater than d, we will find either a fork, or the source of a ray of length greater than d outside of the set S, yielding the desired contradiction and concluding the proof of the claim, and hence the proof of the lemma. □

We are now able to prove the main lemma of this subsection.

Lemma 44 $\oplus\mathsf{Emb}((Q \# \tau, c_\tau) \to (G, c)) = \oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$.

Proof We start with the following claim from [31].

Claim A colour-preserving embedding $\varphi \in \mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c))$ is uniquely defined by its image (which is a subgraph of (G, c)).

For convenience, we provide a proof of the claim: Consider in image (G', c') of φ where G' is a subgraph of G and $c' = c\mid_{V(G')}$. Let $e = \{u, v\}$ be an edge of G' Then $c'(e) = \{c(u), c(v)\}$ is an edge of Q since c is a Q-colouring. Recall that $Q\sharp\tau$ is Q-coloured by the function c_τ that maps w^B to w for each $w \in V(Q)$ and block $B \in \tau_w$. Now recall the definition of fractured graphs (Definition 8) and let B_1 and B_2 be the blocks of $\tau_{c(u)}$ and $\tau_{c(v)}$ that contain $c(e)$. Then, since φ is an embedding, it maps $c(u)^{B_1}$ to u and $c(v)^{B_2}$ to v. Since Q does not have isolated vertices, continuing this process over all edges of G' defines φ. This concludes the proof of the claim. □

By the claim, it is sufficient to construct a bijection b from elements in $\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ to subgraphs (G', c') that are images of embeddings in $\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c))$. Given $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ we set $b(T_{\mathsf{col}}) := (T_{\mathsf{col}}[G], c(T_{\mathsf{col}}))$ where $c(T_{\mathsf{col}})$ is the colouring of vertices of $T_{\mathsf{col}}[G]$ which agrees with $\hat{\gamma}$ on the edges of $T_{\mathsf{col}}[G]$. In the rest of the proof, we show that b is the desired bijection.

First, we have to show that for all T_{col}, $(T_{\mathsf{col}}[G], c(T_{\mathsf{col}}))$ is the image of an embedding in $\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c))$. To this end, recall that $T_{\mathsf{col}}[G]$ induces a fracture $\rho = \rho(T_{\mathsf{col}})$ of Q. By the definition of ρ, $T_{\mathsf{col}}[G]$ and $Q\sharp\rho$ are isomorphic and this isomorphism preserves the colours so c_ρ agrees with $\hat{\gamma}$ on the edges of $Q\sharp\rho$. This implies that c_ρ and $c(T_{\mathsf{col}})$ are the same. So $(T_{\mathsf{col}}[G], c(T_{\mathsf{col}}))$ is the image of an embedding in $\mathsf{Emb}((Q\sharp\rho, c_\rho) \to (G, c))$. Finally, Lemma 43 guarantees that $\rho = \tau$.

Second, we will show that b is injective. To this end, let $T_{\mathsf{col}1} \neq T_{\mathsf{col}2} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$. Since $T_{\mathsf{col}1}$ and $T_{\mathsf{col}2}$ must both fully contain $V(T)\backslash C(H)$, and since both are edge-colourful (see Definition 42), the only possibility for $T_{\mathsf{col}1}$ and $T_{\mathsf{col}2}$ not being equal is that they disagree on G, that is, $T_{\mathsf{col}1}[G] \neq T_{\mathsf{col}2}[G]$. This proves b to be injective.

Finally, we will show that b is surjective: Given any (G', c') that is the image of an embedding $\varphi \in \mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c))$, we construct $T_{\mathsf{col}}(G', c') \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ with $b(T_{\mathsf{col}}(G', c')) = (G', c')$ as follows. Observe first that G' is isomorphic to $T[C(H)]$ since $Q\sharp\tau$ is, by definition of τ, isomorphic to $T[C(H)]$: Splitting the inner paths of length $2d$ in Q at their central vertices yields precisely $T[C(H)]$. Then $T_{\mathsf{col}}(G', c')$ is obtained by adding the remainder of T to (G', c'):

1. We add to (G', c') all vertices in $V(T) \setminus C(H)$ (see (B) in Definition 42).
2. We add all edges between vertices in $V(T) \setminus C(H)$ that are present in $\hat{G}$ (see (C) in Definition 42).
3. Finally, we connect a vertex in z in $V(T)\backslash C(H)$ with a vertex w in G' if and only if z and w are connected in $\hat{G}$ (see (D) in Definition 42).

The resulting subgraph $T_{\mathsf{col}}(G', c')$ of $\hat{G}$ is clearly edge-colourful and isomorphic to T, concluding the proof. □

We are now able to establish hardness of $\oplus$Sub$(\mathcal{T})$ in case of unbounded C-number.

Lemma 45 *Let $\mathcal{T}$ be a recursively enumerable class of trees of unbounded* C*-number. Then* $\oplus$Sub$(\mathcal{T})$ *is* $\oplus$W[1]*-hard.*

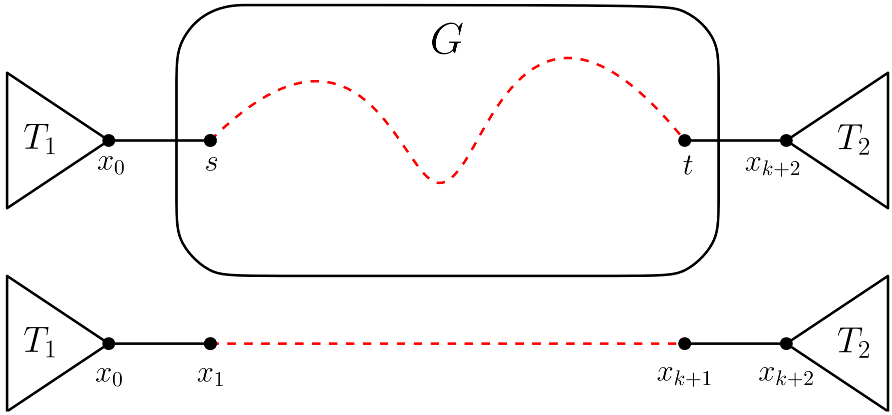

Fig. 6 Reduction from counting s-t-paths of length k, modulo 2, in a graph G to counting copies of a tree T with a 2-path of length at least $k+2$

Proof Assume first that $\mathcal{T}$ contains trees with 2-paths of unbounded length. In this case we reduce from the problem of counting k-cycles, modulo 2, which was shown $\oplus$W[1]-hard in [12]. In the first step, this problem reduces to the problem of counting s-t-paths of length k, modulo 2 as shown in Lemma 5.2 in the full version [29] of [28]. In the second and final step, we can easily reduce from the problem of counting s-t-paths of length k, modulo 2, to $\oplus\textsc{Sub}(\mathcal{T})$, as shown in Fig. 6: Concretely, let (G, s, t, k) be a problem instance. Since $\mathcal{T}$ contains trees with 2-paths of unbounded length, we can find, in time only depending on k, a tree T in $\mathcal{T}$ containing a 2-path $x_0, x_1, \ldots, x_{k+1}, x_{k+2}$ of length $k+2$. Let furthermore T_1 and T_2 be the subtrees of T as depicted in Fig. 6. We construct a graph G' from G in two steps as follows: First, we add fresh vertices x_0 and x_{k+2} and edges $\{x_0, s\}$ and $\{t, x_{k+2}\}$. Second, we add T_1 and T_2 and identify their roots with x_0 and x_{k+2}, respectively. The construction is depicted in Fig. 6 as well. Now let A be the set of subgraphs of G' that are isomorphic to T and that contain all edges of T_1 and T_2. It is easy to see that the cardinality of A is equal to the number of s-t-paths of length k in G. Thus it suffices to compute $|A|$ mod 2, using an oracle for $\oplus\textsc{Sub}(\mathcal{T})$. This can be achieved by a simple application of the inclusion–exclusion principle: Setting $S = E(T_1) \cup E(T_2)$, we have

$$|A| = \sum_{J \subseteq S} (-1)^{|J|} \cdot \#\mathsf{Sub}(T \to G' \setminus J), \tag{7}$$

where $G' \setminus J$ is the graph obtained from G' by deleting all edges in J. We can conclude the reduction by observing that the number of terms in (7) only depends on T and thus on k, and that our oracle to $\oplus\textsc{Sub}(\mathcal{T})$ allows us to evaluate (7) modulo 2.

For the remainder of the proof we can thus assume that the length of any 2-path in any tree in $\mathcal{T}$ is bounded by a constant d. Since $\mathcal{T}$ has unbounded C-number, we obtain that the trees in $\mathcal{T}$ contain C-gadgets of order d of unbounded length. By Corollary 35 we obtain that for any positive integer k, there is a value d' in the range $1 \leq d' \leq d$

such that there is a tree T_k in $\mathcal{T}$ which contains a strong C-gadget of order d' with k junctions.

Let $\mathcal{C}$ be a class of cubic Hamiltonian graphs of unbounded treewidth. Assume w.l.g. that, for each k, the class $\mathcal{C}$ contains at most one graph with k vertices; otherwise we just keep one k-vertex graph with the largest treewidth among all k-vertex graphs in $\mathcal{C}$. For each $\Delta \in \mathcal{C}$ set $T_\Delta := T_{|V(\Delta)|}$, that is T_Δ is contained in $\mathcal{T}$ and contains a strong C-gadget H_Δ with at least $|V(\Delta)|$ junctions. Recall Definition 37 and set

$$\mathcal{Q} := \{Q(\Delta, T_\Delta, H_\Delta) \mid \Delta \in \mathcal{C}\}.$$

Observe that $Q(\Delta, T_\Delta, H_\Delta)$ contains as minor the graph obtained from Δ by removing one edge. Since the removal of a single edge can decrease the treewidth only by a constant, and since treewidth is minor-monotone, we have that $\mathcal{Q}$ has unbounded treewidth.

By Theorem 12 the problem $\oplus$CP- HOM($\mathcal{Q}$) is therefore $\oplus$W[1]-hard. Thus it suffices to show that

$$\oplus\text{CP- HOM}(\mathcal{Q}) \leq^{\text{fpt}}_{\text{T}} \oplus\text{SUB}(\mathcal{T}).$$

In the first step, we reduce the computation of $\oplus\mathsf{Hom}((Q, \mathsf{id}_Q) \to \star)$ to the computation of $\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to \star)$; here, τ is the fracture defined in Definition 37. To this end, it was shown in [30] that

$$\oplus\,\mathsf{Emb}((Q \sharp \tau, c_\tau) \to \star) = \sum_{\rho \geq \tau} \mu(\tau, \rho) \cdot \oplus\mathsf{Hom}((Q \sharp \rho, c_\rho) \to \star), \qquad (8)$$

where the relation "$\geq$" and the Möbius function μ are over the lattice of fractures. We omit introducing these objects in detail, since we only require that the coefficient of the term $\oplus\mathsf{Hom}((Q \sharp \top, c_\top) \to \star)$ (which is equal to $\oplus\mathsf{Hom}((Q, \mathsf{id}_Q) \to \star)$) in the above linear combination was shown in [30] to be equal to

$$\prod_{v \in V(Q)} (-1)^{|\tau_v|-1} \cdot (|\tau_v| - 1)!.$$

Since each partition τ_v has at most two blocks, the above term is odd. Thus, by Lemma 14, we can evaluate the term $\oplus\mathsf{Hom}((Q \sharp \top, c_\top) \to \star)$ if we can evaluate the entire linear combination, that is, if we can evaluate $\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to \star)$. It thus remains to show how we can evaluate $\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to \star)$ using our oracle for $\oplus$SUB($\mathcal{T}$).

To this end, we use Lemma 44: Given any $Q = Q(\Delta, T_\Delta, H_\Delta)$-coloured graph (G, c) for which we want to compute $\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$, we first construct $(\hat{G}, \hat{\gamma})$ as in Definition 42. Then Lemma 44 yields that

$$\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c)) = \oplus\mathsf{ColSub}(T_\Delta \to (\hat{G}, \hat{\gamma})).$$

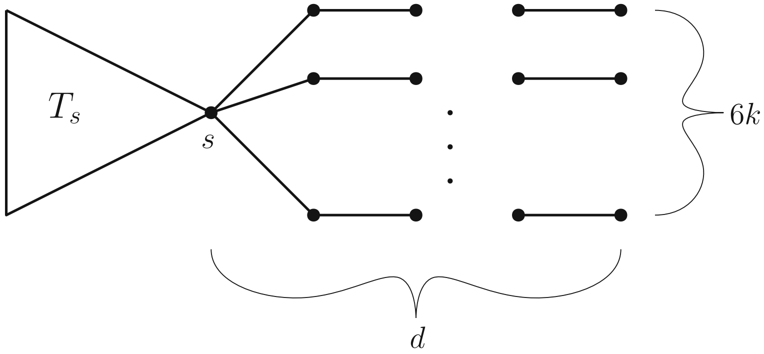

Fig. 7 A tree with $\mathsf{S}_d(T) \geq 6k$

Finally, by Lemma 15 we can compute $\oplus\mathsf{ColSub}(T_\Delta \to (\hat{G}, \hat{\gamma}))$ in FPT time using an oracle for $\oplus\mathsf{Sub}(T_\Delta \to \star)$. Since the size of T_Δ only depends on Q, and since, with input Q we can find T_Δ (recall that $\mathcal{T}$ is recursively enumerable) this yields indeed a parameterised Turing-reduction and the proof is concluded. □

4.2 Unbounded Star Number

We will use the same strategy as in Sect. 4.1: Given a tree T with large star number, we start with a properly chosen cubic graph Δ, and we construct a graph Q depending on Δ and T which contains Δ as a minor. Then we show that for any Q-coloured graph (G, c), we can construct an edge-coloured graph $(\hat{G}, \hat{\gamma})$ such that $\oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ is equal to $\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$ for a particular fracture τ.

To this end, let T be a tree with star number (at least) $6k$ for some positive integer k. By definition of the star number, there is a $d \geq 3$ such that T contains a vertex s which is the source of $6k$ rays $R_1, \ldots, R_{6k}$ of length precisely d. For each $i \in [6k]$, let $R_i = s, r_i^1, \ldots, r_i^d$. Furthermore, let T_s be the subtree of T obtained by deleting the vertices $r_i^1, \ldots, r_i^d$ for each $i \in [6k]$; consider Fig. 7 for an illustration.

Definition 46 (Q) Let Δ be cubic graph on k vertices. We obtain Q from Δ by substituting each vertex v by a gadget depicted in Fig. 8. Afterwards, we connect the gadgets as follows: If $\{v, x\}$ is an edge of Δ, then we identify the vertex v_x in the gadget of v and the vertex x_v in the gadget of x.

Observation 47 *Δ is a minor of Q.*

The fracture τ of Q that we will be interested in is defined as follows; Fig. 9 depicts the fractured graph $Q \sharp \tau$.

Definition 48 (τ) Let Q be the graph defined in Definition 46.

- For each edge $\{v, x\}$ of Δ, the graph Q contains a vertex $v_x (= x_v)$, which has degree 2. We let τ_{v_x} be the partition consisting of 2 singleton blocks.

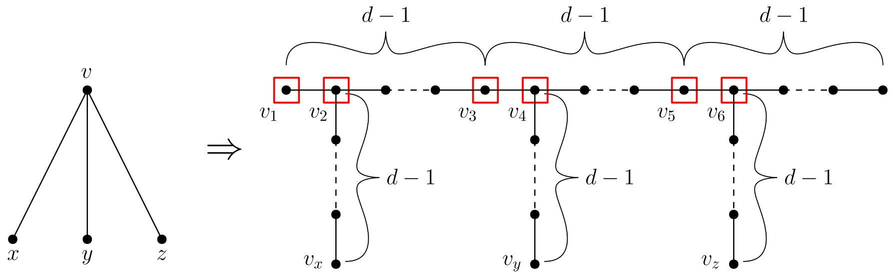

Fig. 8 The construction of Q; the vertices $v_1, \ldots, v_6$ on the gadget of v are emphasized

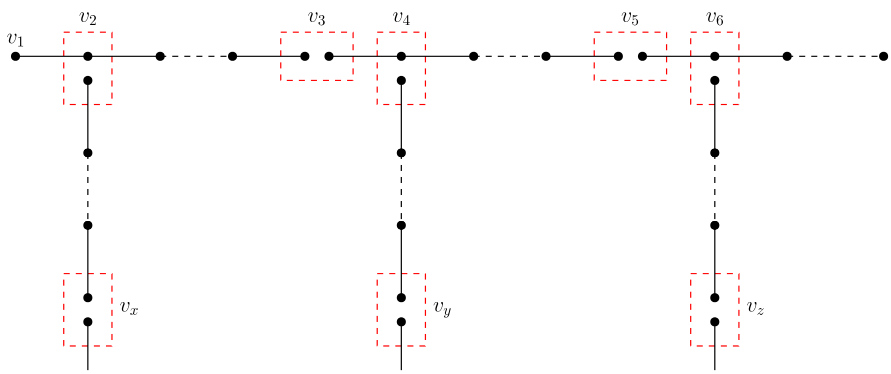

Fig. 9 Illustration of the fractured graph $Q \sharp \tau$ via fracturing the vertex gadgets

- For each vertex v of Δ, the vertices v_3 and v_5 have degree 2 in Q. We let τ_{v_3} and τ_{v_5} be the partitions consisting of 2 singleton blocks.
- For each vertex v of Δ, the vertices v_2, v_4 and v_6 have degree 3 in Q. For each $i \in \{2, 4, 5\}$ we let τ_{v_i} be the partition consisting of one block of size 2 corresponding to the edges incident to v_i from the left and the right, and one block of size 1 corresponding to the edge incident to v_i from below.

For all other vertices u of Q, we let τ_u be the partition consisting only of one block.

Analogously to the notion of a core in the case of unbounded C-number, we will identify a specific subgraph of the tree T and we will use it to define the graph $\hat{G}$ later.

Definition 49 (V') Let V' be the vertex subset of T defined as follows:

$$V' := \left(\bigcup_{i \in [6k]} V(R_i) \right) \setminus \{s\}.$$

Furthermore, we set $E' := E(T[V'])$.

Observe that $T[V']$ is a (disjoint) union of $6k$ paths of length $d-1$, where the vertices of the i-th path are $r_i^1, \ldots, r_i^d$. Observe further that $V(T) = V(T_s) \dot\cup V'$ and that

$$E(T) = E' \dot{\cup} E(T_s) \dot{\cup} \{\{s, r_i^1\} \mid i \in [6k]\}. \quad (9)$$

Next, note that the edges of Q can be decomposed into $6k$ paths, each of length $d-1$: There are k vertices of Δ. For each vertex $v \in V(\Delta)$ the graph Q contains, by definition, a gadget corresponding to v, the edges of which can be decomposed into 6 paths $P_v^1, \dots, P_v^6$ of length $d-1$ (formally, the fractured graph $Q \sharp \tau$ yields precisely this decomposition; see Fig. 9). Additionally, for each $v \in V(\Delta)$ and $i \in [6]$, the first vertex of P_v^i is chosen to be v_i as depicted in Fig. 8.

Definition 50 (γ, γ_E) We define a function $\gamma : T[V'] \to V(Q)$ as follows. Recall that $T[V']$ is the union $6k$ paths $P'_j := r_j^1, \dots, r_j^d$ for $j \in [6k]$. Fix any bijection $b : [6k] \to V(\Delta) \times [6]$. Then γ maps P'_j to P_v^i, where $b(j) = (v, i)$. In particular, we enforce that the first vertices of the paths are mapped onto each other, that is, $\gamma(r_j^1) := v^i$. Additionally, we define $\gamma_E : E' \to E(Q)$ by mapping e to $\gamma(e)$.

Observation 51 *The function γ is an edge-bijective homomorphism from $T[V']$ to Q. Specifically, γ_E is a bijection.*

Now let (G, c) be a Q-coloured graph. We state the following explicitly, since it will be crucial in our reduction.

Observation 52 *Let (G, c) be a Q-coloured graph. The mapping $c_E \circ \gamma_E^{-1}$ is a map from $E(G)$ to E'.*

Let us now construct a graph $\hat{G}$ from a Q-coloured graph G; an illustration is provided in Fig. 10.

Definition 53 $((\hat{G}, \hat{\gamma}))$ Let (G, c) be a Q-coloured graph. The graph $\hat{G}$ is an edge-coloured graph, with colouring $\hat{\gamma} : E(\hat{G}) \to E(T)$, constructed as follows:

(A) The graph $\hat{G}$ contains G as a subgraph. For each $e \in E(G)$ we set $\hat{\gamma}(e) = \gamma_E^{-1}(c_E(e))$.
(B) The vertex set of $\hat{G}$ is the union of $V(G)$ and $V(T_s)$, and pairs of vertices in $V(T_s)$ are connected by an edge in $\hat{G}$ if and only they are adjacent in T. For each such edge e, $\hat{\gamma}(e) = e$.
(C) The remaining edges of $\hat{G}$ are defined as follows. For each edge $e = \{s, r_j^1\} \in E(T)$, we connect s to all vertices in G that are coloured (by c) with $\gamma(r_j^1)$ (see Definition 50), and for each of those newly added edges e' we set $\hat{\gamma}(e') := e$

Observe that $\hat{\gamma}$ colours the edges of $\hat{G}$ with $E(T)$; the cases (A), (B), and (C) correspond, respectively, to the sets E', $E(T_s)$ and $\{\{s, r_i^1\} \mid i \in [6k]\}$ (see Eq. (9)). Similarly to the case of unbounded C-gadgets, for each element $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ the induced subgraph

$$T_{\mathsf{col}}[G] := T_{\mathsf{col}}[V(T_{\mathsf{col}}) \cap V(G)]$$

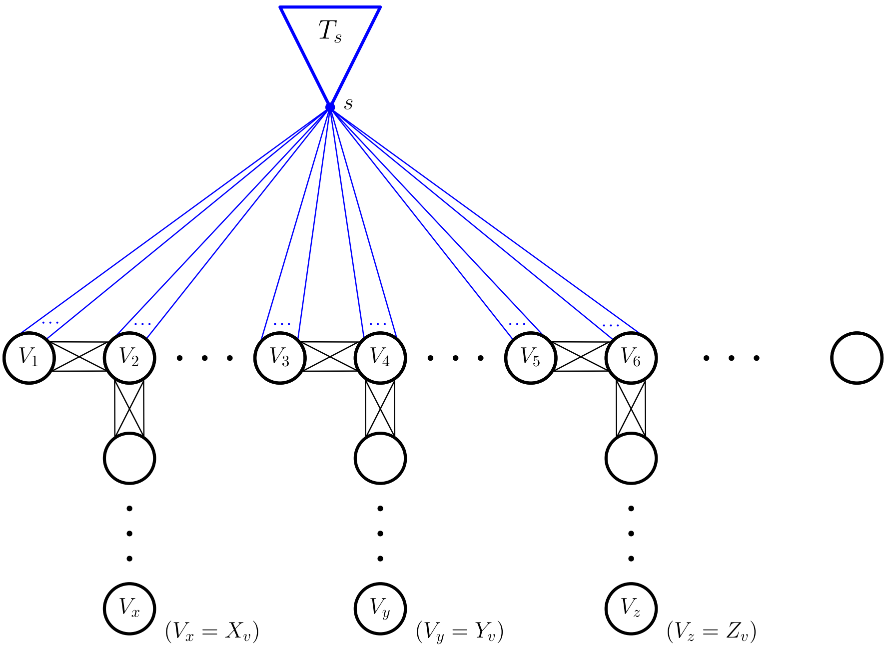

Fig. 10 The construction of $\hat{G}$. The graph G within $\hat{G}$ is depicted in black

of T_{col} is an edge-colourful subgraph in G, that is, $T_{\mathsf{col}}[G]$ contains precisely one edge per edge-colour of G under the edge colouring $\hat{\gamma}$ hence it contains precisely one edge per edge-colour of G under c_{E}. As shown in Section 3 in the full version [31] of [32], $T_{\mathsf{col}}[G]$ thus induces a fracture $\rho = \rho(T_{\mathsf{col}})$ of Q: Two edges $\{v, w\}$ and $\{v, y\}$ of Q are in the same block in the partition ρ_v corresponding to vertex v of Q if and only if the edges of $T_{\mathsf{col}}[G]$ that are coloured $\gamma_{\mathrm{E}}^{-1}(\{v, w\})$ and $\gamma_{\mathrm{E}}^{-1}(\{v, y\})$ are adjacent. In what follows, we show that ρ must always be equal to $\tau(\Delta, T, H)$ (see Definition 48).

Lemma 54 *For every $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ we have that $\rho(T_{\mathsf{col}}) = \tau$.*

Proof Let $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to \hat{G}, \hat{\gamma})$. Since T_{col} must include each of the edge colours given by $\hat{\gamma}$ (precisely) once, we have that T_{col} must fully contain T_s. Note that T_s fully contains T except for $6k$ rays of length d, and the only way to attach those rays in $\hat{G}$ is via the vertex s. Now consider the subgraph $T_{\mathsf{col}}[G + s]$ of T_{col} defined as follows:

$$T_{\mathsf{col}}[G + s] := T_{\mathsf{col}}[(V(T_{\mathsf{col}}) \cap V(G)) \cup \{s\}].$$

Since T_{col} includes all edge colours given by $\hat{\gamma}$, we have that s must have degree $6k$ in $T_{\mathsf{col}}[G + s]$: By (C) in Definition 53, the vertex s must be connected (within $T_{\mathsf{col}}[G + s]$) to one vertex in each of the colour classes $V_i = c^{-1}(v_i)$ for $v \in V(\Delta)$ and $i \in [6]$. Additionally, this implies the following: □

Observation 55 *$T_{\mathsf{col}}[G + s]$ is isomorphic to the d-stretch of $K_{1,6k}$ with s at the centre.*

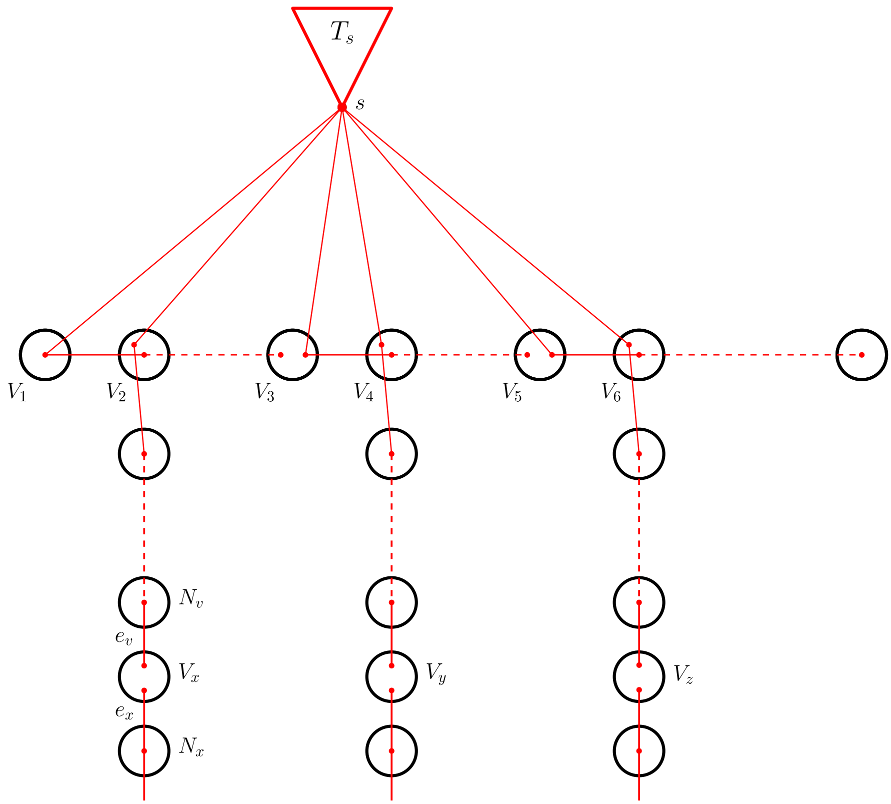

Fig. 11 Illustration of the unique way to colourfully embed T into $\hat{G}$. The induced fracture is τ

In the remainder of the proof, we will show that the only way for T_{col} to (colourfully) embed the $6k$ rays of length d is as depicted in Fig. 11. Note that this will conclude the proof since the induced fracture of the depicted embedding is τ.

Hence we proceed with proving the claim. We first consider, for each edge $\{v, x\} \in E(\Delta)$, the vertex $v_x = (x_v)$ of Q (see Definition 46 and Fig. 8). The vertex v_x has two neighbours n_v and n_x in Q, where n_v denotes the neighbour in the gadget of v and n_x denotes the neighbour in the gadget of x. Recall that we write $V_x = c^{-1}(v_x)$, $N_v = c^{-1}(n_v)$, $N_x = c^{-1}(n_x) \subseteq V(G)$ for their colour class within G (and thus within $\hat{G}$). Since T_{col} is edge-colourful, it must contain precisely one edge e_v between V_x and N_v and one edge e_x between V_x and N_x (see (A) in Definition 53). Now observe that every vertex in V_x has distance (at least) d to s within $\hat{G}$. This has two crucial consequences:

- First, the endpoints of e_v and e_x inside V_x cannot be equal: Otherwise, they could not be part of a ray of length precisely d with source s, and this would contradict the previous observation that $T_{\mathsf{col}}[G + s]$ is isomorphic to the d-stretch of $K_{1,6k}$ with s at the centre (Observation 55).
- Hence, second, the endpoints of e_v and e_x inside V_x both have degree 1. Consequently, they must be the endpoints of two of the rays of length d. However, the

only way for this to be true is them each being connected to s as depicted in Fig. 11; in all other cases, $T_{\mathsf{col}}[G+s]$ cannot be isomorphic to the d-stretch of $K_{1,6k}$ with s at the centre.

The second consequence implies that the edge colours corresponding to the edges in the paths P_v^2, P_v^4, and P_v^6 are covered for each v (recall that T_{col} must include each edge colour precisely once). Thus, the only possibility to include the remaining edge colours corresponding to the paths P_v^1, P_v^3, and P_v^5 while keeping $T_{\mathsf{col}}[G+s]$ being isomorphic to the d-stretch of $K_{1,6k}$, is to embed, for each gadget, the remaining 3 rays of length d as depicted in Fig. 11. This concludes the proof. □

We are now able to prove the main lemma of this section.

Lemma 56 $\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c)) = \oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$.

Proof Thanks to Lemma 54, the proof is similar to the proof of Lemma 44: Colour-preserving embeddings in $\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$ are uniquely identified by their image, and a bijection b from $\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ to images of colour-preserving embeddings in $\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$ is given by $b : T_{\mathsf{col}} \mapsto T_{\mathsf{col}}[G]$. □

Similarly to the proof in Sect. 4.1, Lemma 56 is sufficient for hardness.

Lemma 57 *Let $\mathcal{T}$ be a recursively class of trees of unbounded star number. Then* $\oplus$SUB$(\mathcal{T})$ *is* $\oplus$W[1]-*hard.*

Proof The proof is almost identical to the proof of Lemma 45, with the exception that we use Q, τ, $\hat{G}$, and $\hat{\gamma}$ as defined in the current section, and that we rely on Lemma 56 for the identity

$$\oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c)) = \oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma})).$$

The remainder of the proof transfers verbatim. □

4.3 Unbounded Fork Number

We will rely on the same high-level strategy as the one that we used when the C-number or star number was unbounded: Given a tree T with large a-b-fork number, we start with a properly chosen cubic graph Δ, and we construct a graph Q which depends on T and Δ, and which contains Δ as a minor. Afterwards, we show that for any Q-coloured graph (G, c) we can construct an edge-coloured graph $(\hat{G}, \hat{\gamma})$ where the co-domain of $\hat{\gamma}$ is $E(T)$ such that $\#\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ is equal (modulo 2) to $\#\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$ for a particular fracture τ of Q. However, proving this equality will be more involved than it was in the previous cases: In Sects. 4.1 and 4.2, we were able to prove, implicitly, that $\#\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma})) = \#\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$, that is, we were able to establish equality, rather than equality modulo 2. In the current case, we are not able to prove equality and must therefore rely on parity arguments, which makes the case slightly more involved. We start by fixing the following:

- Positive integers k, a and b with $a \le b$ and $k \ge 2$.

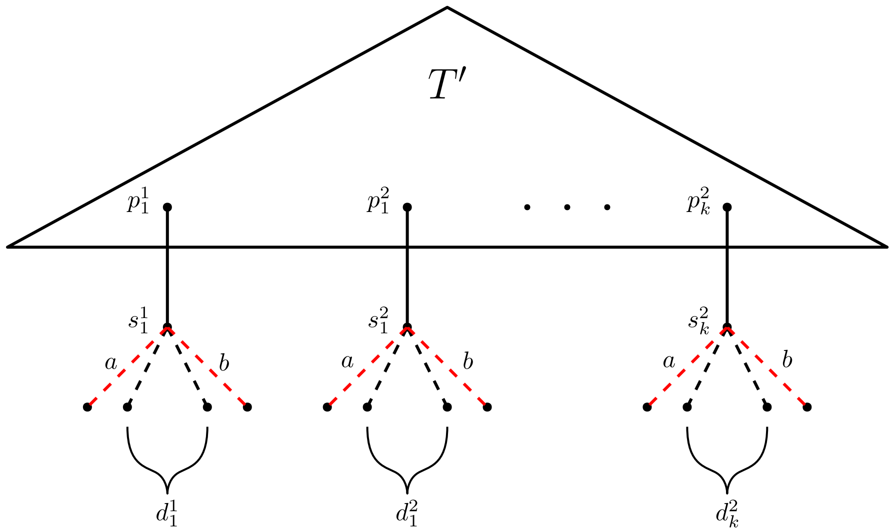

Fig. 12 A tree T with $\mathrm{F}_{a,b}(T) \geq 2k$. Note that the parents of the s_i^j are not necessarily distinct. The rays $F_a(i, j)$ and $F_b(i, j)$ are depicted in red (Color figure online)

- A tree T with $\mathrm{F}_{a,b}(T) \geq 2k$. By definition of forks (Definition 29), T contains *designated sources* $s_1^1, s_1^2, \ldots, s_k^1, s_k^2$ such that for each $(i, j) \in [k] \times [2]$, the source s_i^j is the source of two (distinct) rays $F_a(i, j)$ of length a and $F_b(i, j)$ of length b. Additionally $\mathsf{deg}_{\mathsf{NL}}(s_i^j) = 1$. We assume w.l.o.g. that the designated sources are ordered by their leaf-degrees, that is

$$\mathsf{deg}_{\mathsf{L}}(s_1^1) \geq \mathsf{deg}_{\mathsf{L}}(s_1^2) \geq \cdots \geq \mathsf{deg}_{\mathsf{L}}(s_k^1) \geq \mathsf{deg}_{\mathsf{L}}(s_k^2). \tag{10}$$

  Consider Fig. 12 for an illustration of T, its designated sources, and the rays $F_a(i, j)$ and $F_b(i, j)$.
- A k-vertex *bipartite* cubic graph Δ with vertices $V(\Delta) = \{v_1, \ldots, v_k\}$.
- A proper 3-edge-colouring $C : E(\Delta) \to \{s, m, \ell\}$ of Δ.[6]

We first note that, since there are at least $2k \geq 4$ sources in T, any pair of distinct sources must not be adjacent: Otherwise, the tree T would either be disconnected, or one of the sources would have $\mathsf{deg}_{\mathsf{NL}}$ at least 2, both of which is a contradiction.

Observation 58 *For any distinct pair $(i, j) \neq (i', j')$ we have that s_i^j and $s_{i'}^{j'}$ are not adjacent in T.*

Next, we define the graph Q.

Definition 59 (Q) The graph Q is obtained from Δ and C via substituting v_i by the gadget depicted in Fig. 13 for each $i \in [k]$. Afterwards, for every edge $e = \{v_i, v_j\}$

[6] That is, $C(e_1) \neq C(e_2)$ whenever $e_1 \neq e_2$ share a vertex. Note that every cubic bipartite graph has a 3-edge-colouring by Hall's Theorem.

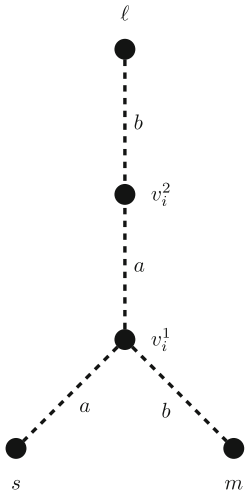

Fig. 13 A vertex gadget in the construction of Q in Definition 59. A dashed line labelled with a (resp. b) depicts a path of length a (resp. b)

of Δ we identify the vertex coloured with $C(e)$ in the gadget of v_i with the vertex coloured with $C(e)$ in the gadget of v_j.

While Definition 59 will be useful in our proofs, we note the following easier equivalent way to define Q.

Observation 60 *The graph Q is obtained from Δ and C by substituting each edge of colour s (of Δ) with a path of length $2a$, each edge of colour m with a path of length $2b$, and each edge of colour ℓ with a path of length $2(a+b)$. Consequently, Δ is a minor of Q.*

The fracture τ of Q that we will be interested in is defined as follows; Fig. 14 depicts the fractured graph $Q \sharp \tau$.

Definition 61 (τ) Let Q be the graph defined in Definition 59.

- For each edge $e = \{v_i, v_j\}$ of Δ, there is a vertex $C(e) \in \{s, m, \ell\}$ of degree 2 that connects the gadgets of v_i and v_j. We let $\tau_{C(e)}$ be the partition consisting of two singleton blocks.
- For each vertex v_i of Δ, the gadget of v_i in Q contains the vertex v_i^1 of degree 3 which is connected to s via a path of length a, to m via a path of length b, and to ℓ via a path of length $a+b$. Let e_s, e_m, and e_ℓ be the first edges on those paths. We set

$$\tau_{v_i} = \{\{e_s, e_m\}, \{e_\ell\}\}.$$

For all other vertices u of Q, we let τ_u be the partition consisting only of one block.

Next we identify specific substructures of T that will be necessary in the construction of $\hat{G}$.

Definition 62 Recall that s_i^j with $(i, j) \in [k] \times [2]$ are the designated sources of T.

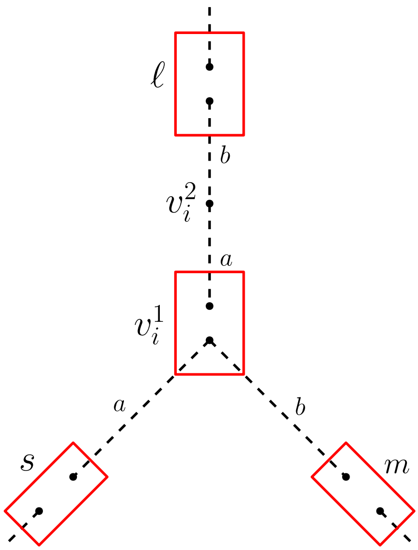

Fig. 14 The fractured graph $Q \sharp \tau$. Note that the illustration only depicts the fracturing of a single vertex gadget

- T' is the graph obtained from T by deleting, for each $(i, j) \in [k] \times [2]$, the designated source s_i^j as well as all rays with source s_i^j.
- For each $(i, j) \in [k] \times [2]$, p_i^j is the neighbour of s_i^j which is not contained in a ray. Note that p_i^j is unique by definition of forks. Note that $p_i^j \in V(T')$ and that the p_i^j are not necessarily pairwise distinct.
- For each $(i, j) \in [k] \times [2]$, $d_i^j = \mathsf{deg}_{\mathsf{L}}(s_i^j) - 2$, that is, d_i^j is the number of rays with source s_i^j minus 2. Note that $d_i^j \geq 0$ since each s_i^j is the source of $F_a(i, j)$ and $F_b(i, j)$.
- $F := \bigcup_{(i,j)\in[k]\times[2]} (F_a(i, j) \cup F_b(i, j))$, that is, F is the subset of $V(T)$ that contains the vertices of the rays $F_a(i, j)$ and $F_b(i, j)$ (which includes s_i^j) for each $(i, j) \in [k] \times [2]$.
- $E' := E(T[F])$.

An illustration of these notions is given in Fig. 12.

Observe that $T[F]$ is a disjoint union of $2k$ paths of length $a + b$. Specifically, for each $(i, j) \in [k] \times [2]$ it contains the path

$$F_i^j := T[F_a(i, j) \cup F_b(i, j)].$$

It turns out that Q is isomorphic to a quotient graph of $T[F]$, since for each vertex v_i of Δ, the vertex gadget of v_i decomposes into two paths of length $a + b$. In fact, this decomposition is given by the fractured graph $Q \sharp \tau$ (see Fig. 14). Formally, we have the following:

Observation 63 $T[F] \cong Q \sharp \tau \cong 2k P_{a+b}$.

Similarly to the previous two cases, we introduce functions γ and γ_{E} which we will need for defining the edge-colours of $\hat{G}$.

Definition 64 ($\gamma, \gamma_{\mathrm{E}}$) We define a function $\gamma : F \to V(Q)$ as follows:

1. For each $i \in [k]$, γ maps F_i^1 to the $(a+b)$-path in the gadget of v_i from s to m, such that $\gamma(s_i^1) = v_i^1$.
2. For each $i \in [k]$, γ maps F_i^2 to the $(a+b)$-path in the gadget of v_i from v_i^1 to ℓ, such that $\gamma(s_i^2) = v_i^2$.

Furthermore, we write $\gamma_{\mathrm{E}} : E' \to E(Q)$ by setting $\gamma_{\mathrm{E}}(\{x, y\}) := \{\gamma(x), \gamma(y)\}$.

Note that the definition of γ_{E} is well-defined since γ is a homomorphism by Observation 63. Concretely, γ can be viewed as the composition of an isomorphism from $T[F]$ to $Q \sharp \tau$ and the Q-colouring c_τ of $Q \sharp \tau$ (see Definition 9). Furthermore, γ_{E} is clearly a bijection. Hence, similarly to the previous sections, we point out the following:

Observation 65 *Let (G, c) be a Q-coloured graph. The mapping $c_E \circ \gamma_E^{-1}$ is a map from $E(G)$ to E'.*

We are now able construct a graph $\hat{G}$ from a Q-coloured graph G; an illustration is provided in Fig. 15.

Definition 66 ($(\hat{G}, \hat{\gamma})$) Let (G, c) be a Q-coloured graph. The pair $(\hat{G}, \hat{\gamma})$ is an edge-coloured graph constructed as follows, where the co-domain of $\hat{\gamma}$ is $E(T)$.

(A) The graph $\hat{G}$ contains G as a subgraph. For each $e \in E(G)$, define $\hat{\gamma}(e) = \gamma_{\mathrm{E}}^{-1}(c_E(e))$.
(B) The vertex set of $\hat{G}$ is the union of $V(G)$ and $V(T) \setminus F$.
(C) Pairs of vertices in $V(T) \setminus F$ are connected by an edge in $\hat{G}$ if and only if they are adjacent in T. For each such edge e, we set $\hat{\gamma}(e) = e$.
(D) The remaining edges of $\hat{G}$ are defined as follows. For each edge $e \in E(T)$ that connects a vertex $z \in V(T) \setminus F$ to a vertex $y \in F$ there are corresponding edges in $\hat{G}$. These edges connect z to all vertices $g \in V(G)$ such that $c(g) = \gamma(y)$ For each such edge e' in $\hat{G}$, $\hat{\gamma}(e') = e$.

In (D), the only edges in T connecting $z \in V(T) \setminus F$ to a vertex $y \in F$ satisfy that y is one of the designated sources s_i^j, and z is either $p_i^j \in V(T')$ or z is contained in one of the d_i^j rays with source s_i^j that are not $F_a(i, j)$ or $F_b(i, j)$ (see Definition 62).

Similarly to the other cases, for each element $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ the induced subgraph $T_{\mathsf{col}}[G] := T_{\mathsf{col}}[V(T_{\mathsf{col}}) \cap V(G)]$ of T_{col} is an edge-colourful subgraph in G. Also, $T_{\mathsf{col}}[G]$ induces a fracture $\rho = \rho(T_{\mathsf{col}})$ of Q as follows. First, recall that G is Q-coloured by c, and that G is contained in $\hat{G}$ (see (A) in Definition 66). Next note that $T_{\mathsf{col}}[G]$ is a subgraph of G that contains each edge colour in the image of $c_E \circ \gamma_{\mathrm{E}}^{-1}$ precisely once. Since γ_{E} is a bijection from E' to $E(Q)$, we can thus equivalently view $T_{\mathsf{col}}[G]$ as a subgraph of G that contains each edge colour in the image of c_E precisely once. This fact allows us to define $\rho(\mathcal{T})$ in terms of the function c_E as follows.

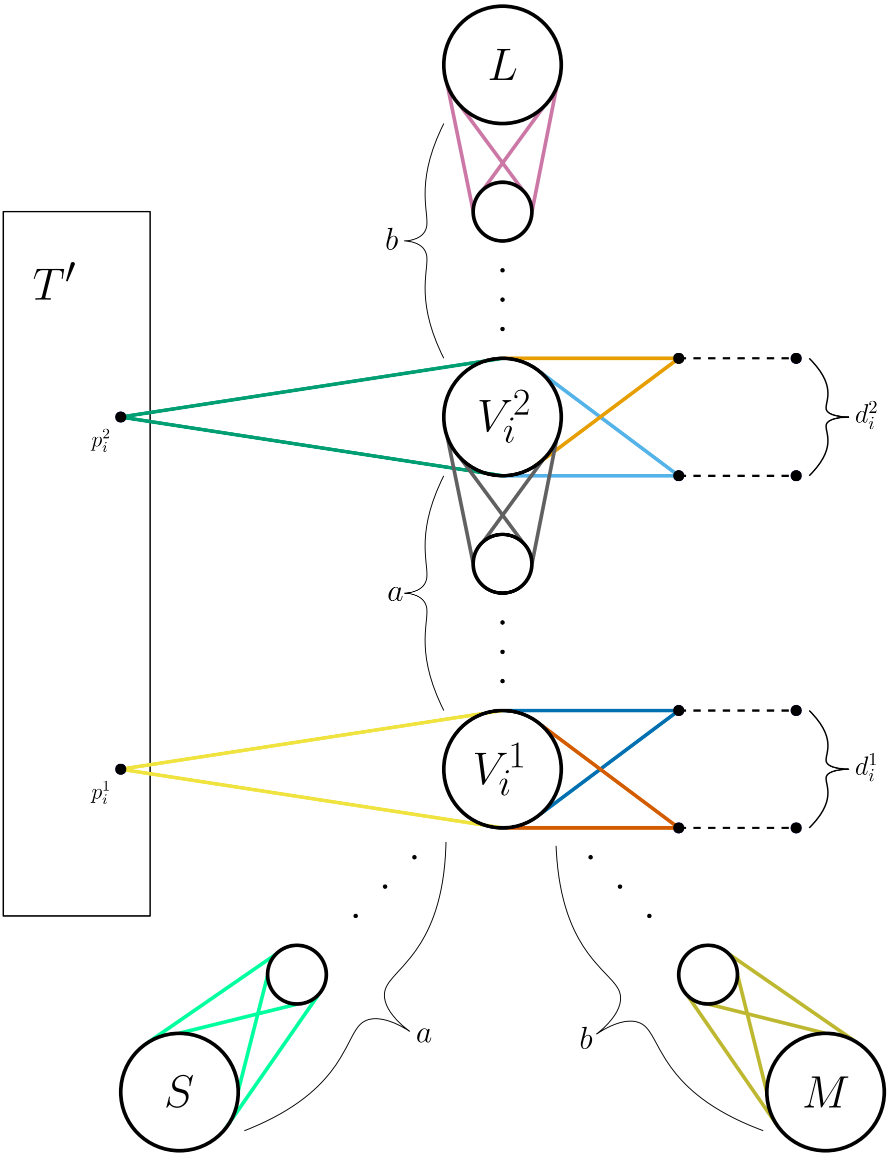

Fig. 15 The graph $\hat{G}$. Depicted in the centre is the part of G (within $\hat{G}$) that is coloured with the vertices of the i-th vertex gadget of Q. Depicted in black are the subtree T' of T (left), and, as dashed lines, the inner edges of the $d_i^1 + d_i^2$ rays incident to s_i^1 and s_i^2 (right)—here, the *inner* edges are those that are not incident to the sources s_i^1 and s_i^2. Each edge of $\hat{G}$ fully contained in the black part has a unique colour w.r.t. $\hat{\gamma}$ (see Definition 66 (C)). Pairs consisting of remaining edges have the same colour (w.r.t. $\hat{\gamma}$) if and only if they are depicted with the same colour (Color figure online)

Definition 67 ($\rho(T_{\mathsf{col}})$) Let T_{col} be an element of $\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$. The fracture $\rho = \rho(T_{\mathsf{col}})$ of Q is defined as follows. Two edges $\{v, w\}$ and $\{v, y\}$ of Q are in the same block in the partition ρ_v corresponding to vertex v of Q if and only if the edges of $T_{\mathsf{col}}[G]$ that are coloured by c_E with $\{v, w\}$ and $\{v, y\}$ are incident.

With $(\hat{G}, \hat{\gamma})$ defined, we can finally state formally the goal of this section. Recall that (G, c) is a Q-coloured graph.

Lemma 68 *Suppose that $|c^{-1}(v)|$ is odd for each $v \in V(Q)$. Then* $\oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma})) = \oplus\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c))$.

The proof requires some additional set-up. In particular, we need the condition that $|c^{-1}(v)|$ is odd to deal with the case in which what we call "invalid trees" arise. To this end, recall that $V_i^j = c^{-1}(v_i^j)$ denotes the set of vertices in G that are coloured by c with v_i^j. Since G is a subgraph of $\hat{G}$ (see Definition 66), we slightly abuse notation and write V_i^j also for the subset of vertices in $\hat{G}$ corresponding to V_i^j in G.

Definition 69 Let $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ and let $(i, j) \in [k] \times [2]$. We call T_{col} *invalid at* (i, j) if the following two conditions are met:

(I) T_{col} contains precisely two vertices x and y in V_i^j.

(II) x is adjacent to p_i^j and not incident in T_{col} to any edge coloured with a colour in E' (see Definition 66 (A)).

Otherwise T_{col} is called *valid* at (i, j). We call T_{col} an *invalid tree* if there exists a pair $(i, j) \in [k] \times [2]$ such that T_{col} is invalid at (i, j). Otherwise, we call T_{col} a *valid tree*. We write $\mathsf{ColSub}_{\mathsf{val}}(T \to (\widehat{G}, \hat{\gamma}))$ for the set of all valid T_{col} in $\mathsf{ColSub}(T \to (\widehat{G}, \widehat{\gamma}))$.

Consider Fig. 16 for an illustration of Definition 69.

Lemma 70 *Suppose that $|c^{-1}(v)|$ is odd for each $v \in V(Q)$. Then the number of invalid trees $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ is even.*

Proof For the proof, given two tuples (i, j) and (i', j') in $[k] \times [2]$ we write $(i', j') < (i, j)$ if (i', j') is lexicographically smaller than (i, j). Write $\mathcal{T}(i, j)$ for the set of all $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ that are invalid at (i, j) but valid on all pairs $(i', j') < (i, j)$. We will prove that $\mathcal{T}(i, j)$ is even for all $(i, j) \in [k] \times [2]$; this is sufficient for the lemma to hold.

Hence fix (i, j), let $T_{\mathsf{col}} \in \mathcal{T}(i, j)$, and let x and y be as in Definition 69. Since $V_i^j = c^{-1}(v_i^j)$ and for $j \in [2]$, v_i^j is a vertex of Q, the assumption in the statement of the lemma implies that $|V_i^j|$ is odd. Since x and y are distinct vertices in V_i^j, V_i^j contains additional vertices other than x and y. Fix a vertex $x' \in V_i^j \setminus \{x, y\}$. Obtain T'_{col} from T_{col} by deleting x (including edges incident to x) and by adding x' and the edge $\{x', u\}$ for every u that was adjacent to x — this is well-defined since x is not incident to any edge coloured with a colour in E', and by construction of $\hat{G}$ (see Definition 66 (C) and (D)) whenever $\{x, u\} \in E(\hat{G})$ is an edge not coloured with a colour in E', then $\{x', u\} \in E(\hat{G})$ for every $x' \in V_i^j$. Additionally, $\{x, u\}$ and $\{x', u\}$

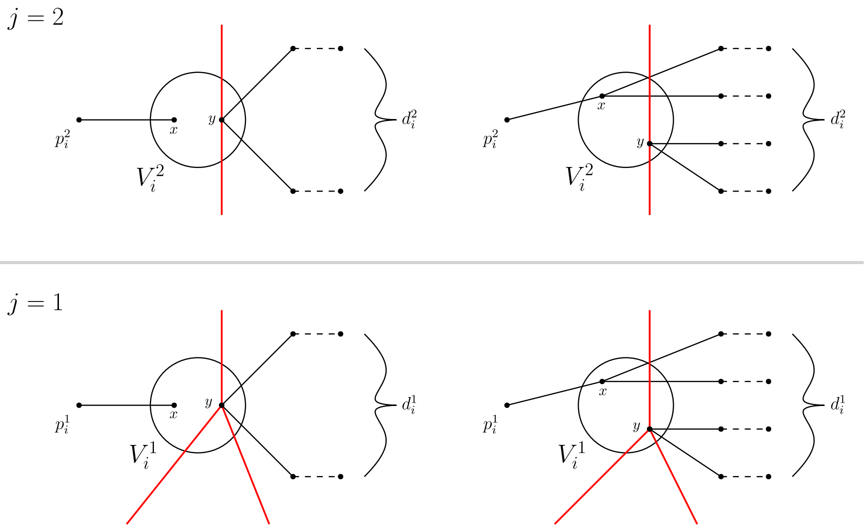

Fig. 16 Illustration of the condition that yields invalid trees at $(i, 1)$ (below) and $(i, 2)$ (above). Edges contained in E' are coloured red (Color figure online)

have the same edge-colour. Hence, clearly, T'_{col} an edge-colourful subgraph of $\hat{G}$ that is isomorphic to T_{col} (and thus to T). For this reason, we obtain that $T'_{\mathsf{col}} \in \mathcal{T}(i, j)$.

More generally, the observation that $T'_{\mathsf{col}} \in \mathcal{T}(i, j)$ allows us to define an equivalence relation on $\mathcal{T}(i, j)$: Let T_{col} and T'_{col} be elements of $\mathcal{T}(i, j)$, and let x and x' be the vertices in T_{col} and T'_{col} that satisfy (II) in Definition 69. We set T_{col} and T'_{col} to be equivalent if and only if one can obtained from the other by switching x with x' as defined above. The size of one equivalence class is precisely $|V_i^j| - 1 = |c^{-1}(v_i^j)| - 1$, which is even by the premise of the lemma. □

For the proof of Lemma 68, we need to establish some facts about rays and 2-paths of elements $T_{\mathsf{col}} \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$, which are those $T_{\mathsf{col}} \in \mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma}))$ that are valid. We encapsulate these facts in the next section.

4.3.1 The Proof of Lemma 68

We first note that, thanks to Lemma 70, it suffices to prove that

$$\#\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma})) = \#\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c)).$$

This requires some preparation. We first fix the following objects (recall the definitions of 2-path, Definition 27 and ray, Definition 28).

- T_{col} is an element of $\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$
- $T_{\mathsf{col}}[G]$ is the graph obtained from $T_{\mathsf{col}}[V(T_{\mathsf{col}}) \cap V(G)]$ with isolated vertices removed. (In fact, our proof will show that, for valid trees $T_{\mathsf{col}} \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$, the induced subgraph $T_{\mathsf{col}}[V(T_{\mathsf{col}}) \cap V(G)]$ cannot have

isolated vertices. However, at the current point of the proof, it is easiest to just remove them.)
- For any positive integer t, $\mathcal{R}^t$ is the set of length-t rays in T. $\mathcal{P}^t$ is the set of length-t 2-paths in T that are not rays.
- For any positive integer t, $\mathcal{R}^t_{\mathrm{col}}$ is the set of length-t rays in T_{col} and $\mathcal{P}^t_{\mathrm{col}}$ is the set of 2-paths in T_{col} that are not rays. Note that $|\mathcal{R}^t| = |\mathcal{R}^t_{\mathrm{col}}|$ and $|\mathcal{P}^t| = |\mathcal{P}^t_{\mathrm{col}}|$ for all t since T and T_{col} are isomorphic.

We will also rely on the following notion of external rays and 2-paths.

Definition 71 A 2-path P of T_{col} is called *external* if the following two conditions are satisfied.

- Except for the endpoints, none of the vertices of P is contained in $V(G)$.
- P does not contain an edge of G.

Definition 71 applies whether or not P is a ray. The following lemmas establish that all 2-paths of T_{col} of length greater than b must be external.

Lemma 72 *Suppose that t is an integer that is greater than b. Suppose that, for all $t' > t$, every 2-path in $\mathcal{R}^{t'}_{\mathrm{col}} \cup \mathcal{P}^{t'}_{\mathrm{col}}$ is external. Then every 2-path in $\mathcal{R}^t_{\mathrm{col}} \cup \mathcal{P}^t_{\mathrm{col}}$ is external.*

Proof We first construct a bijection f from $\mathcal{R}^t$ to $\mathcal{R}^t_{\mathrm{col}}$. We will use this bijection to argue that every ray in $\mathcal{R}^t_{\mathrm{col}}$ is external. In order to define the bijection, consider a ray $R = r_0, r_1, \ldots, r_t$ in $\mathcal{R}^t$. Since $t > b \geq a$, R is not one of the designated rays $F_a(i, j)$ and $F_b(i, j)$. If r_0 is not among the designated sources s_i^j, then, by the construction of $\hat{G}$, R is contained in T' and thus $R \in \mathcal{R}^t_{\mathrm{col}}$. In this case R is external and we set $f(R) := R$. Alternatively, suppose that $r_0 = s_i^j$ for some i and j. Then R must be one of the d_i^j black rays in Fig. 12 (see Definition 62). By the construction of $\hat{G}$ and the fact that T_{col} is edge-colourful, there is a vertex $x \in V_i^j$ such that T_{col} contains the path $x, r_1, \ldots, r_t$. In T_{col}, as in T, the vertices $r_1, \ldots, r_{t-1}$ have degree 2 and the vertex r_t has degree 1. Vertex x cannot have degree 1 in T_{col} since this would disconnect T_{col}. Also, vertex x x cannot have degree 2: To see this, assume for contradiction that x has degree 2. Then there is an integer $t' > t$ and a ray $R' \in \mathcal{R}^{t'}_{\mathrm{col}}$ the last vertices of which are $x, r_1, \ldots, r_t$. Since x is not an endpoint of the ray and since $x \in V(G)$, the ray R' is not external, contradicting the premise of the lemma. Hence x has degree at least 3 and therefore $f(R) := x, r_1, \ldots, r_t$ is an external ray of T_{col}. The function f is injective by construction. Since T_{col} and T are isomorphic, $|\mathcal{R}^t| = |\mathcal{R}^t_{\mathrm{col}}|$ and thus f is a bijection. Since the image of f only contains external rays, we have shown that every element of $\mathcal{R}^t_{\mathrm{col}}$ is external.

Every ray in the image of f has the property that its degree-1 endpoint is not contained in $V(G)$. Since the image of f is $\mathcal{R}^t_{\mathrm{col}}$, we obtain
($*$) Every ray in $\mathcal{R}^t_{\mathrm{col}}$ has the property that its degree-1 endpoint is not contained in $V(G)$.

To complete the proof, we show that every 2-path in $\mathcal{P}^t_{\mathrm{col}}$ is external. Following the same strategy that we used before, we construct a bijection g from $\mathcal{P}^t$ to $\mathcal{P}^t_{\mathrm{col}}$. Every 2-path in the range of g is external, so we will conclude that every element of

$\mathcal{P}^t$ is external. In order to define the bijection, consider a 2-path $P = p_0, \ldots, p_t$ in $\mathcal{P}^t$. If neither of the endpoints of P is among the designated sources s_i^j, then P is contained in T' and thus $P \in \mathcal{P}^t$. In this case, P is external and we set $g(P) := P$. If exactly one endpoint of P is among the designated sources, say $p_0 = s_i^j$, then there is a vertex $x \in V_i^j$ such that $x, p_1, \ldots, p_t$ is a path in T_{col}. The vertices $p_1, \ldots, p_{t-1}$ have degree 2 in T_{col} (as in T) and the vertex p_t has degree at least 3.

If x has degree 1 in T_{col}, the ray $R = p_t, \ldots, p_1, x$ is in T_{col}, and its degree-1 endpoint x is in $V(G)$, contradicting $(*)$. Hence x cannot have degree 1 in T_{col}. Similarly, x cannot have degree 2, since this would create a 2-path longer than t in T_{col} that is not external, which contradicts the premise of the lemma. Hence x has degree at least 3, and thus $g(P) := x, p_1, \ldots, p_t$ is an external 2-path in $\mathcal{P}^t_{\mathsf{col}}$.

For the last case, suppose that both endpoints of P are among the designated sources, say $p_0 = s_i^j$ and $p_t = s_{i'}^{j'}$. Then there are x and y in, respectively, V_i^j and $V_{i'}^{j'}$ such that $x, p_1, \ldots, p_{t-1}, y$ is a path in T_{col}. Again, $p_1, \ldots, p_{t-1}$ must all have degree 2 in T_{col} as well. We show that both x and y have degree at least 3 in T_{col}: If both have degree 1, then T_{col} is disconnected. If one of them has degree 1 and the other one has degree at least 3, then we created a ray of length t whose degree-1 endpoint in in $V(G)$, contradicting $(*)$. If one has degree 1 and the other one has degree 2, then we found a ray longer than t which is not external, contradicting the premise of the lemma. If one has degree 2 and the other has degree at least 2, then there is a non-external 2-path longer than t, again contradicting the premise of the lemma. Thus, as desired, both must have degree at least 3. Therefore, $g(P) := x, p_1, \ldots, p_{t-1}, y$ is an external 2-path in $\mathcal{P}^t_{\mathsf{col}}$. The function g is injective by construction. Since T_{col} and T are isomorphic, $|\mathcal{P}^t| = |\mathcal{P}^t_{\mathsf{col}}|$ and thus g is a bijection. Since the image of g only contains external 2-paths, we have shown that every element of $\mathcal{P}^t_{\mathsf{col}}$ is external, concluding the proof. □

Lemma 73 *Suppose that t is an integer that is greater than b. Then every 2-path in $\mathcal{R}^t_{\mathsf{col}} \cup \mathcal{P}^t_{\mathsf{col}}$ is external.*

Proof Let t_{max} be the maximum integer for which $\mathcal{R}^{t_{\mathsf{max}}} \cup \mathcal{P}^{t_{\mathsf{max}}}$ is nonempty. Let Φ_t be the proposition "$t \leq b$ or every 2-path in $\mathcal{R}^t_{\mathsf{col}} \cup \mathcal{P}^t_{\mathsf{col}}$ is external".

We will show by induction on $t_{\mathsf{max}} - t$ that Φ_t holds. The base case arises when $t_{\mathsf{max}} - t = 0$, so $t = t_{\mathsf{max}}$. If $t_{\mathsf{max}} \leq b$ then Φ_t is satisfied. Otherwise, for each $t' > t$, the set $\mathcal{R}^{t'}_{\mathsf{col}} \cup \mathcal{P}^{t'}_{\mathsf{col}}$ is empty and we can invoke Lemma 72 to conclude that Φ_t holds.

For the induction step, consider t such that $t_{\mathsf{max}} - t \geq 1$. By the induction hypothesis, $\Phi_{t'}$ holds for all $t' \in \{t+1, \ldots, t_{\mathsf{max}}\}$. If $t \leq b$ then Φ_t holds. Otherwise, for all $t' > t > b$, we know from $\Phi_{t'}$ that every 2-path in $\mathcal{R}^{t'}_{\mathsf{col}} \cup \mathcal{P}^{t'}_{\mathsf{col}}$ is external. We can then apply Lemma 72 to conclude that every 2-path in $\mathcal{R}^t_{\mathsf{col}} \cup \mathcal{P}^t_{\mathsf{col}}$ is external. □

Before proceeding with the proof of Lemma 68, we provide an overview of the central steps of the proof. Recall that it suffices to prove that

$$\#\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma})) = \#\mathsf{Emb}((Q \# \tau, c_\tau) \to (G, c))$$

and that we have a fixed an element T_{col} of $\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$ and proved various properties about it.

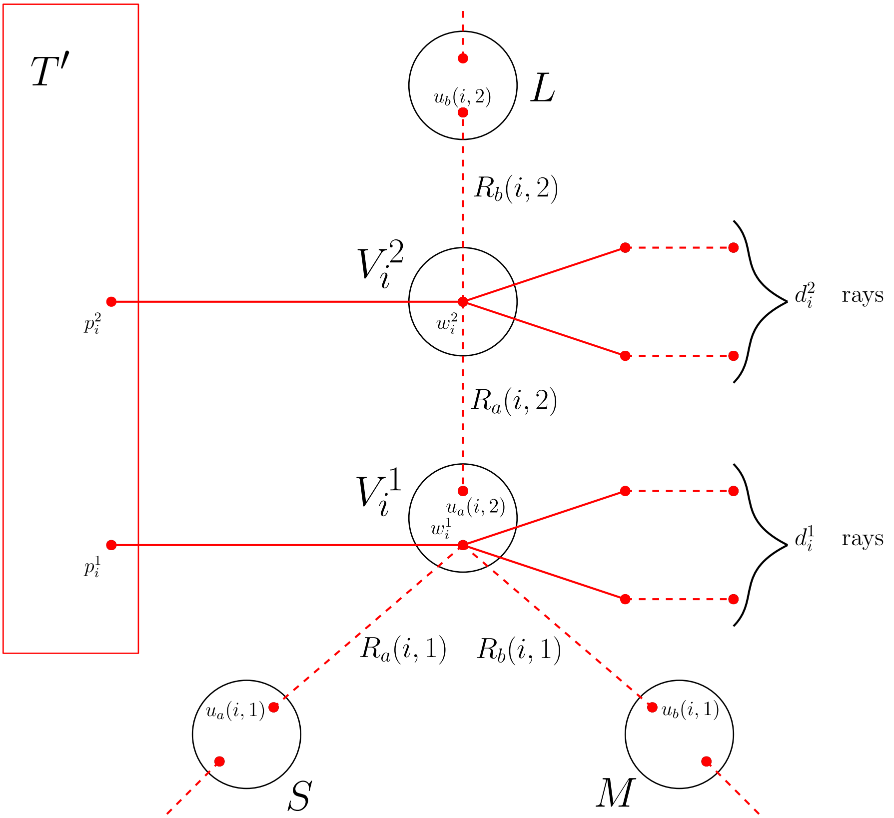

Fig. 17 An embedding T_{col} of T in $\hat{G}$ that yields the fracture τ. We will show that this is the only way to embed T in $\hat{G}$ in such a way that each edge-colour is used precisely once. Note that dashed lines depict paths in T_{col}, and solid lines depict edges in T_{col} (Color figure online)

(1) Our goal is to show that T_{col} is embedded in $\hat{G}$ in the following manner (see Fig. 17). For each $(i, j) \in [k] \times [2]$, T_{col} contains a ray $R_a(i, j)$ of length a and a ray $R_b(i, j)$ of length b; those rays correspond to the designated rays $F_a(i, j)$ and $F_b(i, j)$ in T (recall that T and T_{col} are isomorphic.)

a. T' is part of T_{col}.
b. For every $i \in [k]$ and $j \in [2]$, the vertices p_i^j in T' is connected to a vertex w_i^j of G with $c(w_i^j) = v_i^j = \gamma(s_i^j)$. In T_{col}, the vertex w_i^j is the source of d_i^j rays other than $R_a(i, j)$ and $R_b(i, j)$. The vertices of these d_i^j rays are not in T' and are not in G. The edge colours of the edges in these rays in $\hat{\gamma}$ are the same as the edge-names in T (see Definition 66 (C)).
c. The length-a ray $R_a(i, 1)$ is a path in T_{col} from w_i^1 to the vertex $u_a(i, 1)$ of G with some colour $c(u_a(i, 1))$ (a vertex of Q). This colour $c(u_a(i, 1))$ corresponds to the vertex "s" in the gadget of the vertex v_i of Δ (see Definition 59 and Fig. 13).

d. The length-b ray $R_b(i, 1)$ is a path in T_{col} from w_i^1 to the vertex $u_b(i, 1)$ of G with some colour $c(u_b(i, 1))$ (a vertex of Q). This colour $c(u_b(i, 1))$ corresponds to the vertex "m" in the gadget of the vertex v_i of Δ (see Definition 59 and Fig. 13).
e. The length-b ray $R_b(i, 2)$ is a path in T_{col} from w_i^2 to the vertex $u_b(i, 2)$ of G with some colour $c(u_b(i, 2))$ (a vertex of Q). This colour $c(u_b(i, 2))$ corresponds to the vertex "ℓ" in the gadget of the vertex v_i of Δ (see Definition 59 and Fig. 13).
f. The length-a ray $R_a(i, 2)$ is a path in T_{col} from w_i^2 to the vertex $u_a(i, 2) \neq w_i^1$ of G with some colour $c(u_a(i, 2)) = \gamma(s_i^1) = v_i^1$ (recall that the colour is a vertex of Q).
g. For every edge $e = \{v_i, v_{i'}\}$ in Δ, $u_a(i, 1) \neq u_a(i', 1)$, $u_b(i, 1) \neq u_b(i', 1)$ and $u_b(i, 2) \neq u_b(i', 2)$.

(2) We now make some observations about the fracture $\rho = \rho(T_{\mathsf{col}})$ from Definition 67, given that T_{col} is embedded in $\hat{G}$ as described in Item (1).

- The definition of Q (Definition 59) tells us that, for every edge $e = \{v_i, v_{i'}\}$ in Δ, there is a degree-2 vertex y of Q that connects the gadgets of v_i and $v_{i'}$. Vertex y corresponds to the vertex $C(e) \in \{s, m, \ell\}$ in the two gadgets. Suppose without loss of generality that $C(e) = s$. The other cases are similar. From (1c) the colour $C(e) = s$ is the same as $c(u_a(i, 1))$ and $c(u_a(i', 1))$. From (1b) $c(w_i^1) = v_i^1$ and $c(w_{i'}^1) = v_{i'}^1$. Since T_{col} is colourful and the embedding is as in (1), the edges of the ray from w_i^1 to $u_a(i, 1)$ have different edge colours to the ray from $w_{i'}^1$ to $u_a(i', 1)$. Thus, the edge in G in the first ray that is adjacent to $u_a(i, 1)$ (call it e_i) has a different colour from the edge n G in the second ray that is adjacent to $u_a(i', 1)$ (call it $e_{i'}$). Concretely, we have $c_E(e_i) = \{s, x\}$ and $c_E(e_{i'}) = \{s, x'\}$ where x and x' are the neighbours of s (in Q) in the gadgets of v_i and $v_{i'}$, respectively. By (1 g) we have $u_a(i, 1) \neq u_a(i', 1)$ and thus, by definition of ρ (Definition 67), ρ_y consists of two singleton blocks. Similar arguments show that ρ coincides with τ (see Definition 61) at every vertex of Q that corresponds to vertex "s", "ℓ" or "m" in any gadget corresponding to any vertex v_i of Δ.
- We now continue with the vertices v_i^1 for $i \in [k]$ of Q. See Fig. 13 for the gadget containing v_i^1 in Q and Fig. 17 for the graph $\hat{G}$. We will use "s", "ℓ" and "m" as the names of these vertices in the gadget containing v_i^1. The vertex v_i^1 has degree 3 and is connected to s via a path of length a, to m via a path of length b and to ℓ via a path of length $a + b$. Let y_s, y_m, and y_ℓ be the successors of v_i^1 on those paths, that is, the edges incident to v_i^1 in Q are $e_s := \{v_i^1, y_s\}$, $e_m := \{v_i^1, y_m\}$, an $e_\ell := \{v_i^1, y_\ell\}$. Now, by (1c) and (1d), the edges of T_{col} that are coloured (by c_E) with e_s and e_m are $\{w_i^1, r_a\}$ and $\{w_i^1, r_b\}$, where r_a and r_b are the successors of w_i^1 on the rays $R_a(i, 1)$ and $R_b(i, 1)$, respectively. Furthermore, by (1f), the edge of T_{col} that is coloured (by c_E) with e_ℓ is $\{u_a(i, 2), \hat{r}\}$ where $\hat{r}$ is the vertex in the ray $R_a(i, 2)$ that is adjacent to $u_a(i, 2)$. Since $u_a(i, 2) \neq w_i^1$ (by (1f)), the edge $\{u_a(i, 2), \hat{r}\}$ is

not incident to either $\{w_i^1, r_a\}$ or $\{w_i^1, r_b\}$. Thus $\rho_{v_i^1} = \{\{e_s, e_m\}, \{e_\ell\}\}$ which coincides with $\tau_{v_i^1}$ by Definition 61. So τ and ρ coincide at vertex v_i^1.

- Next are the vertices v_i^2 for $i \in [k]$ (see Fig. 13). This case is easy. If T_{col} is embedded as described in (1) (see Fig. 17), then, for each $i \in [k]$, there is only one vertex of T_{col} which is coloured by c with colour v_i^2. This vertex is w_i^2. Thus every edge of T_{col} whose edge colour includes v_i^2 is incident to w_i^2. Hence $\rho_{v_i^2}$ only consists of one block, which coincides with $\tau_{v_i^2}$ by Definition 61.
- Finally, every remaining vertex of Q (see Fig. 13) has degree 2. Let y be such a vertex and let y_1 and y_2 be the neighbours of y. Then the edges of T_{col} coloured by c_E with $\{y, y_1\}$ and $\{y, y_2\}$ must be successive edges on one of the rays $R_a(i, 1)$, $R_b(i, 1)$, $R_a(i, 2)$, or $R_b(i, 2)$. So these successive edges are both incident to the vertex of the ray that is coloured y by c. Thus ρ_y only consists of one block, which coincides with τ_y.

Since we have shown that the fractures ρ and τ coincide at every vertex of Q, we conclude that $\rho = \tau$.

(3) We next explain why it is useful to have $\rho = \tau$. Recall that our goal is to prove that $\#\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma})) = \#\mathsf{Emb}((Q\#\tau, c_\tau) \to (G, c))$ and that T_{col} is an element of $\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$. Our method will be to show that the function β defined by $\beta(T_{\mathsf{col}}) = T_{\mathsf{col}}[G]$ is a bijection from $\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$ to $\mathsf{Emb}((Q\#\tau, c_\tau) \to (G, c))$. It will turn out that this implies that the embedding ρ coincides with τ.

(4) In order to prove Item (1) we will proceed as follows.

(i) We show that all 2-paths (including rays) of T_{col} are external, *except* for $2k$ rays of length b and $2k$ rays of length a. Note that we already established this claim for 2-paths of lengths greater than b in Lemma 73.

(ii) Then we show that T_{col} contains two degree-1 vertices in each of the vertex sets L and M of G (within $\hat{G}$)—see Fig. 17, recalling that, for each vertex gadget, the sets L and M denote the vertex subsets of G that are coloured by c with ℓ and m. The point of this is that we will also prove that T_{col} has two degree-1 vertices in S (Item (iv))—this will split off the part of T_{col} corresponding to a single gadget, so we will only have to study the embedding of T_{col} within each gadget. We prove the claim about L and M by using the fact that T_{col} is isomorphic to T and that all 2-paths longer than b are external. This implies that if v_i and $v_{i'}$ are the two vertices of Δ sharing this gadget then the 2-paths between V_i^2 and $V_{i'}^2$ are covered by two rays in T_{col}, both of which end in L.

(iii) We next show that the degree-1 vertices in (ii) are the endpoints of $2k$ rays of length b. We have already seen that for each of the k gadgets the endpoints of these rays are in L and M. For the i'th gadget, the sources are in V_i^1 and V_i^2 If $b > a$ then we show that all remaining 2-paths of length b and also all 2-paths with lengths in $a + 1, \dots, b - 1$ are external. The proof of this claim relies on the same arguments as the proof of Lemma 73.

(iv) Next, we show that for each gadget, T_{col} contains two degree-1 vertices in S—see Fig. 17. The proof uses the fact that all 2-paths longer than a that are not covered by (iii) are external.

(v) We next show that the degree-1 vertices in (iv) are the endpoints of $2k$ rays of length a. We have already seen that for each of the k gadgets the endpoints of these rays are in S. For the i'th gadget, the source is in V_i^1.

(vi) The remaining details of the proof rely on the fact that the tree T_{col} is valid.

We now provide the proof in detail; for convenience, we also restate the lemma.

Lemma 68 Suppose that $|c^{-1}(v)|$ is odd for each $v \in V(Q)$. Then $\oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma})) = \oplus\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$.

Proof We will prove that for any $T_{\mathsf{col}} \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$, Item (1) of the proof overview holds.

Using this fact and the argument from Item (2) of the proof overview, we conclude that for any $T_{\mathsf{col}} \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$, $\rho(T_{\mathsf{col}}) = \tau$.

Recall that every edge-colourful subgraph of G induces a fracture of Q.

Let G' be an element of $\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$. This means that G' is an edge-colourful subgraph of G that induces τ. We wish to see how G' can be extended to some ${T_{\mathsf{col}}}' \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$. We know from Item (1) that any ${T_{\mathsf{col}}}'' \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$ can only be embedded in $\hat{G}$ in one way, so G' can only be extended in one way. The details are as follows. We claim that there is only one possible extension because T' has to be included and item (b) of (1) ensures that, for each $j \in [2]$, the vertex p_i^j is connected to w_i^j. The rest of (1) shows the unique way to include the rays, so the extension is unique.

Let β be the function from $\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$ that maps any element T_{col} to $T_{\mathsf{col}}[G]$. Note that $T_{\mathsf{col}}[G] \in \mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$ since $\rho(T_{\mathsf{col}}) = \tau$ and $\rho(T_{\mathsf{col}})$ is a function of $T_{\mathsf{col}}[G]$. Let β' be the function that maps an element of $\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))$ to its unique extension in $\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$. Note that $\beta \circ \beta'$ and $\beta' \circ \beta$ are both the identity. Therefore β is a bijection and $|\mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))| = |\mathsf{Emb}((Q \sharp \tau, c_\tau) \to (G, c))|$. The lemma follows from Lemma 70.

To finish the proof, we will fix $T_{\mathsf{col}} \in \mathsf{ColSub}_{\mathsf{val}}(T \to (\hat{G}, \hat{\gamma}))$ and we will show that Item (1) of the proof overview holds. Part (a) of (1) is trivial since T_{col} is edge-colourful so it contains T'. The first sentence of (b) is also trivial. We will next focus on (c)–(g), noting along the way when the rest of (b) is proved.

Recall from Definition 59 that, for each $i \in [k]$, the graph Q contains

- for each vertex v_j such that Δ has an edge $e = \{v_i, v_j\}$ with $C(e) = m$, a path $P_{i,j}$ of length $2b$ from v_i^1 to v_j^1, and
- for each vertex v_j such that Δ has an edge $e = \{v_i, v_j\}$ with $C(e) = \ell$, a path $P_{i,j}$ of length $2b$ from v_i^2 to v_j^2.

Recall from Definition 6 that c_E maps edges of G to edges of Q. Furthermore, G is a subgraph of $\hat{G}$, see Definition 66 (A). Let $T_{\mathsf{col}}(i, j)$ be the subgraph of $T_{\mathsf{col}}[G]$ induced by edges e of G such that $c_E(e)$ is in the path $P_{i,j}$

By construction, $T_{\mathsf{col}}(i, j)$ is the union of some number of paths. We will next argue that it is the union of exactly two disjoint length-b paths:

- If $T_{\mathsf{col}}(i, j)$ has more than two components then at least one component is disconnected from T' in T_{col}, contradicting the fact that T_{col} is a tree.

- If $T_{\mathsf{col}}(i, j)$ is a single path then it is contained in a 2-path of length at least $2b$. Since this 2-path contains an edge in G, it is not external (Definition 71). This contradicts Lemma 73.
- If $T_{\mathsf{col}}(i, j)$ is the union of exactly two disjoint paths, one of which has length larger than b then this larger 2-path is contained in a 2-path that is not external contradicting Lemma 73

What we have shown is that $T(i, j)$ consists of two length-b paths. For some $t \in \{1, 2\}$, one of these paths is from V_i^t and the other is from V_j^t. To be more precise and to fix the notation for $t = 1$, we have now shown that, for each $i \in [k]$, $T_{\mathsf{col}}[G]$ contains a path $R_b(i, 1)$ of length b that starts at a vertex $w_i^1 \in V_i^1$. We refer to the other end of this path as $u_b(i, 1)$. The vertex $u_b(i, 1)$ has degree 1 and is contained in M (i.e., in $c^{-1}(m)$). We next argue that w_i^1 has degree at least 3 in T_{col}. (See Fig. 18.)

- If w_i^1 has degree 1 in T_{col} then T_{col} is disconnected, contradicting the fact that it is a tree.
- If w_i^1 has degree 2 in T_{col}, then T_{col} has a ray of length at least $b + 1$ that is not external, which is again a contradiction.

By the same reasoning, T_{col} contains a ray $R_b(i, 2)$ of length b that starts at a vertex $w_i^2 \in V_i^2$ and ends at a vertex $u_b(i, 2)$. The ray $R_b(i, 2)$ is contained in $T_{\mathsf{col}}[G]$.

We have just finished parts (d) and (e) of (1) and the part of (g) that concerns length b. So what we have shown corresponds to Fig. 18. We would now like to prove parts (c) and (f) but unfortunately these are more difficult because we have to show where the rays with lengths between a and b are embedded so that we can argue about where the length-a rays are embedded.

Define $\hat{\mathcal{R}} := \bigcup_{i=1}^{k} \{R_b(i, 1), R_b(i, 2)\}$. Recall that k, a, and b are positive integers with $a \le b$ and $k \ge 2$ and that T has $\mathrm{F}_{a,b}(T) \ge 2k$ and $T_{\mathsf{col}} \cong T$. Also, $\mathcal{R}_{\mathsf{col}}^b$ is the set of length-b rays in T_{col} and $\mathcal{R}^b$ is the set of length-b rays in T. (See Fig. 12.) Using the notation that we have established, we will prove the following claims.

Claim 1 Let $P \in (\mathcal{R}_{\mathsf{col}}^b \backslash \hat{\mathcal{R}}) \cup \mathcal{P}_{\mathsf{col}}^b$. If $a < b$ then P is external.

We prove Claim 1 for the case where $P \in \mathcal{R}_{\mathsf{col}}^b \backslash \hat{\mathcal{R}}$. The other case is similar but easier.

Observe that $|\mathcal{R}^b| \ge 2k$ since $\mathrm{F}_{a,b}(T) \ge 2k$. So $\mathcal{R}^b$ can be partitioned as follows

- $\mathcal{R}^b[S]$ is the set of the $2k$ length-b rays $F_b(i, j)$ whose sources are $s_1^1, \ldots, s_k^2$ and which are depicted as red dashed lines in Fig. 12.
- $\mathcal{R}^b[T] = \mathcal{R}^b \setminus \mathcal{R}^b[S]$ contains the remaining rays of length b.

Our goal is to show that all rays in $\mathcal{R}_{\mathsf{col}}^b \setminus \hat{\mathcal{R}}$ are external. To do this, we first show that $|\mathcal{R}^b[T]| = |\mathcal{R}_{\mathsf{col}}^b \backslash \hat{\mathcal{R}}|$ and we then provide an injection from $\mathcal{R}^b[T]$ to $\mathcal{R}_{\mathsf{col}}^b \backslash \hat{\mathcal{R}}$ in which all elements of the range are external rays.

To show that $|\mathcal{R}^b[T]| = |\mathcal{R}_{\mathsf{col}}^b \setminus \hat{\mathcal{R}}|$, first note that $|\mathcal{R}^b| = |\mathcal{R}_{\mathsf{col}}^b|$ because T and T_{col} are isomorphic. We further have $|\mathcal{R}^b[S]| = |\hat{\mathcal{R}}| = 2k$.

We next define the (injective) map from $\mathcal{R}^b[T]$ to $\mathcal{R}_{\mathsf{col}}^b \setminus \hat{\mathcal{R}}$. For any ray $R = r_0, r_1, \ldots, r_b \in \mathcal{R}^b[T]$ we proceed as follows.

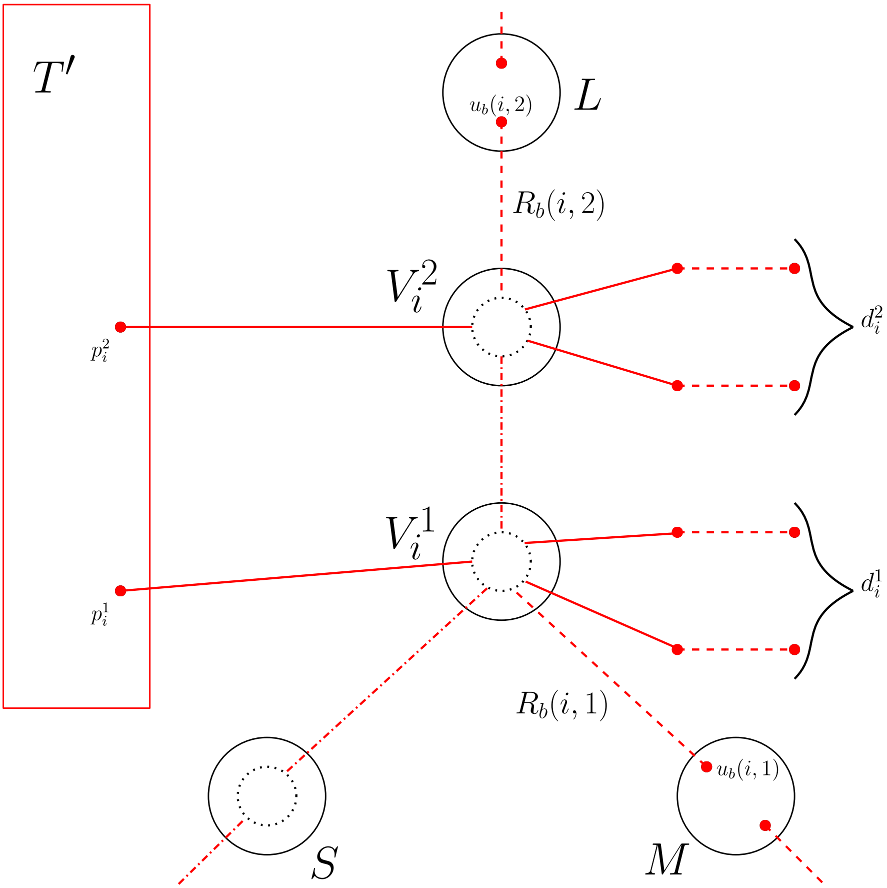

Fig. 18 Illustration of the embedding of T_{col} after the rays of length b are analysed. Solid lines depict edges, dashed lines depict paths, and dash-dotted lines depict sequences of edges (the identification of the endpoints of which we have not yet been determined). Note that both $R_b(i, 1)$ and $R_b(i, 2)$ must be of length b. Except for those two rays, the identification of endpoints of the remaining edges that are incident to G (within $\hat{G}$) has not been determined yet either; this is depicted by the dotted circles inside the colour classes. The fracture ρ induced by T_{col} will depend on the identification of the edges of T_{col}, both endpoints of which lie in G. The goal is to show that the endpoints have to be identified precisely as depicted in Fig. 17 (Color figure online)

- If r_0 is not among the designated sources s_i^j, then R is fully contained in T' (see Fig. 12) and thus R is a ray in T_{col}. We map R to itself. Note that R is external since it is fully contained in T'.
- Otherwise, $r_0 = s_i^j$ and R is one of the rays depicted as black dashed lines in Fig. 12. Since T_{col} is edge-colourful, and by construction of $\hat{G}$, T_{col} contains a path $R' = x, r_1, \ldots, r_b$ where $x \in V_i^j$. (See Fig. 15.) If x has degree 1 in T_{col} then T_{col} is disconnected, which is not true. If x has degree 2 in T_{col} then T_{col} has a non-external ray which is longer than b, which is also a contradiction by Lemma 73. Thus, x has degree at least 3 in T_{col}, and R' is an external ray. We map R to R'.

This concludes the proof of Claim 1 for the case where $P \in \mathcal{R}^b_{\mathsf{col}} \setminus \hat{\mathcal{R}}$. □

Claim 2 Suppose that there is an integer t' such that $a < t' < b$. Suppose that $P \in \mathcal{R}^{t'}_{\mathsf{col}} \cup \mathcal{P}^{t'}_{\mathsf{col}}$. Then P is external.

In order to explain the proof of Claim 2, recall that we have established the following facts about 2-paths in T_{col} in Lemma 73 and Claim 1.

- Every 2-path of length greater than b is external.
- Every 2-path of length b is either a ray in $\hat{\mathcal{R}}$ or is external.

With those 2-paths covered, the proof of Claim 2 is analogous to the proof of Lemma 73. □

Using Claims 1 and 2 we will now prove parts (c) and (f) of (1). For each 2-path whose length is larger than a, we have already shown that it is in $\hat{\mathcal{R}}$ or we have shown that it is external. In order to prove (c) we will show that, for each edge $\{v_i, v_{i'}\}$ of Δ with colour s, the sequence of edges in T_{col} between V^1_i and $V^1_{i'}$ is the union of two disjoint length-a rays. This is formalised as follows.

Note that for each edge $\{v_i, v_j\}$ of Δ coloured by the 3-edge-colouring C with s, there is a path $P_{i,j}$ of length $2a$ from v^1_i to v^1_j. Recall that c_E maps edges of G to edges of Q. We write $T_{\mathsf{col}}(i, j)$ for the subgraph of $T_{\mathsf{col}}[G]$ induced by edges e of G such that $c_E(e)$ is in the path $P_{i,j}$. By construction, $T_{\mathsf{col}}(i, j)$ is the union of some number of paths. We will next argue that it is the union of exactly two disjoint length-a paths:

- If $T_{\mathsf{col}}(i, j)$ has more than two components then at least one component is disconnected from T' in T_{col}, contradicting the fact that T_{col} is a tree.
- If $T_{\mathsf{col}}(i, j)$ is a single path then it is contained in a 2-path of length at least $2a$. Since this 2-path contains an edge in G, it is not external (Definition 71). Additionally, it is not included in $\hat{\mathcal{R}}$. This contradicts the aforementioned fact that each 2-paths of length at least $a + 1$ is external or included in the set $\hat{\mathcal{R}}$.
- If $T_{\mathsf{col}}(i, j)$ is the union of exactly two disjoint paths, one of which has length larger than a, then this larger path yields a contradiction similarly to the previous case.

What we have shown is that $T(i, j)$ consists of two length-a paths. One of these paths is from V^1_i and the other is from V^1_j. To be more precise and to fix the notation, we have now shown that, for each $i \in [k]$, $T_{\mathsf{col}}[G]$ contains a path $R_a(i, 1)$ of length a that starts at a vertex $\hat{w}^1_i \in V^1_i$. We refer to the other end of this path as $u_a(i, 1)$. The vertex $u_a(i, 1)$ has degree 1 and is contained in S (i.e., in $c^{-1}(s)$). So we have established Part (c) of item (1). Consider Fig. 19 for an illustration of all the information we gathered so far. (The vertices labelled z^j_i and the edge set E^a_i in the figure will be discussed below).

To finish the proof of item (1) we will show part (f) and the rest of part (b). We take these together. Recall that for every $i \in [k]$ there is a path $P^a_i = v^1_i, y_1, \ldots, y_{a-1}, v^2_i$ of length a in Q from v^1_i to v^2_i. Since T_{col} is edge-colourful, it includes each of the colours of the edges on this path exactly once—these colours are $\gamma_{\mathrm{E}}^{-1}(\{v^1_i, y_1\})$,$\gamma_{\mathrm{E}}^{-1}(\{y_1, y_2\})$, $\ldots$,$\gamma_{\mathrm{E}}^{-1}(\{y_{a-1}, v^2_i\})$. Under the edge colouring c_E, the same edges of T_{col} are coloured with the colours $\{v^1_i, y_1\}, \{y_1, y_2\}, \ldots, \{y_{a-1}, v^2_i\}$.

Let $e_1, \ldots, e_a$ be the edges of T_{col} with those colours; we write E^a_i for this set of edges (as is depicted in Fig. 19). We let x^1_i be the vertex of T_{col} which is contained in

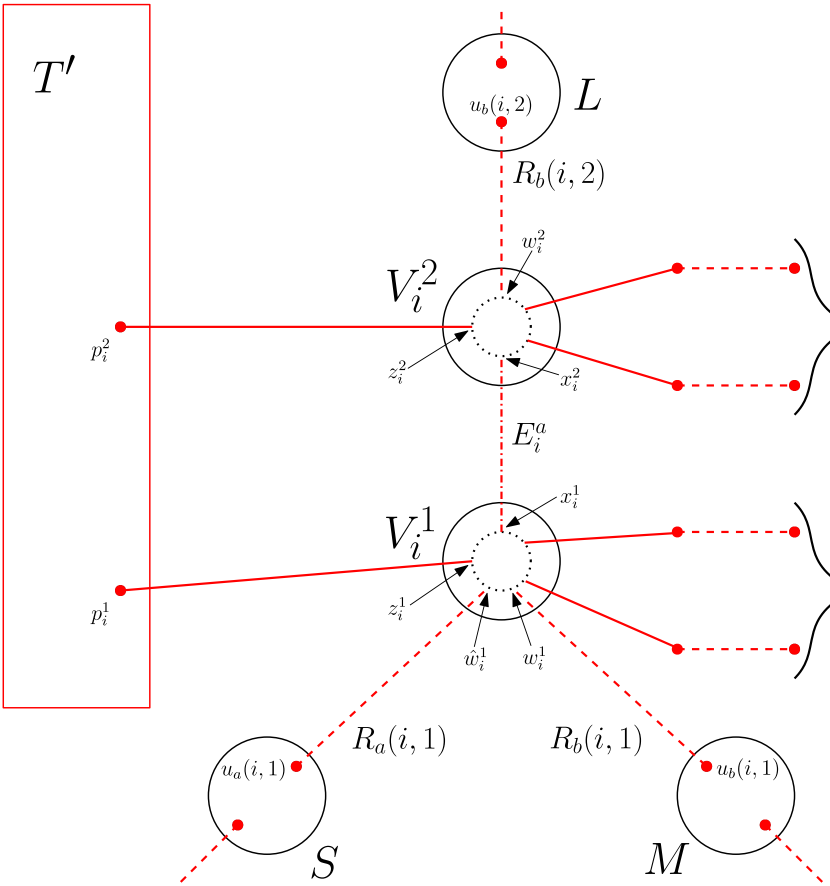

Fig. 19 Depiction of the embedding of T_{col} as established after Claim 2 (in the proof of Lemma 68). Solid lines depict edges, dashed lines depict paths, and dash-dotted lines depict sequences of edges (the identification of the endpoints of which has not yet been determined). Note that we have not yet determined how the endpoints inside of the colour classes V_i^1 and V_i^2 are identified either; this is depicted by the dotted circles inside these colour classes. Proving that the embedding of T_{col} is as depicted in Fig. 17 requires us to show that all endpoints in V_i^2 are identified, and that all endpoints in V_i^1, *except for* x_i^1, are identified (Color figure online)

V_i^1 and incident to e_1, and we let x_i^2 be the vertex of T_{col} which is contained in V_i^2 and incident to e_a. Let z_i^1 and z_i^2 be the vertices of T_{col} in V_i^1 and V_i^2 that are adjacent to p_i^1 and p_i^2—those vertices are depicted in Fig. 19 and we point out that, a priori, x_i^1 might be equal to to z_i^1 and x_i^2 might be equal to z_i^2.

Claim 3 There are no vertices in $V(T_{\mathsf{col}}) \cap V_i^1$ other than z_i^1, x_i^1, w_i^1, $\hat{w}_i^1$ and vertices in the d_i^1 rays.

To prove Claim 3, assume for contradiction that z is such a vertex. Recall that V_i^1 is an independent set (because vertices in V_i^1 all receive the same colour under c.) Since

T_{col} is connected, z has a neighbour outside of V_i^1 but all of the edge colours incident to V_i^1 are already used. □

The proof of the following claim is similar.

Claim 4 There are no vertices in $V(T_{\mathsf{col}}) \cap V_i^2$ other than z_i^2, x_i^2, w_i^2, and vertices in the d_i^2 rays. □

Claim 5 Both z_i^1 and z_i^2 have degree at least 3 in T_{col}. We prove the claim for z_i^1; an analogous argument applies for z_i^2. Assume first for contradiction that z_i^1 has degree 1. Since T_{col} is connected, Claim 2.5 implies that $|V(T_{\mathsf{col}}) \cap V_i^1| = 2$ so $x_i^1 = w_i^1 = \hat{w}_i^1$ and the depicted vertices in the d_i^1 rays are also identified with this vertex. By Definition 69, T_{col} is invalid, giving a contradiction.

Now assume for contradiction that z_i^1 has degree 2. We consider two subcases:

- z_i^1 is identical to x_i^1. Then T_{col} is disconnected, which yields a contradiction.
- z_i^1 is identical to w_i^1 or $\hat{w}_i^1$. This is an immediate contradiction since sources cannot have degree 2 (recall that we already established $R_a(i, 1)$ and $R_b(i, 2)$ to be rays).
- z_i is incident to the first edges of one of the additional d_i^1 outgoing paths. However, in this case, T_{col} can only be connected if there is precisely one further vertex of T_{col} in V_i^1 that is incident to all outgoing edges not covered by z_i^1. However, in this case, T_{col} is an invalid tree, yielding the desired contradiction.

Since the three cases above are exhaustive, the proof of Claim 5 is concluded. □

Next we need the following property:

Claim 6 Let t be a positive integer. If $t < a$ then each ray in $\mathcal{R}^t_{\mathsf{col}}$ is external.

For the proof, recall that $|\mathcal{R}^t| = |\mathcal{R}^t_{\mathsf{col}}|$ since T and T_{col} are isomorphic. Note that each ray R of length t of T is either fully contained in T', or it is one of the d_i^j black rays for some $(i, j) \in [k] \times [2]$. (See Fig. 12) If R is fully contained in T', then R is also contained in $\mathcal{R}^t_{\mathsf{col}}$ and it is external.

If $R = r_0, r_1, \ldots, r_t$ is one of the d_i^j black rays, then T_{col} contains a path $R' = y_0, r_1, \ldots, r_t$ for some $y_0 \in V_i^j$. Suppose that y_0 has degree at least 3 in T_{col}. Then, as in Claim 1, R' is then an external ray, and we are finished. We next consider the case where y_0 has degree 1 or 2 in T_{col}.

If the degree is 1, then T_{col} is disconnected, leading to a contradiction. If the degree is 2, then $y_0 \neq z_i^j$ by Claim 5. Thus, the only way for T_{col} not being disconnected is $y_0 = x_i^j$ and $T_{\mathsf{col}}[E_i^a]$ is a path. However, then we obtained a ray of length at least $a + t$ which is neither external, nor in the set $\hat{\mathcal{R}}$. Thus, we obtain a contradiction by either Claim 2 ($a + t < b$), or by Claim 1 ($a + t = b$), or by Lemma 73 ($a + t > b$). This concludes the proof of Claim 6. □

Next, observe that T_{col} cannot connect z_i^1 and z_i^2 via a path within G, that is, via a path containing the edges E_i^a: Otherwise T_{col} would contain a cycle since p_i^1 and p_i^2 are connected by a path within T'. We will see that z_i^1 and z_i^2 are sources of T_{col}.

Let $\mathcal{S}$ be the set of all sources of T. Consider the *multi-set* of leaf-degrees of T

$$\mathsf{deg}_{\mathsf{L}}(\mathcal{S}) := \{\{\mathsf{deg}_{\mathsf{L}}(s) \mid s \in \mathcal{S}\}\}.$$

Let $\mathcal{S}_{\mathsf{col}}$ be the set of all sources of T_{col} and let $\mathsf{deg}_{\mathsf{L}}(\mathcal{S}_{\mathsf{col}})$ be the muti-set of leaf-degrees T_{col}. Since T_{col} and T are isomorphic, the multi-sets $\mathsf{deg}_{\mathsf{L}}(\mathcal{S})$ and $\mathsf{deg}_{\mathsf{L}}(\mathcal{S}_{\mathsf{col}})$ are equal.

Suppose that $s \in \mathcal{S}$ is a source of T not among the designated sources s_i^j. Then s is contained in T', and it is also a source of T_{col}. Since all of the z_i^j have degree at least 3 in T_{col} (by Claim 5), they cannot be part of further rays with source s in T_{col} so s has the same leaf-degree in T and in T_{col}.

We next show that for each $i \in [k]$, the set $V_i^1 \cup V_i^2$ contains at least 2 sources of T_{col}: Either z_i^1 is a source or it is connected by a 2-path within $T_{\mathsf{col}}[G]$ to another source. However, the only vertices reachable in $T_{\mathsf{col}}[G]$ from z_i^1 that can have degree at least 3 are contained in V_i^2. Similarly, either x_i^2 is a source or it is connected by a 2-path within $T_{\mathsf{col}}[G]$ to a source in V_i^1. We have already seen that z_i^1 cannot be connected to z_i^2 within $T_{\mathsf{col}}[G]$. Thus the sources reachable from z_i^1 and z_i^2 within $T_{\mathsf{col}}[G]$ must be distinct, and we have shown that for each $i \in [k]$, the set $V_i^1 \cup V_i^2$ contains at least 2 sources of T_{col}.

Since T_{col} and T have the same number of sources, and since $2k$ sources of T are not contained in T', we have thus shown that for each $i \in [k]$, the set $V_i^1 \cup V_i^2$ contains *precisely* 2 sources of T_{col}; let us denote those 2 sources by $\hat{z}_i^1$ and $\hat{z}_i^2$.

Now, consider the following subsets of $\mathcal{S}$ and $\mathcal{S}_{\mathsf{col}}$:

- $\mathcal{S}' := \{s_1^1, s_1^2, \ldots, s_k^1, s_k^2\}$ is the set of designated sources.
- $\mathcal{S}'_{\mathsf{col}} := \{\hat{z}_1^1, \hat{z}_1^2, \ldots, \hat{z}_k^1, \hat{z}_k^2\}$ is the set of sources of T_{col} in G (within $\hat{G}$).

Since we already know that $\mathsf{deg}_{\mathsf{L}}(\mathcal{S} \backslash \mathcal{S}') = \mathsf{deg}_{\mathsf{L}}(\mathcal{S}_{\mathsf{col}} \backslash \mathcal{S}'_{\mathsf{col}})$ (those are the sources in T'), we require $\mathsf{deg}_{\mathsf{L}}(\mathcal{S}') = \mathsf{deg}_{\mathsf{L}}(\mathcal{S}'_{\mathsf{col}})$ for T and T_{col} to be isomorphic.

What follows is the final claim within the proof of this lemma.

Claim 7 For all $i \in [k]$, the following five conditions are satisfied:

- $\{z_i^1, z_i^2\} = \{\hat{z}_i^1, \hat{z}_i^2\}$, that is, z_i^1 and z_i^2 are the two sources in $V_i^1 \cup V_i^2$.
- T_{col} contains precisely 2 vertices in V_i^1: One is z_i^1 and one is x_i^1.
- x_i^1 has degree 1. Further, z_i^1, w_i^1, $\hat{w}_i^1$ and all the endpoints of the d_i^1 rays are the same.
- T_{col} contains precisely 1 vertex in V_i^2. Further, z_i^2, x_i^2, w_i^2 and all endpoints of the d_i^2 rays are the same.
- $T_{\mathsf{col}}[E_i^a]$ is a ray with source $z_i^2 (= x_i^2 = w_i^2)$.

Before proving Claim 7, we point out that (1b) and (1f) follow immediately from Claim 7; see Fig. 17 and observe that $T_{\mathsf{col}}[E_i^a]$ becomes the ray $R_a(i, 2)$, and x_i^1 becomes the endpoint $u_a(i, 2)$ of $R_a(i, 2)$ for each $i \in [k]$. Thus the proof of this lemma is concluded if Claim 7 is proved, which is done below:

- We first show that $\{z_i^1, z_i^2\} = \{\hat{z}_i^1, \hat{z}_i^2\}$ for each $i \in [k]$. Let $\Phi = \sum_{s \in \mathcal{S}'} \mathsf{deg}_{\mathsf{L}}(s)$ and $\Phi_{\mathsf{col}} = \sum_{s \in \mathcal{S}'_{\mathsf{col}}} \mathsf{deg}_{\mathsf{L}}(s)$. Observe that $\mathsf{deg}_{\mathsf{L}}(\mathcal{S}') = \mathsf{deg}_{\mathsf{L}}(\mathcal{S}'_{\mathsf{col}})$ implies $\Phi = \Phi_{\mathsf{col}}$. We start by observing that

$$\mathsf{deg}_{\mathsf{L}}(\hat{z}_i^1) + \mathsf{deg}_{\mathsf{L}}(\hat{z}_i^2) \leq (d_i^1 + 2) + (d_i^2 + 1) + 2.$$

There are d_i^1 rays from V_i^1 and also $R_a(i, 1)$ and $R_b(i, 1)$. There are d_i^2 rays from V_i^2 and also $R_b(i, 2)$. There is also E_i^a which could form two rays.
We next show that E_i^a cannot form two rays. Assume for contradiction that is does. Since T_{col} is connected, $z_i^1, w_i^1, \hat{w}_i^1, x_i^1$ and all the endpoints of the d_i^1 rays are identical, and z_i^2, w_i^2, x_i^2 and all the endpoints of the d_i^2 rays are identical.
Now, if $T_{\mathsf{col}}[E_i^a]$ would be the disjoint union of two rays of length less than a with sources z_i^1 and z_i^2 then those rays are non-external rays of length less than a, contradicting Claim 6. We have now shown

$$\mathsf{deg}_{\mathsf{L}}(\hat{z}_i^1) + \mathsf{deg}_{\mathsf{L}}(\hat{z}_i^2) \leq (d_i^1 + 2) + (d_i^2 + 2). \tag{11}$$

Next, note that by definition of the d_i^j (see Fig. 12), the following holds:

$$(d_i^1 + 2) + (d_i^2 + 2) = \mathsf{deg}_{\mathsf{L}}(s_i^1) + \mathsf{deg}_{\mathsf{L}}(s_i^2) \tag{12}$$

We have now shown that

$$\mathsf{deg}_{\mathsf{L}}(\hat{z}_i^1) + \mathsf{deg}_{\mathsf{L}}(\hat{z}_i^2) \leq \mathsf{deg}_{\mathsf{L}}(s_i^1) + \mathsf{deg}_{\mathsf{L}}(s_i^2).$$

Finally, we will show that z_i^1 and z_i^2 are sources to finish the first bullet point.
Consider z_i^1, and recall that is has degree at least 3 by Claim 5, and assume for contradiction that it is not a source of T_{col}. Then $z_i^1 = x_i^1$, and $T_{\mathsf{col}}[E_i^a]$ is a path, and x_i^2 is source (since it is the only vertex in $V(T_{\mathsf{col}}) \cap V_i^2$ that might have degree at least 3, except for z_i^2). Note that this also implies that z_i^2 is a source. Thus $\{\hat{z}_i^1, \hat{z}_i^2\} = \{x_i^2, z_i^2\}$. In this case, we have

$$\mathsf{deg}_{\mathsf{L}}(\hat{z}_i^1) + \mathsf{deg}_{\mathsf{L}}(\hat{z}_i^2) \leq d_i^2 + 1 < \mathsf{deg}_{\mathsf{L}}(s_i^1) + \mathsf{deg}_{\mathsf{L}}(s_i^2).$$

Consequently, using (11) and (12), we have $\Phi_{\mathsf{col}} < \Phi$, which is a contradiction. Thus z_i^1 is a source of T_{col}, and a similar argument shows that z_i^2 is a source of T_{col} as well.

- We now prove the remaining items. In what follows, using the previous bulleted item, we can assume that w.l.o.g. $\hat{z}_i^1 = z_i^1$ and $\hat{z}_i^2 = z_i^2$ for all $i \in [k]$. First, recall that we ordered the s_i^j by their leaf-degrees, that is

$$\mathsf{deg}_{\mathsf{L}}(s_1^1) \geq \mathsf{deg}_{\mathsf{L}}(s_1^2) \geq \cdots \geq \mathsf{deg}_{\mathsf{L}}(s_k^2) \geq 2.$$

If x_1^1 were equal to z_1^1, then T_{col} can only be connected if there is only one vertex in V_1^1, that is, all edges incident to V_1^1 are in fact incident to z_1^1. However, in that case, we have $\mathsf{deg}_{\mathsf{L}}(z_1^1) = \mathsf{deg}_{\mathsf{L}}(s_1^1) + 1$ (by construction of $\hat{G}$), and thus the multi-sets cannot be equal anymore. Hence $x_1^1 \neq z_1^1$.
If x_1^1 had degree 2, then there would have been a ray of length at least $a + 1$ that originates in V_1^2 (otherwise T_{col} would have been disconnected). However, this ray would neither be external, nor among the rays in $\hat{\mathcal{R}}$, contradicting either Lemma 73

or the previous sequence of claims. Finally, if x_1^1 had degree at least 3, then T_{col} would have contained more sources than T, which also yields a contradiction. This shows that x_1^1 has degree 1. However, this implies that T_{col} can only contain one vertex in V_i^2; otherwise T_{col} would be disconnected. Note that we have just proved the remaining items of Claim 7 for $i = 1$. Additionally, we have shown that $\mathsf{deg}_{\mathsf{L}}(z_1^1) = \mathsf{deg}_{\mathsf{L}}(s_1^1)$ and $\mathsf{deg}_{\mathsf{L}}(z_1^2) = \mathsf{deg}_{\mathsf{L}}(s_1^2)$ Hence we can remove those two numbers from the multi-sets and continue recursively with $i = 2$. This concludes the proof of Claim 7, and thus the proof of the overall lemma. □

We are now ready to conclude the case for trees of unbounded fork number.

Lemma 74 *Let $\mathcal{T}$ be a recursively enumerable class of trees of unbounded fork number. Then $\oplus\textsc{Sub}(\mathcal{T})$ is $\oplus\mathsf{W}[1]$-hard.*

Proof We proceed similarly to Lemma 44. However, we have to take care of some subtleties. First, we start with a class $\mathcal{C}$ of cubic *bipartite* graphs of unbounded treewidth. Next, we wish to rely on Lemma 68 to obtain the identity

$$\oplus\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c)) = \oplus\mathsf{ColSub}(T \to (\hat{G}, \hat{\gamma})),$$

where τ is the fracture defined in Definition 61. Unfortunately, Lemma 68 only yields the above identity if, for each $v \in V(Q)$, $|c^{-1}(v)|$ is odd, that is, each colour class of vertices of G has odd cardinality. However, this property can easily be achieved. Let (G', c') be the Q-coloured graph obtained from (G, c) by adding to each even colour class one fresh isolated vertex. Since $Q\sharp\tau$ does not have isolated vertices, this operation does not change the number of colour-preserving embeddings. In combination with Lemma 68 we thus obtain

$$\begin{aligned}\oplus\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G, c)) &= \oplus\mathsf{Emb}((Q\sharp\tau, c_\tau) \to (G', c')) \\ &= \oplus\mathsf{ColSub}(T \to (\hat{G}', \hat{\gamma})).\end{aligned}$$

From here on, we can proceed analogously to the proof of Lemma 44. □

4.4 The Dichotomy Theorem for Trees

We are now able to prove Theorem 5, i.e., an exhaustive and explicit parameterised complexity classification for counting trees modulo 2:

Theorem 5 Let $\mathcal{T}$ be a recursively enumerable class of trees. If $\mathcal{T}$ is matching splittable, then $\oplus\textsc{Sub}(\mathcal{T})$ is fixed-parameter tractable. Otherwise $\oplus\textsc{Sub}(\mathcal{T})$ is $\oplus\mathsf{W}[1]$-complete.

Proof The fixed-parameter tractability result, as well as the fact that $\oplus\textsc{Sub}(\mathcal{T})$ is always contained in $\oplus\mathsf{W}[1]$ were both shown in [12]. Hence, it remains to prove $\oplus\mathsf{W}[1]$-hardness if $\mathcal{T}$ is not matching splittable.

By Lemma 32 each class $\mathcal{T}$ of trees that is not matching splittable has unbounded C-number, unbounded star number, or unbounded fork number. Finally, each of these three cases yields $\oplus\mathsf{W}[1]$-hardness as established by Lemmas 44, 56, and 74. □

5 Conclusion and Open Questions

Given a class $\mathcal{H}$ of patterns, the problem $\oplus$SUB($\mathcal{H}$) asks, given as input a graph $H \in \mathcal{H}$ and an arbitrary graph G, to count the subgraphs of G that are isomorphic to H.

This work is motivated by the conjecture of Curticapean, Dell and Husfeldt (Conjecture 1) that $\oplus$SUB($\mathcal{H}$) is FPT if and only if $\mathcal{H}$ is *matching splittable*.

Recall that the *matching-split number* of H is the minimum size of a set $S \subseteq V(H)$ such that $H \setminus S$ is a matching. The class $\mathcal{H}$ is *matching splittable* if there is a positive integer B such that the matching-split number of any $H \in \mathcal{H}$ is at most B.

In this work, Theorem 4 proves the conjecture for every hereditary class of graphs. Theorem 5 proofs the conjecture for every class $\mathcal{H}$ of trees.

Clearly, the most important task for related future work is to fully resolve the conjecture. Our work has shown that the use of edge-colours, formalised by the framework of fractured graphs, makes it possible to bypass the problem caused by automorphism groups that have even cardinality. We think that this approach will be useful for future work.

5.1 Related Version

An extended abstract of this work, not containing the proof of the classification for trees, is accepted for publication at the 50th EATCS International Colloquium on Automata, Languages and Programming (ICALP 2023). https://doi.org/10.4230/LIPIcs.ICALP.2023.68.

Acknowledgements We want to thank Radu Curticapean, Holger Dell and Thore Husfeldt for insightful discussions on an early draft of this work.

Author Contributions LG and MR worked in equal parts on the manuscript.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

1. Abboud, A., Feller, S., Weimann, O.: On the fine-grained complexity of parity problems. In: Czumaj, A., Dawar, A., Merelli, E. (eds.) 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8–11, 2020, Saarbrücken, Germany (Virtual Conference), LIPIcs, vol.

168, pp. 1–19. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2020). https://doi.org/10.4230/LIPIcs.ICALP.2020.5
2. Alon, N., Dao, P., Hajirasouliha, I., Hormozdiari, F., Cenk Sahinalp, S.: Biomolecular network motif counting and discovery by color coding. Bioinformatics **24**(13), i241–i249 (2008). https://doi.org/10.1093/bioinformatics/btn163
3. Alon, N., Yuster, R., Zwick, U.: Color-coding. J. ACM **42**(4), 844–856 (1995). https://doi.org/10.1145/210332.210337
4. Bera, S.K., Gishboliner, L., Levanzov, Y., Seshadhri, C., Shapira, A.: Counting subgraphs in degenerate graphs. J. ACM **69**(3), 1–21 (2022). https://doi.org/10.1145/3520240
5. Björklund, A., Dell, H., Husfeldt, T.: The parity of set systems under random restrictions with applications to exponential time problems. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) Automata, Languages, and Programming—42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6–10, 2015, Proceedings, Part I, Lecture Notes in Computer Science, vol. 9134, pp. 231–242. Springer (2015). https://doi.org/10.1007/978-3-662-47672-7_19
6. Bressan, M.: Faster algorithms for counting subgraphs in sparse graphs. Algorithmica **83**(8), 2578–2605 (2021). https://doi.org/10.1007/s00453-021-00811-0
7. Bulatov, A.A., Kazeminia, A.: Complexity classification of counting graph homomorphisms modulo a prime number. In: Leonardi, S., Gupta, A. (eds.) STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20–24, 2022, pp. 1024–1037. ACM (2022). https://doi.org/10.1145/3519935.3520075
8. Chen, J., Chor, B., Fellows, M., Huang, X., Juedes, D.W., Kanj, I.A., Xia, G.: Tight lower bounds for certain parameterized NP-hard problems. Inf. Comput. **201**(2), 216–231 (2005). https://doi.org/10.1016/j.ic.2005.05.001
9. Chen, J., Huang, X., Kanj, I.A., Xia, G.: Strong computational lower bounds via parameterized complexity. J. Comput. Syst. Sci. **72**(8), 1346–1367 (2006). https://doi.org/10.1016/j.jcss.2006.04.007
10. Chen, Y., Thurley, M., Weyer, M.: Understanding the complexity of induced subgraph isomorphisms. In: Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP), pp. 587–596. Springer (2008). https://doi.org/10.1007/978-3-540-70575-8_48
11. Curticapean, R.: Counting matchings of size k is w[1]-hard. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M.Z., Peleg, D. (eds.) Automata, Languages, and Programming—40th International Colloquium, ICALP 2013, Riga, Latvia, July 8–12, 2013, Proceedings, Part I, Lecture Notes in Computer Science, vol. 7965, pp. 352–363. Springer (2013). https://doi.org/10.1007/978-3-642-39206-1_30
12. Curticapean, R., Dell, H., Husfeldt, T.: Modular counting of subgraphs: matchings, matching-splittable graphs, and paths. In: Mutzel, P., Pagh, R., Herman, G. (eds.) 29th Annual European Symposium on Algorithms, ESA 2021, September 6–8, 2021, Lisbon, Portugal (Virtual Conference), LIPIcs, vol. 204, pp. 1–17. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2021). https://doi.org/10.4230/LIPIcs.ESA.2021.34
13. Curticapean, R., Dell, H., Marx, D.: Homomorphisms are a good basis for counting small subgraphs. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC), pp. 210–223. ACM (2017). https://doi.org/10.1145/3055399.3055502
14. Curticapean, R., Marx, D.: Complexity of counting subgraphs: only the boundedness of the vertex-cover number counts. In: 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18–21, 2014, pp. 130–139. IEEE Computer Society (2014). https://doi.org/10.1109/FOCS.2014.22
15. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer (2015). https://doi.org/10.1007/978-3-319-21275-3
16. Dalmau, V., Jonsson, P.: The complexity of counting homomorphisms seen from the other side. Theor. Comput. Sci. **329**(1–3), 315–323 (2004). https://doi.org/10.1016/j.tcs.2004.08.008
17. Dörfler, J., Roth, M., Schmitt, J., Wellnitz, P.: Counting induced subgraphs: An algebraic approach to #w[1]-hardness. Algorithmica **84**(2), 379–404 (2022). https://doi.org/10.1007/s00453-021-00894-9
18. Downey, R.G., Fellows, M.R.: Fundamentals of Parameterized Complexity. Texts in Computer Science. Springer (2013). https://doi.org/10.1007/978-1-4471-5559-1
19. Flum, J., Grohe, M.: The parameterized complexity of counting problems. SIAM J. Comput. **33**(4), 892–922 (2004). https://doi.org/10.1137/S0097539703427203

20. Flum, J., Grohe, M.: Parameterized Complexity Theory. Texts in Theoretical Computer Science. An EATCS Series. Springer (2006). https://doi.org/10.1007/3-540-29953-X
21. Grohe, M., Marx, D.: On tree width, bramble size, and expansion. J. Comb. Theory Ser. B **99**(1), 218–228 (2009). https://doi.org/10.1016/j.jctb.2008.06.004
22. Harvey, D.J., Wood, D.R.: The treewidth of line graphs. J. Comb. Theory Ser. B **132**, 157–179 (2018). https://doi.org/10.1016/j.jctb.2018.03.007
23. Impagliazzo, R., Paturi, R.: On the complexity of k-sat. J. Comput. Syst. Sci. **62**(2), 367–375 (2001). https://doi.org/10.1006/jcss.2000.1727
24. Jansen, B.M.P., Marx, D.: Characterizing the easy-to-find subgraphs from the viewpoint of polynomial-time algorithms, kernels, and turing kernels. In: Indyk, P. (ed.) Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015, pp. 616–629. SIAM (2015). https://doi.org/10.1137/1.9781611973730.42
25. Kuchaiev, O., Milenković, T., Memišević, V., Hayes, W., Pržulj, N.: Topological network alignment uncovers biological function and phylogeny. J. R. Soc. Interface **7**(50), 1341–1354 (2010). https://doi.org/10.1098/rsif.2010.0063
26. Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., Alon, U.: Network motifs: simple building blocks of complex networks. Science **298**(5594), 824–827 (2002). https://doi.org/10.1126/science.298.5594.824
27. Milo, R., Itzkovitz, S., Kashtan, N., Levitt, R., Shen-Orr, S., Ayzenshtat, I., Sheffer, M., Alon, U.: Superfamilies of evolved and designed networks. Science **303**(5663), 1538–1542 (2004). https://doi.org/10.1126/science.1089167
28. Peyerimhoff, N., Roth, M., Schmitt, J., Stix, J., Vdovina, A.: Parameterized (modular) counting and Cayley graph expanders. In: Bonchi, F., Puglisi, S.J. (eds.) 46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23–27, 2021, Tallinn, Estonia, LIPIcs, vol. 202, pp. 1–15. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2021). https://doi.org/10.4230/LIPIcs.MFCS.2021.84
29. Peyerimhoff, N., Roth, M., Schmitt, J., Stix, J., Vdovina, A.: Parameterized (modular) counting and Cayley graph expanders. CoRR arXiv:2104.14596 (2021)
30. Peyerimhoff, N., Roth, M., Schmitt, J., Stix, J., Vdovina, A., Wellnitz, P.: Parameterized counting and Cayley graph expanders. SIAM J. Discrete Math., to appear
31. Roth, M., Schmitt, J., Wellnitz, P.: Detecting and counting small subgraphs, and evaluating a parameterized Tutte polynomial: lower bounds via toroidal grids and Cayley graph expanders. CoRR arXiv:2011.03433 (2020)
32. Roth, M., Schmitt, J., Wellnitz, P.: Detecting and counting small subgraphs, and evaluating a parameterized Tutte polynomial: lower bounds via toroidal grids and Cayley graph expanders. In: Bansal, N., Merelli, E., Worrell, J. (eds.) 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12–16, (2021), Glasgow, Scotland (Virtual Conference), LIPIcs, vol. 198, pp. 1–16. Schloss Dagstuhl—Leibniz-Zentrum für Informatik (2021). https://doi.org/10.4230/LIPIcs.ICALP.2021.108
33. Schiller, B., Jager, S., Hamacher, K., Strufe, T.: StreaM—a stream-based algorithm for counting motifs in dynamic graphs. In: Proceedings of the 2nd International Conference on Algorithms for Computational Biology (AlCoB), pp. 53–67 (2015). https://doi.org/10.1007/978-3-319-21233-3_5
34. Toda, S.: PP is as hard as the polynomial-time hierarchy. SIAM J. Comput. **20**(5), 865–877 (1991). https://doi.org/10.1137/0220053
35. Tran, N.H., Choi, K.P., Zhang, L.: Counting motifs in the human interactome. Nat. Commun. **4**(1), 1–8 (2013). https://doi.org/10.1038/ncomms3241
36. Tsourakakis, C.E., Pachocki, J., Mitzenmacher, M.: Scalable motif-aware graph clustering. In: Proceedings of the 26th International Conference on World Wide Web (WWW), pp. 1451–1460 (2017). https://doi.org/10.1145/3038912.3052653
37. Williams, V.V., Wang, J.R., Williams, R.R., Yu, H.: Finding four-node subgraphs in triangle time. In: Indyk, P. (ed.) Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015, pp. 1671–1680. SIAM (2015). https://doi.org/10.1137/1.9781611973730.111

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.