

A Logic-based Algorithmic Meta-Theorem for Treedepth: Single Exponential FPT Time and Polynomial Space*

Benjamin Bergougnoux¹, Vera Chekan², and Giannos Stamoulis³

¹Aix Marseille Université, CNRS, LIS, Marseille, France
benjamin.bergougnoux@lis-lab.fr

²Humboldt-Universität zu Berlin, Germany
vera.chekan@hu-berlin.de

³Université Paris Cité, CNRS, IRIF, F-75013, Paris, France
stamoulis@irif.fr

Abstract

For a graph G , the parameter treedepth measures the minimum depth among all forests F , called elimination forests, such that G is a subgraph of the ancestor-descendant closure of F . We introduce a logic, called *neighborhood operator logic with acyclicity, connectivity and clique constraints* ($\text{NEO}_2[\text{FRec}]+\text{ACK}$ for short), that captures all NP-hard problems—like INDEPENDENT SET or HAMILTONIAN CYCLE—that are known to be tractable in time $2^{\mathcal{O}(\text{td})}n^{\mathcal{O}(1)}$ and space $n^{\mathcal{O}(1)}$ on n -vertex graphs provided with elimination forests of depth td . We provide a model checking algorithm for $\text{NEO}_2[\text{FRec}]+\text{ACK}$ with such complexity that unifies and extends these results. For $\text{NEO}_2[\text{FRec}]+\text{K}$, the fragment of the above logic that does not use acyclicity and connectivity constraints, we get a strengthening of this result, where the space complexity is reduced to $\mathcal{O}(\text{td} \log(n))$.

With a similar mechanism as the distance neighborhood logic introduced in [Bergougnoux, Dreier and Jaffke, *SODA 2023*], the logic $\text{NEO}_2[\text{FRec}]+\text{ACK}$ is an extension of the fully-existential MSO_2 with predicates for (1) querying generalizations of the neighborhoods of vertex sets, (2) verifying the connectivity and acyclicity of vertex and edge sets, and (3) verifying that a vertex set induces a clique. Interestingly, $\text{NEO}_2[\text{FRec}]$, the fragment of $\text{NEO}_2[\text{FRec}]+\text{K}$ that does not use clique constraints, is equivalent (up to minor features) to a variant of modal logic—introduced in [Pilipczuk, *MFCS 2011*—that captures many problems known to be tractable in single exponential time when parameterized by treewidth. Our results provide $2^{\mathcal{O}(\text{td})}n^{\mathcal{O}(1)}$ time and $n^{\mathcal{O}(1)}$ space algorithms for problems for which the existence of such algorithms was previously unknown. In particular, $\text{NEO}_2[\text{FRec}]$ captures CNF-SAT via the incidence graphs associated to CNF formulas, and it also captures several modulo counting problems like ODD DOMINATING SET.

To prove our results, we extend the applicability of algebraic transforms such as the inclusion-exclusion principle and the discrete Fourier transform. To our knowledge, this is the first time, the discrete Fourier transform have been used to obtain space-efficient algorithms on graphs of bounded treedepth. To achieve the logspace complexity for $\text{NEO}_2[\text{FRec}]+\text{K}$, we also use the technique from [Pilipczuk and Wrochna, *ACM Trans. Comput. Theory 2018*] based on Chinese remainder theorem.

*G.S. was supported by the project BOBR that is funded from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme with grant agreement No. 948057. In particular, a majority of work on this manuscript was done while B.B. and G.S. were affiliated with University of Warsaw. This project was initiated during the research stay of V.C. at the University of Warsaw.

Contents

1	Introduction	1
2	Technical Overview	6
3	Preliminaries	11
3.1	Graphs	11
3.2	Elimination trees and treedepth	11
3.3	Discrete Fourier transforms	12
3.4	Inclusion-exclusion	13
4	Neighborhood Operator Logics	14
4.1	Problems captured by our logics	16
4.2	Core logic	16
5	Partial Solutions and Related Notions	18
5.1	Formula parameters and auxiliary digraph	18
5.2	Neighborhood functions, mod-degrees and partial interpretations	19
5.3	Partial solutions, indexes, and compatibility	21
5.4	Relations between partial solutions and the formula satisfiability	24
6	Recursive Equalities	27
6.1	The definitions of polynomials	27
6.2	Root equalities	29
6.3	Leaf equality	30
6.4	Forget equality	33
6.5	Join equality	41
7	Model Checking	52
7.1	Getting to logarithmic space complexity	52
7.2	Wrapping up the model checking algorithm	56
8	Extensions of the Logic	57
8.1	Clique predicate	57
8.2	Connectivity and acyclicity	58
9	Conclusions	63
A	Discrete Fourier Transforms for Polynomials	68

1 Introduction

Many NP-hard problems—like INDEPENDENT SET and HAMILTONIAN CYCLE—are tractable in time $2^{\mathcal{O}(\text{tw})}n^{\mathcal{O}(1)}$ with tw and n being the treewidth and the number of vertices of the input graph, respectively (see [8, 17, 19, 35, 40, 41] for examples). However, the space usage of these algorithms is exponential in the width of the decomposition. In practical applications, this is often a prohibitive factor, because a machine is likely to simply run out of space much before the elapsed time exceeds tolerable limits. Therefore, a natural question arises: Can we reduce the space complexity of these algorithms while keeping their time complexity? Lokshtanov, Mnich, and Saurabh [29] sketched that many problems admitting a $2^{\mathcal{O}(\text{tw})}n^{\mathcal{O}(1)}$ time algorithm can be solved in polynomial space and time $2^{\mathcal{O}(\text{tw}^2)}n^{\mathcal{O}(1)}$. Unfortunately, some complexity-theoretical evidences [1, 37] suggest that this increase in the running time might be inevitable, and problems like INDEPENDENT SET are probably not solvable by any algorithm using $2^{\mathcal{O}(\text{tw})}n^{\mathcal{O}(1)}$ time and $2^{o(\text{tw})}n^{\mathcal{O}(1)}$ space.

Polynomial space and treedepth. This reduction of space complexity is possible when we consider *treedepth*: a shallow restriction of treewidth. Formally, the treedepth of a graph G is the minimum depth of a rooted forest F —called an *elimination forest* of G —in which G can be embedded by bijectively mapping the vertices of G to the nodes of F such that the endpoints of every edge of G are in an ancestor-descendant relation in F . This parameter was introduced in [33] by Nešetřil and Ossona de Mendez; see e.g. [34, Chapter 6] for a general discussion on the notion of treedepth. The treedepth of a graph is never smaller than its treewidth, but it is also never larger than its treewidth times $\log n$ [9] where n denotes the number of vertices in G . Several NP-hard problems such as INDEPENDENT SET or HAMILTONIAN CYCLE can be solved in time $2^{\mathcal{O}(\text{td})}n^{\mathcal{O}(1)}$ and polynomial space (i.e., $n^{\mathcal{O}(1)}$) where td is the treedepth of the input graph, assuming that we are given a decomposition certifying this value of td , i.e., an elimination tree of G of depth td [22, 27, 32, 37]. In certain cases, the space complexity can even be as low as $\mathcal{O}(\text{td} + \log n)$ or $\mathcal{O}(\text{td} \log n)$ [37]. Typically, the main idea is to reformulate the classic bottom-up dynamic programming approach so that it can be replaced by a simple top-down recursion. This reformulation is by no means easy—it often involves a highly non-trivial use of algebraic transforms or other tools of algebraic flavor, such as inclusion-exclusion branching. These algorithmic applications of treedepth lead naturally to the following question:

What families of problems can be solved in time $2^{\mathcal{O}(\text{td})}n^{\mathcal{O}(1)}$ and polynomial space, assuming that an elimination forest of depth td is provided with input?

We provide a wide family of such problems by introducing a logic called *neighborhood operator logic with acyclicity* (A), *connectivity* (C), and *clique* (K) *constraints* ($\text{NEO}_2[\text{FRec}] + \text{ACK}$ for short) and we show that the model checking problem, for every fixed $\text{NEO}_2[\text{FRec}] + \text{ACK}$ formula, can be solved in time $2^{\mathcal{O}(\text{td})}n^{\mathcal{O}(1)}$ and polynomial space, provided an elimination tree of depth td is part of the input. For $\text{NEO}_2[\text{FRec}] + \text{K}$, i.e., the same logic without the acyclicity and connectivity constraints, we can further reduce the space complexity of our model checking algorithm to $\mathcal{O}(\text{td} \log n)$.

Before proceeding to the description of our logic and a more detailed discussion about our results, let us mention here that the assumption of having an elimination tree as part of the input is an essential part of this question. To our best knowledge, the current state-of-the-art in polynomial space algorithms for computing an elimination tree of depth td of a graph of treedepth td is the one of Nadara, Pilipczuk, and Smulewicz [31], which runs in time $2^{\mathcal{O}(\text{td}^2)}n^{\mathcal{O}(1)}$ (and polynomial space). Recently, Bonnet, Neuen, and Sokolowski [10] showed that the 1.0003-approximation of treedepth is NP-hard. And it is a major open question whether one can obtain a constant-factor approximation of treedepth in time $2^{o(\text{td}^2)}n^{\mathcal{O}(1)}$.

Neighborhood operator logics. The logic $\text{NEO}_2[\text{FRec}]$ extends fully-existential MSO_2 with the additional ability to directly reason about the generalized neighborhoods of sets. The central building block in our neighborhood logic is the following neighborhood operator:

Definition 1.1. *Given a graph G , $\sigma \subseteq \mathbb{N}$, $U \subseteq V(G)$ and $F \subseteq E(G)$, we define the (σ, F) -neighborhood of U in G as*

$$N^\sigma(U, F) := \{v \in V(G) : |\{u \in U : uv \in F\}| \in \sigma\}.$$

Moreover, we define the σ -neighborhood of U , denoted by $N^\sigma(U)$, as the set $N^\sigma(U, E(G))$.

Observe that $N^{\mathbb{N}^+}(U)$ with $\mathbb{N}^+ = \{1, 2, 3, \dots\}$ corresponds to the usual open neighborhood of a set of vertices and $N^{\mathbb{N}^+}(U, F)$ corresponds to the open neighborhood of U in the graph $(V(F), F)$. In this paper, we deal with neighborhood operators $N^\sigma(U, F)$ where σ is *finitely recognizable*:

Definition 1.2. *A subset σ of $\mathbb{N}$ is finitely recognizable if there exist integers $d, p \in \mathbb{N}$ with $p > 0$ such that for every $n \geq d$, we have*

$$n \in \sigma \text{ if and only if } n + p \in \sigma. \quad (1)$$

We denote by $p(\sigma)$ and $d(\sigma)$ the smallest p and d satisfying (1), respectively. Moreover, we call $p(\sigma)$ the period of σ .

We use this neighborhood operator to define so-called *neighborhood terms*. These are built from vertex set variables, vertex set constants (i.e., unary relations) or other neighborhood terms by applying neighborhood operators of form $N^\sigma(\cdot)$ as well as standard set operators such as intersection ($\cap$), union ($\cup$), subtraction ($\setminus$) or complementation (denoted by a bar on top of the term). The second arguments of our neighborhood operators are *edge set terms*, those are built from edge set variables and edge set constants (i.e., binary relations) by applying the above-mentioned standard set operators. Then $\text{NEO}_2[\text{FRec}]$ is the extension of the fully-existential MSO_2 by allowing the following:

- *Size measurement of terms:* We can write, for example, $|t| \leq m$ to express that a neighborhood term or an edge set term t should have size at most m .
- *Comparison between terms:* We can write, for example, $t_1 = t_2$ or $t_1 \subseteq t_2$ to express that t_1 should be equal or contained, respectively, in t_2 where t_1 and t_2 are two neighborhood terms or two edge set terms.

We define $\text{NEO}_2[\text{FRec}]+\text{K}$ as the extension of $\text{NEO}_2[\text{FRec}]$ with predicates $\text{clique}(t)$ that verify that a neighborhood term t induces a clique. Furthermore, we define $\text{NEO}_2[\text{FRec}]+\text{ACK}$ as the extension of $\text{NEO}_2[\text{FRec}]+\text{K}$ with predicates $\text{con}(t)$ and $\text{acy}(t)$ that verify the connectivity and acyclicity, respectively, of a neighborhood term or an edge set term t . We refer to Section 4 for the formal definition of this logic. Let us give some example of $\text{NEO}_2[\text{FRec}]$ formulas that capture some well-known problems. The very simple formula $\text{independent}(X) := N^{\mathbb{N}^+}(X) \cap X = \emptyset$ expresses that no vertex from X intersects the neighborhood of X , i.e., that X forms an independent set. Thus, the $\text{NEO}_2[\text{FRec}]$ formula $\exists X |X| \geq t \wedge \text{independent}(X)$ expresses the existence of an independent set of size at least t . Similarly, the formula

$$\exists Y |Y| \leq t \wedge \text{independent}(V \setminus N^{\mathbb{N}^+}(V, Y))$$

(where V stands for the whole vertex set) expresses the existence of an edge dominating set of size at most t , where $V \setminus N^{\mathbb{N}^+}(V, Y)$ is the set of vertices that are not incident to an edge of Y (where

V denotes the set of all vertices). With $\text{NEO}_2[\text{FRec}]+\text{ACK}$, we can trivially express the **CLIQUE** problem. Moreover, we can express connectivity problems, for example, for **HAMILTONIAN CYCLE** we can use the formula

$$\exists Y \text{ con}(Y) \wedge N^{\{2\}}(V, Y) = V,$$

where $N^{\{2\}}(V, Y) = V$ expresses the property that every vertex of the graph is incident to exactly two edges of Y . For **MAX CUT** to express the existence of a cut of size at least k , we can use the formula:

$$\exists Y \exists A \exists B \ A \cup B = V \wedge A \cap B = \emptyset \wedge N^{\mathbb{N}^+}(A, Y) \subseteq B \wedge N^{\mathbb{N}^+}(B, Y) \subseteq A \wedge |Y| \geq k.$$

For any fixed integer r we can also express the existence of an induced r -regular subgraph with at least k vertices using the following formula:

$$\exists X \exists Y \ N^{\mathbb{N}^+}(V, Y) = X \wedge N^{\mathbb{N}^+}(X, E \setminus Y) = \emptyset \wedge N^{\{r\}}(X, Y) = X \wedge |X| \geq k.$$

Here, X and Y are the vertex and the edge set of the sought subgraph, respectively. We first ensure that the set of the endpoints of the edges in Y is precisely X , then that all edges induced by X belong to Y , and finally that the graph (X, Y) is r -regular and contains at least k vertices.

Observe that finitely recognizable sets allow us to express problems based on modulo counting such as the ones defined in [3, 24] such as **ODD DOMINATING SET**, or more generally the existence of a (σ, ρ) -set [39] where σ and ρ are finitely recognizable:

$$\exists X \ X \subseteq N^\sigma(X) \wedge \bar{X} \subseteq N^\rho(X).$$

And even more generally, we can model a generalization of this problem called (σ, ρ) -**MINPARDOMSET** introduced by Greilhuber and Marx [23] (they only consider it for (co)finite sets σ and ρ) where one allows a set of at most k vertices to violate the (σ, ρ) -constraints, i.e.,

$$\exists X \ |(X \setminus N^\sigma(X)) \cup (\bar{X} \setminus N^\rho(X))| \leq k.$$

In Subsection 4.1, we provide more examples of problems captured by our logics: *locally checkable partitioning* problem, **CNF-SAT** and **MAX-CNF-SAT** (via the incidence graph of a formula).

Our results. In order to describe the running time of our model checking algorithms, we need to introduce the following notation. Given a $\text{NEO}_2[\text{FRec}]+\text{ACK}$ formula φ , we define $d(\varphi)$ as the maximum of $d(\sigma)$ over all sets $\sigma \subseteq \mathbb{N}$ that occur in φ , and we define the value p_φ as the least common multiple of $p(\sigma)$ over the sets $\sigma \subseteq \mathbb{N}$ that occur in φ . First, we show that given a $\text{NEO}_2[\text{FRec}]+\text{K}$ formula, the model checking problem can be solved in single-exponential time and logarithmic space. Formally, we prove the following theorem:

Theorem 1.3. *There is an algorithm that given a formula $\varphi \in \text{NEO}_2[\text{FRec}]+\text{K}$ with ℓ size measurements, a graph G and an elimination tree T of G of depth at most td , runs in time $((d_\varphi + 1) \cdot p_\varphi)^{\mathcal{O}(|\varphi| \cdot \text{td})} \cdot n^{\mathcal{O}(\ell)}$ and space $\mathcal{O}(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$ and decides whether $G \models \varphi$ holds.*

When the formula contains acyclicity and connectivity constraints, we prove the existence of the following model checking algorithm:

Theorem 1.4. *There is a Monte-Carlo algorithm that given a formula $\varphi \in \text{NEO}_2[\text{FRec}]+\text{ACK}$ with ℓ size measurements and q acyclicity constraints, a graph G and an elimination tree T of G of depth at most td , runs in time $((d_\varphi + 1) \cdot p_\varphi)^{\mathcal{O}(|\varphi| \cdot \text{td})} \cdot n^{\mathcal{O}(q + \ell)}$ and space polynomial in $n + |\varphi| + d_\varphi + p_\varphi$ and decides whether $G \models \varphi$ holds. The algorithm cannot give false positives and may give false negatives with probability at most $1/2$.*

Let us remark that to handle connectivity and acyclicity constraints we rely on the Cut&Count technique by Cygan et al. [17]. This technique works under the assumption that a solution (if it exists) is unique. The standard way of ensuring this is the classic Isolation lemma [30]: one samples the weights in a certain probabilistic way so that, with high probability, the solution becomes unique. The application of Isolation lemma is the source of both (1) randomization and (2) the increase in space complexity as storing the sampled weights requires polynomial space. As soon as the weights are fixed, our algorithm is deterministic and runs in logarithmic space similarly to Theorem 1.3.

Another aspect we would like to remark is the dependency on the input formula φ in the exponent of n . As stated in both theorems, the exponents of n depend on ℓ or $\ell + q$ where ℓ denotes the number of size measurements and q the number of acyclicity constraints in the formula φ . In all natural problems captured by our logic, the numbers of size measurements and acyclicity constraints are bounded by very small constants, like 1 or 2, and therefore, the degree of n is independent of the natural problem we solve.

Algorithmic meta-theorems and related logics. We would like to point out that our main result reads as an *algorithmic meta-theorem*, i.e., a general statement of the form “All problems with property P , restricted to a class of inputs I , can be solved efficiently”. The archetypal, and possibly most celebrated algorithmic meta-theorem is commonly known as *Courcelle’s theorem* [13–15] (see also [2, 11]) which states that all problems expressible in MSO_2 logic are linear time solvable on graphs of bounded treewidth.

Designing and proving algorithmic meta-theorems is a very active area of research bringing close topics related to logic, combinatorics (especially graphs), and algorithms; see [25, 26, 28, 36, 38] for related surveys. A remarkable feature of our meta-theorem is that it has a quite moderate run time dependence on the length of the input formula (assuming a constant number of size measurements). In stark contrast, the run time of model checking algorithms for FO and MSO often heavily depends on the formula length and the width parameter. Even non-elementary tower functions $2^{2^{\dots}}$, where the height depends on the length of the formula, are quite common and necessary [21]. The efficiency of our model checking algorithm is shared with the model checking algorithm for the *existential counting modal logic* (ECML for short)—introduced by Pilipczuk [35]—and the model checking algorithm for the *distance neighborhood logic* (DN for short)—introduced by Bergougnoux, Dreier and Jaffke [6]—whose model checking problem admits a single-exponential algorithm for treewidth and mim-width, respectively (both model checking algorithms can handle connectivity and acyclicity constraints).

In fact, $\text{NEO}_2[\text{FRec}]$, DN and ECML are closely related, let us elaborate on that. Our logic $\text{NEO}_2[\text{FRec}]$ relies on the same neighborhood operators mechanism as DN introduced by Bergougnoux, Dreier and Jaffke [6]. Let $\text{NEO}_1[\text{FcoF}]$ be the restriction of $\text{NEO}_2[\text{FRec}]$ without edge set terms and where every set $\sigma \subseteq \mathbb{N}$ is finite or co-finite (i.e., they forbid the “periodicity”). Then, DN is equivalent to the extension of $\text{NEO}_1[\text{FcoF}]$ with neighborhood operators $N_r^\sigma(U)$ with $r \in \mathbb{N}$ that evaluate as the σ -neighborhood of U in the r -th power of G .

In [6], the authors provide a model checking algorithm for DN running in time $n^{\mathcal{O}(d \cdot w \cdot |\varphi|^2)}$ where $|\varphi|$ is the size of the input formula, d is the maximum $d(\sigma)$ over the sets σ that occur in φ , and w is the mim-width of a given decomposition. The running time of their model checking algorithm is also upper bounded by efficient FPT functions in terms of the treewidth, cliquewidth, and rankwidth of the given decomposition.¹

¹Assuming that the input formula has a constant number of size measurements (e.g., $t \leq m$) involving a non-constant integer (e.g., m).

On the other hand, ECML is variant of modal logic which captures a good number of problems known to be tractable in time $2^{\mathcal{O}(\text{tw})}n^{\mathcal{O}(1)}$ where tw is the treewidth of the input graph. The evaluation of formulas in modal logic is tied to an active vertex that changes over time and is not explicitly represented by a variable of the formula. Roughly, an ECML formula is of the form $\exists \bar{X} \xi \wedge \forall v \psi$ where:

- $\exists \bar{X}$ is a sequence of existentially quantified vertex set and edge sets variables,
- ξ an arbitrary quantifier-free arithmetic formula over the parameters and the cardinalities of fixed and quantified sets,
- $\forall v$ is simply a way of initializing the active vertex, and
- ψ is a quantifier-free MSO_2 formula where we can use the central operators of ECML that are $\Box^\sigma$ and $\Diamond^\sigma$ with σ a finitely recognizable set.

When the active vertex is v , a formula $\Box^\sigma \varphi$ shall be read as “the number of neighbors of w that satisfy ψ as an active vertex belongs to σ ”, while $\Diamond^\sigma \varphi$ is equivalent to $\neg \Box^\sigma \neg \varphi$.

Somewhat surprisingly, the authors of [6] prove that ECML and DN are semantically close in the sense that the operators $\Box^\sigma$ and $\Diamond^\sigma$ are equivalent to our neighborhood operator $N^\sigma(\cdot)$. In fact, their arguments prove that ECML is equivalent to $\text{NEO}_2[\text{FRec}]$ (assuming that the arithmetic formulas are restricted to what we can do with our size measurements). Nonetheless, we introduce $\text{NEO}_2[\text{FRec}]$ because the neighborhood operators is a simpler mechanism to deal with in both the design of the model checking algorithm and expression of the problems.

Comparing our logic to CMSO_2 and the case of universal quantifiers. Let us briefly mention here how the expressive power of $\text{NEO}_2[\text{FRec}]$ compares to the expressive power of FO, MSO_2 , and CMSO_2 ; the latter being the enhancement of MSO_2 with modular counting predicates (i.e., predicates expressing that for a given set variable X and a fixed integer $p > 1$, $|X| \equiv 0 \pmod p$). Note that every formula φ of $\text{NEO}_2[\text{FRec}]$ can be rewritten to an equivalent formula φ' of CMSO_2 but this transformation comes with a size blow-up that depends on the maximum of values $d(\sigma)$, over all σ appearing in φ . Moreover, each size measurement $|X| \leq m$ can also be expressed in CMSO_2 , but the size blow-up will depend on m which can be prohibitively big, e.g., $O(n)$. From another viewpoint, while in bounded treedepth graphs the expressive power of MSO_2 and FO coincide [18], it is unclear to which variant of FO our logic $\text{NEO}_2[\text{FRec}]$ maps to. Also, we would like to stress that $\text{NEO}_2[\text{FRec}]$ is an extension of fully-existential MSO_2 that allows limited use of universal quantifiers, “hidden” in the neighborhood operators, the size measurements, and the comparison between terms. The extension of the logic with acyclicity, connectivity and clique constraints, captures even more problems and all these extra features allow further use of universal quantifiers in restricted forms. This extension, while allowing to encode more problems and get further algorithmic results, also investigates the potential of our techniques to the regime of universal quantification.

Organization of the paper In Section 2 we present the main ideas behind our algorithm on a concrete formula. Then in Section 3 we introduce notions and present some algebraic results used in the paper. In Section 4 we formally define our logic, provide an equivalent core fragment of this logic on which we focus in the remainder, and present some problems expressible in the logic. In Section 5 we define the partial solutions, indices and polynomials on which our algorithm relies on. In Section 6 we present the recursive equalities and in Section 7 we put them together to obtain our algorithm for $\text{NEO}_2[\text{FRec}]$. After that in Section 8 we explain how to extend the algorithm to

handle clique, acyclicity, and connectivity constraints. Finally, in Section 9 we conclude with some open questions.

2 Technical Overview

In this section, we briefly present our logic as well as the main techniques behind our model-checking algorithm. We first explain how our algorithm works with a concrete formula, and then, we briefly discuss what changes in the general case. So fix a graph G on n vertices and an elimination tree T of G of depth td and the following formula

$$\varphi = \exists Y \ N^\sigma(\mathbf{v}, Y) = \mathbf{v} \wedge |Y| \leq k$$

where σ is the set

$$\sigma = \{3\} \cup \{i : i \geq 4 \text{ and } i \text{ is even}\}.$$

By Definition 1.1, the set $N^\sigma(\mathbf{v}, Y)$ with $\mathbf{v} = V(G)$ is the set of all vertices v in G such that the number of edges in Y incident with v is at least 4 and even, or it is precisely 3. So the property $N^\sigma(\mathbf{v}, Y) = \mathbf{v}$ states that for every vertex v in G , the number of incident edges from Y is at least 4 and even, or it is precisely 3. Finally, $|Y| \leq k$ bounds the size of Y by k .

According to Definition 1.2 we have $d(\sigma) = 4$ and $p(\sigma) = 2$. Observe that, for any integer $i \geq 4$, whether i belongs to σ depends solely on its parity, i.e., rest modulo 2. The definition of σ implies that for every vertex v , two values are relevant in order to know whether $v \in N^\sigma(\mathbf{v}, Y)$. First, we are interested in the number of incident edges in Y *counting up to* $d(\sigma) = 4$, i.e., whether v has *precisely* i incident edges for some value $0 \leq i \leq d(\sigma) - 1 = 3$, or it has *at least* $d(\sigma) = 4$ such edges. We will use the so-called *neighborhood functions* for this, such a function maps every vertex to a value between 0 and $d(\sigma) = 4$. Second, it is useful to know the *parity* of the number of edges in Y incident to v . We will use the so-called *mod-degrees* for this, such a function maps every vertex to a parity from $\mathbb{F}_{p(\sigma)} = \{0, 1\}$.

Our goal is to develop a space- and time-efficient top-down recursion along T to check whether a set Y with desired properties exists. Each recursive call of our algorithm is associated with a node u of T , a set of edges $\text{sheaf}(u)$ and two sets of vertices: a *tail* and a *subtree*. The edge set of G is partitioned into certain sets $\text{sheaf}(w)$ among the leaves w of T in such a way that both end-points of every edge in $\text{sheaf}(w)$ belong to the path from w to the root (see Subsection 3.2 for a formal definition). Then, $\text{sheaf}(u)$ is defined as the union of $\text{sheaf}(w)$ over the leaves w of T in $\text{subtree}[u]$. The subtree of a recursive call represents the set of already *processed* vertices. On the other hand, the vertices of the tail represent *active* vertices that may be incident to *still unprocessed* edges outside $\text{sheaf}(u)$. With a node u , two types of recursive calls are associated, namely when u is treated as processed and when it is not processed. In the former case, the subtree of a recursive call associated with u is $\text{subtree}[u]$, namely the set of descendants of u in T including u , and it is $\text{subtree}(u) := \text{subtree}[u] \setminus \{u\}$ in the latter case. Similarly, its tail is $\text{tail}(u) := \text{tail}[u] \setminus \{u\}$, namely the set of proper ancestors of u in T , or $\text{tail}[u]$, respectively.

Every recursive call takes as input an *index* that stores the information about the tail, and each index is *compatible* with a set of *partial solutions*. We will gradually present the components of an index and a compatible partial solution in the context of a recursive call associated with a vertex u of T , a tail tail and a subtree subtree . The first component of a partial solution is a set of edges $\hat{\mathbf{g}}(Y) \subseteq \text{sheaf}(u)$ corresponding to a partial interpretation of the edge set variable Y .

Let us explain why an index I cannot simply use a function $h: \text{tail} \rightarrow \{0, 1, 2, 3, 4\}$ to keep track of the number (up to $d(\sigma) = 4$) of edges in $\hat{\mathbf{g}}(Y)$ incident to each vertex of tail . If u is an internal node of T with children $v_1, \dots, v_t$, then $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ is a partition of $\text{sheaf}(u)$, and for

each child v_i of u and each $v \in \text{tail}$, there are $d(\sigma) + 1$ options for the number (up to $d(\sigma) = 4$) of edges in $\hat{\mathbf{g}}(Y) \cap \text{sheaf}(v_i)$ incident to v . In this case, the recursive call associated with I would require to consider $(d(\sigma) + 1)^{|\text{tail}| \cdot t}$ combinations of indexes of the children of u . Since t can be in $\Theta(n)$, this is a prohibitively large for the running time we aim at. One might, instead, process the children of u one after another but then keeping track of all indexes of the children processed so far requires $\mathcal{O}((d(\sigma) + 1)^{\text{td}})$ space.

So we have to proceed differently in order to control the number of edges in Y incident to each vertex. For this, the second component of a partial solution is a neighborhood function $\hat{\mathbb{E}}: \text{tail} \cup \text{subtree} \rightarrow \{0, \dots, 4\}$, and each value $\hat{\mathbb{E}}(v)$ is an *expectation* on the number of edges of Y incident to v when the whole graph is processed, and not among the edges processed so far. For this reason, the value $\hat{\mathbb{E}}(v)$ remains fixed during subsequent recursive calls. And when the whole graph is processed, our algorithm ensures that we *filtered out* (via counting tricks) all partial solutions where there exists at least one vertex $v \in V(G)$ with $\hat{\mathbb{E}}(v) < d(\sigma)$ and the expectation $\hat{\mathbb{E}}(v)$ is not equal the number of edges in $\hat{\mathbf{g}}(Y)$ incident to v . Due to the special meaning of the value $d(\sigma) = 4$, for vertices v with $\hat{\mathbb{E}}(v) = d(\sigma)$ the filtering only ensures that after the whole graph is processed, there are at least edges in $\hat{\mathbf{g}}(Y)$ incident to v . Each index has a neighborhood function $\mathbb{E}: \text{tail} \rightarrow \{0, \dots, 4\}$ to keep track of the values $\hat{\mathbb{E}}(v)$ on the vertices $v \in \text{tail}$. Note that, in this concrete example, the domains of $\hat{\mathbb{E}}$ and $\mathbb{E}$ could be restricted to $\{3, 4\}$ because every vertex needs to be incident to exactly 3 or at least 4 edges in Y but we stick to the set $\{0, \dots, 4\}$ to remain consistent with our algorithm.

We first focus on ensuring that each $v \in \text{tail}$ has at least $\hat{\mathbb{E}}(v)$ incident edges when v is processed. For this, we use an auxiliary graph in which every vertex w of G is replaced by $d(\sigma) = 4$ copies $w^1, \dots, w^4$ of itself. Then if we choose, in a certain way, for every edge from Y exactly one copy of this edge in the auxiliary graph (such choices will be parts of our *partial solutions*), then the fact that v has at least $\hat{\mathbb{E}}(v)$ incident edges is equivalent to the fact that the first $\hat{\mathbb{E}}(v)$ copies of v are incident to at least one edge in this auxiliary graph. The following idea is based on *inclusion-exclusion* and generalizes a technique used in [32] for HAMILTONIAN CYCLE. To obtain the number of *partial solutions* in which the first $\hat{\mathbb{E}}(v)$ copies of v are incident to at least one edge:

- We first take the number of partial solutions where all of the first $\hat{\mathbb{E}}(v)$ copies of v are *allowed* to have incident edges.
- Then we subtract, for each $j \in \{1, \dots, \hat{\mathbb{E}}(v)\}$, the number of partial solutions where the j -th copy of v is *forbidden* to have incident edges. Observe that we subtracted the partial solutions that *forbid* edges incident with, say, the first and the second copy of v at least twice.
- So now for each pair $j_1 \neq j_2 \in \{1, \dots, \hat{\mathbb{E}}(v)\}$, we add back the number of partial solutions where both the j_1 -th and the j_2 -th copy are *forbidden* to have incident edges.
- The process is repeated for every possible choice of *not allowed* copies.

This motivates the second part of the index, namely, a neighborhood function $\text{alw}: \text{tail} \rightarrow \{0, \dots, 4\}$ which determines which copies of the vertices in tail are *allowed* to have incident edges. Finally, as already motivated above, the third component of the index is a mod-degree $\text{pd}: \text{tail} \rightarrow \mathbb{F}_{p(\sigma)} = \{0, 1\}$. Unlike the *expectation* $\mathbb{E}$, the value $\text{pd}(u)$ is the *real* parity of the number of edges in Y processed so far and incident with u , we will see below that an algebraic discrete Fourier transform allows us to handle mod-degrees efficiently. Altogether, an *index* for tail is a tuple $(\mathbb{E}, \text{alw}, \text{pd})$ where $\mathbb{E}$ and alw are neighborhood functions and pd is a mod-degree on tail .

Now we are ready to formally define partial solutions as well as the notion of *compatibility* between partial solutions and indexes. A *partial solution* for $\text{tail}(u)$ is a triple $(\hat{\mathbf{g}}(Y), \hat{\mathbb{E}}, \vec{W})$ where

- $\hat{\mathbf{g}}(Y)$ is a subset of edges from $\text{sheaf}(u)$.
- $\hat{\mathbb{E}}$ is a neighborhood function on $\text{tail}(u) \cup \text{subtree}[u]$.
- $\vec{\mathcal{W}}$ is a set of edges of the auxiliary graph such that for every pair $\{v, w\}$ of distinct vertices of G , there exists a pair $v^i w^j \in \vec{\mathcal{W}}$ if and only if we have $vw \in \hat{\mathbf{g}}(Y)$. Furthermore, we require that for every pair $\{v, w\}$, there exists at most one pair $v^i w^j \in \vec{\mathcal{W}}$. Informally speaking, $\vec{\mathcal{W}}$ *assigns* every edge from $\hat{\mathbf{g}}(Y)$ to precisely one edge between the copies of its endpoints.
- Additionally, we require the following for every *processed* vertex $v \in \text{subtree}[u]$. First, v has at least $\hat{\mathbb{E}}(v)$ incident edges in $\hat{\mathbf{g}}(Y)$ —this ensures that the *expectation* $\hat{\mathbb{E}}(v)$ is a lower bound on the number of incident edges in $\hat{\mathbf{g}}(Y)$. Secondly, either $\hat{\mathbb{E}}(v) = 3$, or $\hat{\mathbb{E}}(v) = 4$ and v has an even number of incident edges in $\hat{\mathbf{g}}(Y)$ —i.e., assuming that the expectation is correct, the vertex v has a number of incident edges in $\hat{\mathbf{g}}(Y)$ from σ as desired.

Let now $I := (\mathbb{E}, \text{alw}, \text{pd})$ be an index of $\text{tail}(u)$. We say that the partial solution $P := (\hat{\mathbf{g}}(Y), \hat{\mathbb{E}}, \vec{\mathcal{W}})$ of u is compatible with I if the following properties are satisfied:

- $\hat{\mathbb{E}}$ and $\mathbb{E}$ are equal on $\text{tail}(u)$.
- For every $v \in \text{tail}(u)$ and every edge $v^i w^j \in \vec{\mathcal{W}}$, we have $i \leq \text{alw}(v)$, i.e., there are only edges incident to *allowed* copies.
- For every vertex $v \in \text{tail}(u)$, the number $\text{pd}(v)$ is equal, modulo $p(\sigma) = 2$, to the total number of edges in $\vec{\mathcal{W}}$ incident to the copies of v . Observe that by definition of a partial solution, this is simply the parity of the number of edges in $\hat{\mathbf{g}}(Y)$ incident to v .

To keep track of the size of $\vec{\mathcal{W}}$, we define the value $c_u(P) = |\vec{\mathcal{W}}|$. Furthermore, we define two following values:

$$a_u(P) = \sum_{\substack{v \in \text{subtree}[u]: \\ \hat{\mathbb{E}}(v) < d(\sigma) = 4}} \hat{\mathbb{E}}(v) \text{ and } b_u(P) = \sum_{\substack{v \in \text{tail}(u) \cup \text{subtree}[u]: \\ \hat{\mathbb{E}}(v) < d(\sigma) = 4}} |\{\{v^i, w^j\} \in \vec{\mathcal{W}}\}|.$$

The following observation is crucial. The definition of a partial solution ensures that every vertex in $\text{subtree}[u]$ has at least $\hat{\mathbb{E}}(u)$ incident edges in $\hat{\mathbf{g}}(Y)$. Also recall that the edges in $\hat{\mathbf{g}}(Y)$ are in bijection with the edges in $\vec{\mathcal{W}}$, and the bijection maps the endpoints of every edge to their copies. Therefore, we always have $a_u(P) \leq b_u(P)$. Furthermore, the equality is achieved precisely if every vertex $v \in \text{subtree}[u]$ with $\hat{\mathbb{E}}(v) < d(\sigma) = 4$ has *exactly* $\hat{\mathbb{E}}(v)$ incident edges in $\hat{\mathbf{g}}(Y)$. Therefore, keeping track of the values $a_u(P)$ and $b_u(P)$ permits us to keep the partial solutions for which $\hat{\mathbb{E}}$ correctly guesses the number of incident edges (except for the special value $d(\sigma) = 4$ for which we only ensure the existence of at least $d(\sigma) = 4$ incident edges). We now define the polynomial

$$\mathcal{P}_u(I) = \sum_P \alpha^{a_u(P)} \beta^{b_u(P)} \gamma^{c_u(P)}$$

where the sum is over all partial solutions of $\text{tail}(u)$ compatible with I . Crucially, computing the polynomial $\mathcal{P}_r(\emptyset)$ for the root r of T suffices to solve the problem. Indeed, we have $G \models \varphi$ if and only if there is a monomial $q \cdot \alpha^i \beta^j \gamma^k$ in $\mathcal{P}_r(\emptyset)$ with $j \leq k$ and such that for the coefficient q we have $q \neq 0$. Here choosing the same exponent of α and β ensures, as argued above, that the *expectations* made by $\hat{\mathbb{E}}$ reflect the correct number of incident edges in $\hat{\mathbf{g}}(Y)$ when it is strictly smaller than $d(\sigma)$. And the exponent of γ reflects the cardinality of $\hat{\mathbf{g}}(Y)$ so we only allow the sets of size $j \leq k$.

We can similarly define partial solutions and indexes for $\text{tail}(u)$ by replacing $\text{tail}(u)$ by $\text{tail}[u]$ and $\text{subtree}[u]$ by $\text{subtree}(u)$, respectively, in the definitions above, we then denote the corresponding polynomials by $\mathcal{P}_u[\cdot]$.

Let us now explain the recursive computation of this root polynomial. The following equality holds for every (non-leaf) node u

$$\mathcal{P}_u[I] = \sum_{\substack{\text{pd}_1, \dots, \text{pd}_t: \text{tail}[u] \rightarrow \{0,1\}: \\ \text{pd}_1 + \dots + \text{pd}_t \equiv 2\text{pd}}} \prod_{i=1}^t \mathcal{P}_{v_i}(\mathbb{E}, \text{alw}, \text{pd}_i)$$

if $v_1, \dots, v_t$ denote the children of u in the elimination tree T . Vaguely speaking, this holds mainly for the following reasons: (1) the sets $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$ partition the set $\text{subtree}(u)$, (2) the sets $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ partition the set $\text{sheaf}(u)$, (3) to obtain degree (modulo 2) of a certain vertex in $\text{tail}[u]$ in the edge set $\hat{\mathbf{g}}(Y)$, we can simply sum up its degrees restricted to the edges sets $\hat{\mathbf{g}}(Y) \cap \text{sheaf}(v_1), \dots, \hat{\mathbf{g}}(Y) \cap \text{sheaf}(v_t)$, and (4) when multiplying the polynomials, their degrees are added so that the values $a_u[\cdot]$, $b_u[\cdot]$, and $c_u[\cdot]$ are computed correctly. The above formula has the problem that it requires to branch over $2^{|\text{tail}[u]| \cdot t}$ options with t possibly in $\Theta(n)$. But we can show that for the discrete Fourier transforms $\mathcal{Q}_u^p[\cdot]$ and $\mathcal{Q}_{v_i}^p(\cdot)$ of $\mathcal{P}_u[\cdot]$ and $\mathcal{P}_{v_i}(\cdot)$, respectively, (see Subsection 3.3 and Definition 6.3 for a formal definition) we have

$$\mathcal{Q}_u^p[I] \equiv_p \prod_{i=1}^t \mathcal{Q}_{v_i}^p(I)$$

for a prime number p satisfying certain technical preconditions. For this reason, our algorithm works with discrete Fourier transforms and thus quickly combines the children of u without any branching.

To compute the polynomial $\mathcal{P}_u(I)$ from some polynomials of form $\mathcal{P}_u[\cdot]$, we use the following equality:

$$\begin{aligned} \mathcal{P}_u(\mathbb{E}, \text{alw}, \text{pd}) &= \sum_{0 \leq i \leq 3} S_{\pm}(3, i) \cdot \mathcal{P}_u[\mathbb{E}[u \mapsto 3], \text{alw}[u \mapsto i], \text{pd}[u \mapsto 1]] \cdot \alpha^3 + \\ &\quad \sum_{0 \leq i \leq 4} S_{\pm}(4, i) \cdot \mathcal{P}_u[\mathbb{E}[u \mapsto 4], \text{alw}[u \mapsto i], \text{pd}[u \mapsto 0]]. \end{aligned}$$

Let us sketch what all of those values are and why the formula is correct. For an integer i and a function $f \in \{\mathbb{E}, \text{alw}, \text{pd}\}$, the function $f[u \mapsto i]$ has the domain $\text{tail}[u]$, agrees with f on $\text{tail}(u)$, and maps u to i . Since the vertex u moves to the *processed* part, by definition of a partial solution we want to ensure that either (1) u has precisely $j = 3$ incident edges in $\hat{\mathbf{g}}(Y)$, in particular this number is odd, (this corresponds to the first summand); or (2) u has at least $j = 4$ incident edges in $\hat{\mathbf{g}}(Y)$ and this number is even (this corresponds to the second summand). And in both cases we need to ensure that u has at least j incident edges in $\hat{\mathbf{g}}(Y)$, i.e., each of the j copies of u is incident to a least one edge in $\vec{\mathcal{W}}$. The function $S_{\pm}(j, i)$ is the coefficient used to apply the inclusion-exclusion sketched earlier: to ensure that all j copies of u are incident to an edge in $\vec{\mathcal{W}}$, we *allow* edges to be incident to various subsets of the copies of u , and properly combine those options via the coefficients $S_{\pm}(k, i)$. Finally, the first summand is multiplied by α^3 to correctly update the value $a_u(\cdot)$, i.e., reflect that u now belongs to the subtree and $\hat{\mathbb{E}}(u) = 3 < 4 = d(\sigma)$. In terms of running time, the branching factor of this computation is only constant.

The last missing piece for the recursive algorithm is the non-recursive computation of the polynomial $\mathcal{P}_u[I]$ for some leaf u . We do not provide the details here due to their technicality but

we can show that for the discrete Fourier transformed polynomial $\mathcal{Q}_u^p[I]$, the computation can be carried out in polynomial time by considering every edge independently and correctly initializing the exponents of α , β , and γ depending on whether the edge is taken into $\hat{\mathbf{g}}(Y)$. Altogether, the ideas presented so far suffice to obtain a recursive algorithm with only a constant branching factor and the recursion depth bounded by the depth td of the elimination tree T , i.e., with the running time of $2^{\mathcal{O}(\text{td})}n^{\mathcal{O}(1)}$. At every level of recursion, we keep track of at most $n^{\mathcal{O}(1)}$ coefficients of polynomials (this follows from the upper-bounds on the degrees of the formal variables α , β , and γ) each of length upper-bounded by $n^{\mathcal{O}(1)}$ so the space complexity is polynomial.

General case We now briefly summarize which of the techniques from the paper were not covered by this example.

1. As done in [6], we define a simpler fragment $\text{coreNEO}_2[\text{FRec}]+\text{ACK}$ of $\text{NEO}_2[\text{FRec}]+\text{ACK}$ and show that it has the same expressive power. Moreover, we observe that the conversion of any $\text{NEO}_2[\text{FRec}]+\text{ACK}$ formula to a $\text{coreNEO}_2[\text{FRec}]+\text{ACK}$ formula can be done with only a linear increase of the formula length, increasing the number of size measurements by at most two, and without altering the values d_φ and p_φ . In particular, $\text{coreNEO}_2[\text{FRec}]+\text{ACK}$ has no nested neighborhood operators, in fact, all of them are of the form $N^\sigma(X, Y)$ where X and Y are a vertex set variable and an edge set variable, respectively.
2. The neighborhood operators with various sets σ can occur in the formula. We show that we can deal with all of them by counting the number of neighbors (1) up to d_φ defined as the maximum of all $d(\sigma)$ and (2) modulo p_φ where p_φ is defined as the least common multiple of all $p(\sigma)$. Further, we use “private” functions $\mathbb{E}$, pd , alw for every pair (X, Y) such that $N^\sigma(X, Y)$ occurs in the formula for some $\sigma \subseteq \mathbb{N}$.
3. To keep track of vertex set variables we extend the index by a so-called *vertex interpretation*, usually denoted by $\mathbf{f}$, which stores, for every vertex on the tail, to which vertex set variables it belongs.
4. Dealing with neighborhood operators where the first operand is a variable creates asymmetries: whether a vertex u belongs to $N^\sigma(X, Y)$ depends only on the edges $uv \in Y$ where $v \in X$. To deal with these asymmetries, the auxiliary graph becomes an auxiliary digraph and $\vec{W}$ becomes a function mapping each pair (X, Y) such that $N^\sigma(X, Y)$ occurs in the formula to a subset of arcs of this digraph.
5. The formula, in general, might be not in conjunctive normal form (unlike our example), and some of the vertex or edge set equalities might, in general, not hold in a solution. So we extend the index by a subset $\mathbf{F}$ of set equalities occurring in the formula, it then determines which equalities are *allowed to be falsified*, and this value remains fixed during the recursive calls. We then use another inclusion-exclusion to obtain partial solutions where *precisely* the equalities from $\mathbf{F}$ are falsified.
6. The formula might contain multiple size measurements, we introduce one formal variable γ_i per measured variable X_i .
7. To decrease the space complexity to logarithmic in n , we make use of the technique of [37] based on the Chinese remainder theorem. This technique reconstructs a polynomial by evaluating it in many points modulo certain primes of length logarithmic in n .
8. To handle connectivity and acyclicity constraints, we employ Cut&Count by Cygan et al. [17].

3 Preliminaries

We use the convention that 0 belongs to $\mathbb{N}$ and by $\mathbb{N}^+$ we denote the set $\mathbb{N} \setminus \{0\}$. For a non-negative integer $i \in \mathbb{N}$, by $[i]$ we denote the set $\{j \in \mathbb{N}^+ : j \leq i\}$, in particular, it holds that $[0] = \emptyset$. By $[i]_0$ we then denote the set $[i] \cup \{0\}$. We will also use the Iverson bracket notation: for Boolean predicate p , the value $[p]$ is equal to 1 if p is true and equal to zero otherwise.

3.1 Graphs

We consider finite, simple, vertex and edge colored graphs G where $V(G)$ and $E(G)$ denote the vertex and edge set, respectively. We assume that a color is used either on vertices or edges. For a graph G and color class P , we denote the set of elements in G with color P by $P(G)$. Note that color classes may be empty and that a vertex or an edge may have zero, one, or multiple colors. We denote by n the order of $V(G)$, where the corresponding graph is always clear from the context. For every vertex set $A \subseteq V(G)$ we denote by $\overline{A}$ the set $V(G) \setminus A$.

The subgraph of G *induced* by a subset $A \subseteq V(G)$, whose vertex set is A and whose edge set is $\{uv \in E(G) : u, v \in A\}$, is denoted by $G[A]$.

3.2 Elimination trees and treedepth

Let $G = (V, E)$ be a graph. A forest T is an *elimination forest* of G if the vertex set of T coincides with V and for every edge $uv \in E$, either u is an ancestor of v , or v is an ancestor of u in T . We refer to the vertices of T as *nodes*. The *depth* of T is the maximum number of vertices on any root-to-leaf path. And the *treedepth* of G is the minimum depth over all elimination forests of G . Every node in a tree of depth at most d has at most d ancestors so the following well-known property holds.

Observation 3.1. *If a graph on n vertices admits an elimination forest of depth d , this graph has at most $d \cdot n$ edges.*

Along this work, we assume that the elimination forest consists of a single tree (and we then refer to it as *elimination tree*) as we can choose one dedicated root in an elimination forest and attach the roots of the remaining trees as its new children: the arising rooted tree is then elimination tree and its depth is at most one larger than the depth of the forest we started with. The elimination tree of a graph will be fixed and we will use $\mathbf{td}$ to denote its depth.

Let T be a elimination tree of a graph G . Following the standard terminology (see [27, 32, 37]) for algorithms on elimination trees, we define the following sets for every node u of T :

$$\begin{aligned} \text{subtree}[u] &:= \{v : v \text{ is descendant of } u\}, & \text{subtree}(u) &:= \text{subtree}[u] \setminus \{u\}, \\ \text{tail}[u] &:= \{v : v \text{ is ancestor of } u\}, & \text{tail}(u) &:= \text{tail}[u] \setminus \{u\}, \\ \text{broom}[u] &:= \text{tail}[u] \cup \text{subtree}[u]. \end{aligned}$$

We remark that we assume the convention that u is both an ancestor and a descendant of itself, i.e., both $u \in \text{tail}[u]$ and $u \in \text{subtree}[u]$ hold. We fix an arbitrary left-to-right ordering of the children of every node. Similarly to [32] we define the following. For a node u by $\text{left}(u)$ we denote the leftmost leaf descendant of u , that is, the leaf obtained by starting at u and iteratively moving to the leftmost child of the current vertex until a leaf is reached. For a leaf w of T , we define

$$\text{sheaf}(w) := \{ab : ab \in E(G), a \text{ is an ancestor of } b, \text{left}(b) = w\}$$

and finally, for a node u we define

$$\text{sheaf}(u) := \bigcup_{w \in \text{subtree}[u], w \text{ is a leaf of } T} \text{sheaf}(w).$$

See Figure 1 for an illustration.




Figure 1: Illustration of an elimination tree of a graph with two distinguished leaves w and w' ; the edges of the tree are depicted in light grey. The vertices b with $\text{left}(b) = w$ are depicted in blue. The (thick) blue edges are the edges of $\text{sheaf}(w)$ and the (thick) magenta edges are the edges of $\text{sheaf}(w_2)$.

First, observe that $(\text{sheaf}(w))_{w \text{ is a leaf}}$ is a partition of $E(G)$. Second, for every non-leaf node u with children $v_1, \dots, v_t$, the sets $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ partition $\text{sheaf}(u)$ and the sets $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$ partition $\text{subtree}(u)$. And finally, the following holds:

Observation 3.2. *For every node u of T , we have $V(\text{sheaf}(u)) \subseteq \text{tail}[u] \cup \text{subtree}[u]$.*

The notion of partial solutions we use in our proofs is defined with respect to a triple $\text{tss} = (\text{tail}, \text{subtree}, \text{sheaf})$ —called a *triad* (in G)—where tail and subtree are disjoint subsets of $V(G)$ and sheaf is a subset of $E(G[\text{tail} \cup \text{subtree}])$. The elements of this triad are then referred to as respectively the *tail*, the *subtree*, and the *sheaf* of tss . We define the *broom* of the triad tss as the set $\text{tail} \cup \text{subtree}$. In our model checking algorithm, we deal with the triads that are naturally induced by the elimination tree. For every node u of T , we define the triads $\text{tss}[u] = (\text{tail}[u], \text{subtree}(u), \text{sheaf}(u))$ and $\text{tss}(u) = (\text{tail}(u), \text{subtree}[u], \text{sheaf}(u))$. Observe that $\text{tss}[u]$ and $\text{tss}(u)$ have the same sheaf and the same broom which we refer to as $\text{broom}[u]$. Notice that, for the root r of T , we have $\text{tss}(r) = (\emptyset, V(G), E(G))$.

3.3 Discrete Fourier transforms

For integers a, b, c with $c > 0$ we write $a \equiv_c b$ to denote that two integers a and b leave the same rest when divided by c , i.e., are equal modulo c . Also if $a \in [c-1]_0$, we sometimes write $a = b \pmod{c}$ to denote that a is the rest of b modulo c . We naturally extend this notation to polynomials and functions. For a positive integer r , by $\mathbb{Z}_r$ we denote the ring of integers modulo r . For a prime number p , by $\mathbb{F}_p$ we denote the field of integers modulo p , and by $\mathbb{F}_p^*$ we denote the multiplicative group of this field. If $\mathbb{F}_p^*$ contains a number ω such that $\omega^r \equiv_p 1$ and $\omega^j \not\equiv_p 1$ for all $1 \leq j < r$, then ω is called the r -th root of unity. The r -th root of unity of $\mathbb{F}_p^*$ does not necessarily exist and it is, in general, not unique.

For the remainder of this subsection we fix a finite set D , a positive integer r , and a prime number p such that in $\mathbb{F}_p^*$ the r -th root of unity, denoted by $\omega = \omega_{r,p}$, exists. If multiple such roots of unity exist, we fix one arbitrary (it will be the smallest one in our applications). For a mapping $h: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ the *discrete Fourier transform* (also called *DFT*) of h $\text{DFT}_{D,r,p}(h): (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ and

the *inverse discrete Fourier transform* (also called *inverse DFT*) of h $\text{DFT}_{D,r,p}^{-1}(h): (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ are defined as

$$\text{DFT}_{D,r,p}(h)(y) \equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{\sum_{z \in D} q(z) \cdot y(z)} \cdot h(q).$$

and

$$\text{DFT}_{D,r,p}^{-1}(h)(y) \equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{\sum_{z \in D} -q(z) \cdot y(z)} \cdot h(q).$$

for every $y: D \rightarrow \mathbb{Z}_r$, here $\frac{1}{r^{|D|}}$ denotes the inverse of $r^{|D|}$ in $\mathbb{F}_p^*$. For $y, q: D \rightarrow \mathbb{Z}_r$ we use the notation $q \cdot y$ for $\sum_{z \in D} q(z) \cdot y(z)$. With these shortcuts we have

$$\text{DFT}_{D,r,p}(h)(y) \equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} \cdot h(q) \text{ and } \text{DFT}_{D,r,p}^{-1}(h)(y) \equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} \cdot h(q).$$

The following two classical results will be useful for us. We refer, for example, to the survey by van Rooij [40] for the proofs and applications in the area of fast algorithms on tree decompositions. First, the inverse DFT is indeed the inverse of the DFT:

Theorem 3.3. *It holds that $\text{DFT}_{D,r,p}^{-1}(\text{DFT}_{D,r,p}(h)) \equiv_p h$ and $\text{DFT}_{D,r,p}(\text{DFT}_{D,r,p}^{-1}(h)) \equiv_p h$.*

For $y, q: D \rightarrow \mathbb{Z}_r$ we define $y + q: D \rightarrow \mathbb{Z}_r$ as $(y + q)(d) \equiv_r y(d) + q(d)$ for every $d \in D$. Further, for $h, g: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ we define $h \cdot g: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ via $(f \cdot g)(y) \equiv_p h(y) \cdot g(y)$ for every $y \in D \rightarrow \mathbb{Z}_r$. Then the following holds

Theorem 3.4. *Let $t \in \mathbb{N}$, let $h^1, \dots, h^t: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$, and let $y: D \rightarrow \mathbb{Z}_r$. Then it holds that*

$$\sum_{\substack{y^1, \dots, y^t: D \rightarrow \mathbb{Z}_r \\ y^1 + \dots + y^t \equiv_r y}} \prod_{i=1}^t h^i(y^i) \equiv_p (\text{DFT}_{D,r,p}^{-1} \left(\prod_{i=1}^t \text{DFT}_{D,r,p}(h^i) \right))(y),$$

Now we extend the notions and results to the case where the codomain of h is not $\mathbb{F}_p$ but rather the set of polynomials with coefficients from $\mathbb{F}_p$. Let $\ell \in \mathbb{N}$ and let $\gamma_1, \dots, \gamma_\ell$ be pairwise distinct formal variables. The set $\mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$ consists of all polynomials in formal variables $\gamma_1, \dots, \gamma_\ell$ with coefficients from $\mathbb{F}_p$. For a mapping $h: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$, we define the *DFT* of h $\text{DFT}_{D,r,p}(h): (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$ and the *inverse DFT* of h $\text{DFT}_{D,r,p}^{-1}(h): (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$ as

$$\text{DFT}_{D,r,p}(h)(y) \equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} \cdot h(q) \text{ and } \text{DFT}_{D,r,p}^{-1}(h)(y) \equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} \cdot h(q)$$

for all $y: D \rightarrow \mathbb{Z}_r$. Observe that in the case $\ell = 0$ the codomain of h is simply $\mathbb{F}_p$ so we indeed extend the above case. Thanks to the linearity of DFT and its inverse, Theorems 3.3 and 3.4 hold for polynomials as well. A formal proof of this claim is available in Section A.

3.4 Inclusion-exclusion

We will also make use of the folklore inclusion-exclusion principle (we refer to [16] for a proof and applications). In particular, we will rely on its intersection version:

Theorem 3.5 ([16], Theorem 10.2). *Let U and Y be finite sets and let $A_y \subseteq U$ for every $y \in Y$. Then it holds that:*

$$\left| \bigcap_{y \in Y} A_y \right| = \sum_{X \subseteq Y} (-1)^{|X|} \cdot |U \setminus (\cup_{x \in X} A_x)|.$$

4 Neighborhood Operator Logics

In the following, we introduce the syntax and semantics of NEO logics. As explained in the introduction, these logics are based on the DN logic introduced in [6]. These logics are obtained by extending fully-existential monadic second-order (MSO_2) logic. Remember that MSO_2 allows quantification over vertices, edges, sets of vertices and sets of edges; together with an adjacency relation $E(\cdot, \cdot)$ and equality relation $=$ between vertices, unary vertex relations (i.e., colors) as well as the containment relation $\in$ of vertices/edges in sets. Fully-existential MSO_2 is the restriction of MSO_2 to existential (first-order and second-order) quantifiers, while furthermore requiring that only quantifier-free formulas may be negated. Set variables are denoted by upper-case letters ($X, Y, Z, \dots$) while unary and binary relations (or colors) are denoted by typewriter letters ($P, Q, R, \dots$). The letters V and E are special unary and binary relations reserved for the set of all vertices and all edges, respectively.

Syntax. We first define *edge set terms* using the following rules:

1. Every edge set variable Y is an edge set term.
2. Every binary relational symbol P is an edge set term.
3. $\emptyset$ and E are edge set terms.
4. If t_1 and t_2 are edge set terms then $\overline{t_1}$, $t_1 \cap t_2$, $t_1 \cup t_2$ and $t_1 \setminus t_2$ are also edge set terms.

Secondly, we define *neighborhood terms* using the following rules:

1. Every vertex set variable X is a neighborhood term.
2. Every unary relational symbol U is a neighborhood term.
3. $\emptyset$ and V are neighborhood terms.
4. $N^\sigma(t)$, $N^\sigma(t, Y)$ and $N^\sigma(t, F)$ are neighborhood terms for every $\sigma \subseteq \mathbb{N}$, neighborhood term t , edge set variable Y and binary relational symbol F .
5. If t_1 and t_2 are neighborhood terms then $\overline{t_1}$, $t_1 \cap t_2$, $t_1 \cup t_2$ and $t_1 \setminus t_2$ are also neighborhood terms.

Then NEO_2 is the extension of fully-existential MSO_2 by the following two rules:

6. If t is a neighborhood term or an edge set term and $m \in \mathbb{N}$, then $|t| = m$, $|t| \leq m$ and $|t| \geq m$ are formulas called *size measurements*.
7. If t_1 and t_2 are both neighborhood terms or edge set terms, then $t_1 = t_2$, $t_1 \subseteq t_2$ and $t_1 \supseteq t_2$ are formulas.

We define NEO_1 as the restriction of NEO_2 where we forbid variables for edges and edge sets. Given $i \in \{1, 2\}$, and a family F of subsets of $\mathbb{N}$, we denote by $\text{NEO}_i[F]$ the restriction of NEO_i such that, for all neighborhood operators $N^\sigma(\cdot)$ we have $\sigma \in F$. In this paper, we consider the following families of subsets of $\mathbb{N}$:

- FcoF consists of all the non-empty subsets of $\mathbb{N}$ that are finite or co-finite.
- FRec consists of all the subsets σ of $\mathbb{N}$ that are *finitely recognizable* (see Definition 1.2).

On top of this, given a fragment of $\mathbf{L}$ of $\mathbf{NEO}_2$, we define $\mathbf{L+AC}$ as the extension of $\mathbf{L}$ with the following additional rule:

8. If t is a neighborhood term or an edge term, then $\text{con}(t)$ and $\text{acy}(t)$ are formulas.

Similarly, we define $\mathbf{L+K}$ as the extension of $\mathbf{L}$ with the following additional rule:

9. If t is a neighborhood term, then $\text{clique}(t)$ is a formula.

And we define $\mathbf{L+ACK}$ as the extension of $\mathbf{L}$ with the rules 8. and 9.

Let ξ be a formula or a term of $\mathbf{NEO}_2+\mathbf{ACK}$. We denote the set of variables that occur in ξ by the ordered tuple $\text{var}(\xi)$ and we denote by $\text{free}(\xi)$ the set of free variables of ξ . We denote by $\text{var}^V(\xi)$ the set of variables for vertices and vertex sets, and we denote by $\text{var}^E(\xi)$ the set of variables for edges and edge sets. Moreover, we denote by $\text{color}^V(\xi)$ and $\text{color}^E(\xi)$ the sets of unary and respectively binary relational symbols that occur in ξ . We define the *length* $|\varphi|$ of a $\mathbf{NEO}_2+\mathbf{ACK}$ formula φ to be the number of symbols of φ . Note that every number or set (as occurring for example in a size measurement $t \leq m$ or in a superscript of a neighborhood term $N^\sigma(\cdot)$) is *one* symbol.

Semantics. Next, we define the semantics of $\mathbf{NEO}$ logic. We consider *vertex and edge colored* graphs. This means, each vertex (respectively edge) of a graph may be in zero, one or more unary (resp. binary) relations $(P, Q, R, \dots)$. Our notion of interpretation of a logical formula uses two functions, one for interpreting the variables for vertices and vertex sets, and another for the variables for edges and edge sets. We do so to improve the readability of the model checking section.

Definition 4.1. *Given a graph G and $A \subseteq V(G)$. A vertex interpretation of a formula φ on A is a function mapping each vertex variable of φ to a vertex in A and each vertex set variable of φ to a subset of A . Given $F \subseteq E(G)$, an edge interpretation of φ on F is a function mapping each edge variable of φ to an edge in F and each edge set variable of φ to a subset of F .*

An *interpretation* of a formula φ is a tuple $(G, \mathbf{f}, \mathbf{g})$ consisting of a graph G , a vertex interpretation $\mathbf{f}$ of φ on $V(G)$ (mapping every vertex set variable to a subset of $V(G)$, and every vertex variable to a vertex of G) and an edge interpretation $\mathbf{g}$ of φ on $E(G)$ (mapping every edge set variable to a subset of $E(G)$, and every edge variable to an edge of G). Given an interpretation $(G, \mathbf{f}, \mathbf{g})$, we define the semantics of neighborhood and edge set terms.

1. $\llbracket X \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = \mathbf{f}(X)$ if $X \in \text{var}^V(\varphi)$ or $\llbracket X \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = \mathbf{g}(X)$ if $X \in \text{var}^E(\varphi)$,
2. $\llbracket P \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = P(G)$,
3. $\llbracket \emptyset \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = \emptyset$, $\llbracket V \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = V(G)$, and $\llbracket E \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = E(G)$
4. $\llbracket N^\sigma(t, \star) \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = N^\sigma(\llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})}, \llbracket \star \rrbracket^{(G, \mathbf{f}, \mathbf{g})})$, and $\llbracket N^\sigma(t) \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = N^\sigma(\llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})})$, where the neighborhood operators on the right are evaluated in G (see Definition 1.1),
5. $\llbracket \bar{t} \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = V(G) \setminus \llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})}$, $\llbracket t_1 \star t_2 \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = \llbracket t_1 \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \star \llbracket t_2 \rrbracket^{(G, \mathbf{f}, \mathbf{g})}$ for $\star \in \{\cap, \cup, \setminus\}$,

$\mathbf{NEO}_2$ inherits the semantics from $\mathbf{MSO}_1$, with the following semantics of the additional rules.

6. $\llbracket |t| < m \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = 1$ if $|\llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})}| < m$ and $\llbracket |t| < m \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = 0$ otherwise, for $< \in \{=, \leq, \geq\}$,

7. $\llbracket t_1 \prec t_2 \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = 1$ if $\llbracket t_1 \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \prec \llbracket t_2 \rrbracket^{(G, \mathbf{f}, \mathbf{g})}$ and $\llbracket t_1 \prec t_2 \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = 0$ otherwise, for $\prec \in \{=, \subseteq, \supseteq\}$.

For NEO+ACK, the semantics of the additional rule are as follows.

8. $\llbracket \text{con}(t) \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = [\text{the subgraph of } G \text{ induced by } \llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \text{ is connected}]$,
 $\llbracket \text{acy}(t) \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = [\text{the subgraph of } G \text{ induced by } \llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \text{ is acyclic}]$,
 $\llbracket \text{clique}(t) \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = [\text{the subgraph of } G \text{ induced by } \llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \text{ forms a clique}]$.

4.1 Problems captured by our logics

We provide here some examples of problems that are captured by NEO₂[FRec]+ACK and not yet mentioned in the introduction including SAT and MAX-SAT on CNF formulas.

As demonstrated in the introduction, we can express any (σ, ρ) -set property—as introduced by Telle and Proskurowski [39]—where σ and ρ are finitely recognizable sets. With the same arguments as used in [6] for the DN logic, we can generalize the latter formula to express the more general *locally checkable vertex partitioning* problems from [39], that ask, given a $q \times q$ matrix D whose entries are finitely recognizable sets, for a partition $(X_1, \dots, X_q)$ of $V(G)$ such that for every $i, j \in [q]$, we have $X_i \subseteq N^{D[i, j]}(X_j)$.

Now, we show that we can express SAT via the *signed incidence graph* of a CNF formula. Given a CNF formula ξ , we define its signed incidence graph G_ξ as follows. The vertices of G_ξ are the variables and the clauses of ξ , and the latter are colored with $\mathbf{C}$. The edges of G_ξ are defined as follows. For every variable x and every clause C of ξ , we have

- If x appears positively in C , then xC is an edge of G_ξ colored with $\mathbf{P}$.
- If x appears negatively in C , then xC is an edge of G_ξ colored with $\mathbf{N}$.

Then, we can express SAT, i.e., the existence of a model for ξ with the following formula

$$\exists X \subseteq \bar{\mathbf{C}} \wedge \mathbf{C} = N^{\mathbf{N}^+}(X, \mathbf{P}) \cup N^{\mathbf{N}^+}(\bar{X}, \mathbf{N}).$$

Basically, in this formula X is the set of variables that a model of ξ sets to true, and every clause needs to be adjacent to either a variable in X via an edge labeled $\mathbf{P}$ or a variable not in X via an edge labeled $\mathbf{N}$. Moreover, we can express MAX-SAT, i.e., the existence of an interpretation of ξ that satisfies at least k clauses with the following formula

$$\exists X \subseteq \bar{\mathbf{C}} \wedge |N^{\mathbf{N}^+}(X, \mathbf{P}) \cup N^{\mathbf{N}^+}(\bar{X}, \mathbf{N})| \geq k.$$

4.2 Core logic

Some of the operations in our logic can be understood as “syntactic sugar”, that does not increase the expressiveness, but merely reduces some friction when expressing problems. To facilitate our proofs, we consider a smaller fragment **coreNEO₂** of NEO₂, which we call *core neighborhood operator logic*, that has the same expressive power as NEO₂, and a similar equivalent fragment **coreNEO₂+AC** for NEO₂ +AC. Similarly, we also consider the fragments **coreNEO₂+K** and **coreNEO₂+ACK**.

Definition of core logics. We first define *primitive formulas*.

1. If X is a vertex set variable and $\mathbf{U}$ is a unary relational symbol, $\emptyset$, or $\mathbf{V}$, then $\mathbf{U} = X$ is a *primitive formula*.

2. If Y is an edge set variable and B is a binary relational symbol, $\emptyset$, or E then $B = Y$ is a *primitive formula*.
3. If X, Y and Z are vertex set variables, then $X = Y$, $\overline{X} = Y$, $X \cup Y = Z$, and $X \cap Y = Z$ are *primitive formulas*.
4. If X and W are vertex set variables, Y an edge set variable and, $\sigma \subseteq \mathbb{N}$, then $N^\sigma(X, Y) = W$ is a *primitive formula*.
5. If X is a vertex or edge set variable and $m \in \mathbb{N}$ then $|X| \leq m$, $|X| \geq m$, and $|X| = m$ are *primitive formulas*.
6. If X is a vertex or edge set variable then $\text{acy}(X)$ and $\text{con}(X)$ are *primitive formulas*.
7. If X is a vertex set variable then $\text{clique}(X)$ is a *primitive formula*.

Let coreNEO_2 be the fragment of NEO_2 containing all formulas of the form $\exists X_1 \dots \exists X_k \psi$, where ψ is a Boolean combination of primitive formulas as described by the items 1. to 5. and such that $\text{var}(\psi) = \{X_1, \dots, X_k\}$. We define $\text{coreNEO}_2 + \text{AC}$ (resp. $\text{coreNEO}_2 + \text{K}$) to be the fragment of $\text{NEO}_2 + \text{AC}$ (resp. $\text{NEO}_2 + \text{K}$) containing all formulas of the form $\exists X_1 \dots \exists X_k \psi$, where ψ is a Boolean combination of primitive formulas as described by the items 1. to 6. (resp. 1. to 5. as well as 7.) and such that $\text{var}(\psi) = \{X_1, \dots, X_k\}$. Finally, we define $\text{coreNEO}_2 + \text{ACK}$ to be the fragment of $\text{NEO}_2 + \text{ACK}$ containing all formulas of the form $\exists X_1 \dots \exists X_k \psi$, where ψ is a Boolean combination of primitive formulas as described by the items 1. to 7. and such that $\text{var}(\psi) = \{X_1, \dots, X_k\}$. The following observation proves that $\text{coreNEO}_2 + \text{ACK}$ is equivalent to $\text{NEO}_2 + \text{ACK}$, a similar statement for the DN logic can be found in [6]. We only provide a sketch of proof as most of the arguments can be found in [6, Subsection 3.1].

Lemma 4.2. *For every formula $\varphi \in \text{NEO}_2 + \text{ACK}$ one can compute in time $O(|\varphi|^2)$ an equivalent formula $\varphi' \in \text{coreNEO}_2 + \text{ACK}$ such that*

- $|\varphi'| \leq 10|\varphi|$,
- the sets $\sigma \subseteq \mathbb{N}$ occurring in φ' are exactly the ones occurring in φ , and
- φ' has at most two size measurements more than φ .

Moreover, if $\varphi \in \text{NEO}_2$ then $\varphi' \in \text{coreNEO}_2$; and if $\varphi \in \text{NEO}_2 + \text{K}$ then $\varphi' \in \text{coreNEO}_2 + \text{K}$.

Proof. We start by converting φ into prenex-normal form, i.e., of the form $\exists \bar{X} \xi$, where $\exists \bar{X}$ is a sequence of existentially quantified variables and ξ is quantifier-free.

Next, we get rid of every vertex variable. Let $x_1, \dots, x_t$ be the vertex variables of φ , we transform $\exists \bar{X} \xi$ by replacing each x_i by a new vertex set variable X_i , and we make the following translations: $x_i \in Y$ translates to $X_i \subseteq Y$, $x_i = x_j$ translates to $X_i = X_j$, and $E(x_i, x_j)$ translates to $X_i \subseteq N^{\{1\}}(X_j)$. Then we replace ξ by

$$\exists \bar{X} \xi \wedge X_1 \neq \emptyset \wedge \dots \wedge X_t \neq \emptyset \wedge |X_1 \cup \dots \cup X_t| \leq t.$$

This will make sure that every model of the new formula interprets the new variables $X_1, \dots, X_t$ as singletons. Similarly, we can get rid of edge variables.

Now, we replace every $t_1 \supseteq t_2$ by $t_2 \subseteq t_1$. After that, replace every $t_1 \subseteq t_2$ with $t_1 \cap t_2 = t_1$. Then, we replace every neighborhood term $N^\sigma(t)$ by $N^\sigma(t, E)$. The above operations conserve a formula in prenex-normal form that we denote by $\exists \bar{X} \psi$, where ψ is quantifier-free. Then we can

exhaustively apply the following simplification: If ψ contains a term $N^\sigma(t, t')$, $N^\sigma(t', t)$, $t \cup t'$, $t' \cup t$, $t \cap t'$, $t' \cap t$, $\bar{t}$, $\text{con}(t)$, $\text{acy}(t)$, $t' = t$, or $|t| < m$ for $< \in \{\leq, =, \geq\}$, where t is not a variable then we can replace ψ by $\exists Y t = Y \wedge \psi'$, where Y is some unused variable and ψ' is obtained from ψ by replacing t with Y . For example, the formula $N^{\sigma_1}(N^\sigma(Y, B), W) = \bar{Z}$ is replaced with

$$\exists X_1 \exists X_2 \exists X_3 N^{\sigma_1}(X_1, W) = X_2 \wedge N^\sigma(Y, X_3) = X_1 \wedge \bar{Z} = X_2 \wedge B = X_3.$$

Let φ' be resulting formula. The final transformation iteratively remove nested neighborhood operators, and equalities that are not primitive formulas. We deduce that φ' is a **coreNEO₂+ACK** formula. It can be checked that this procedure can be carried out in time quadratic in $|\varphi|$ and $|\varphi'|$ is at most a factor of 10 longer than the original one. Moreover, we did not change the set of σ -s occurring in the formula, and we introduce at most two new size measurements (for the vertex variables and edge variables). Finally, since we did not introduce new **con**, **acy**, **clique** constraints, the last statement of the lemma is trivially satisfied. $\square$

5 Partial Solutions and Related Notions

Let φ be a **coreNEO₂[FRec]** formula, we fix it for the rest of this paper. We also fix a vertex and edge colored graph G and we will show how to check whether $G \models \varphi$ given an elimination tree of depth td of G in time $2^{\mathcal{O}(\text{td})} n^{\mathcal{O}(1)}$ and space $\mathcal{O}(\text{td} \log(n))$. Here $\mathcal{O}$ hides some factors depending on the formula φ . In Theorem 7.7 we will explicitly provide the dependence of time and space complexity on φ .

Definition 5.1. We define $\boxtimes$ as the set of all pairs (X, Y) with $X \in \text{var}^V(\varphi)$ and $Y \in \text{var}^E(\varphi)$ such that $N^\sigma(X, Y)$ occurs in φ .²

We denote by $\text{EQ}^V(\varphi)$ the set of equalities in φ between neighborhood terms and by $\text{EQ}^E(\varphi)$ the set of those between edge set terms. And we use $\text{EQ}(\varphi)$ to denote the set $\text{EQ}^V(\varphi) \cup \text{EQ}^E(\varphi)$. Given an interpretation $(G, \mathbf{f}, \mathbf{g})$ of φ , we denote by $\perp(G, \mathbf{f}, \mathbf{g})$ the set of equalities $t_1 = t_2$ in $\text{EQ}(\varphi)$ such that $\llbracket t_1 = t_2 \rrbracket^{(G, \mathbf{f}, \mathbf{g})} = 0$. And we let $X_1, \dots, X_\ell$ be all (vertex set and edge set) variables of φ that occur in at least one size measurement of φ .

5.1 Formula parameters and auxiliary digraph

In this subsection, based on G and φ we define two crucial integers d_φ and p_φ as well as an auxiliary digraph $\vec{G}_\varphi$ relying on which we will later define our notion of partial solutions.

We define d_φ as the maximum over all $d(\sigma)$ for the sets σ occurring in φ ; moreover, we define p_φ as the least common multiple of all $p(\sigma)$ for all sets σ occurring in φ (see Definition 1.2 for the definitions of $d(\sigma)$ and $p(\sigma)$). To simplify our arguments, we assume that $d_\varphi \geq 1$. This is safe as we can always increase d_φ artificially without changing the correctness of any of the following statements including the following key observation.

Observation 5.2. Let $\sigma \subseteq \mathbb{N}$ be a set that occurs in φ . And let $U \subseteq V(G)$, $F \subseteq E(G)$, $v \in V(G)$, and $i \in [0, p_\varphi - 1]$ be such that $|N(v, F) \cap U| \geq d_\varphi$ and $i \equiv_{p_\varphi} |N(v, F) \cap U|$. Then we have $|N(v, F) \cap U| \equiv_{p(\sigma)} i + (d_\varphi \cdot p_\varphi)$. And therefore, $v \in N^\sigma(U, F)$ if and only if $i + (d_\varphi \cdot p_\varphi) \in \sigma$.

Our algorithm makes use of the following auxiliary digraph associated with G and φ .

²This notation is explained by the lack of imagination of the authors during the last day of a research visit. It stuck. We adore cats!

Definition 5.3. We define $\vec{G}_\varphi$ as the directed graph with vertex set $V(\vec{G}_\varphi) = \{v^1, \dots, v^{d_\varphi} : v \in V(G)\}$ and arc set $\vec{A}(\vec{G}_\varphi) = \{(u^1, v^i) : uv \in E(G) \wedge i \in [d_\varphi]\}$; see Figure 2 for an illustration. For a set $W \subseteq E(G)$ of edges, we denote by $\vec{W}$ the set of arcs $\{(u^1, v^i) : uv \in W \wedge i \in [d_\varphi]\}$.

We say that a set $\vec{S} \subseteq \vec{A}(\vec{G}_\varphi)$ of arcs is simple if for every edge $uv \in E(G)$, there is at most one $i \in [d_\varphi]$ such that $(u^1, v^i) \in \vec{S}$ and at most one $j \in [d_\varphi]$ such that $(v^1, u^j) \in \vec{S}$.



Figure 2: Illustration of the transformation from a graph G to the graph $\vec{G}_\varphi$. The edge $uv \in E(G)$ on the left corresponds to the arcs $(u^1, v^1), \dots, (u^1, v^{d_\varphi})$ and $(v^1, u^1), \dots, (v^1, u^{d_\varphi})$ on the right.

5.2 Neighborhood functions, mod-degrees and partial interpretations

The following definition introduces two important notions—neighborhood functions and mod-degrees—that play a key role in our algorithm. Indeed, they are important parts of our partial solutions, but also of the indexes of our recursive algorithm.

Definition 5.4. Let $\text{dom} \subseteq V(G)$ and³ $\text{sheaf} \subseteq E(G)$. A mod-degree on dom is a function $\text{pd}: \text{dom} \times \mathfrak{X} \rightarrow \mathbb{Z}_{p_\varphi}$. We denote by $\text{PD}(\text{dom})$ the set of all mod-degrees on dom . A neighborhood function on dom is a function $\mathbb{E}: \text{dom} \times \mathfrak{X} \rightarrow [0, d_\varphi]$.

Let us give some intuitions about how we use these two kinds of functions. Roughly, mod-degrees are used to handle the periodic part of the sets σ occurring in φ while neighborhood functions handle the part between 0 and d_φ . We use mod-degree to control, for each $v \in \text{dom}$ and $(X, Y) \in \mathfrak{X}$, the number of edges in Y between v and X modulo p_φ . We will apply the discrete Fourier transform to mod-degrees to accelerate the so-called join-computation combining the subtrees rooted at the children of some node of an elimination forest (see Subsection 6.5). We use neighborhood functions for two purposes:

- The *guessing* neighborhood functions denoted by $\mathbb{E}$ or $\hat{\mathbb{E}}$ are used to guess, for each $v \in \text{dom}$ and $(X, Y) \in \mathfrak{X}$, the number of expected edges in Y between v and X . In fact, when $\mathbb{E}(v, (X, Y)) < d_\varphi$, we expect that there will be **exactly** $\mathbb{E}(v, (X, Y))$ edges in Y between v and X . However, when $\mathbb{E}(v, (X, Y)) = d_\varphi$, we expect **at least** d_φ such edges—we emphasize this special meaning of the value d_φ as it will be used several times in definitions and intuitions. In both cases, this number is an expectation and we will first consider some partial solutions that do not respect this at all, but then we will filter these renegade partial solutions with inclusion-exclusion principle and counting.
- The *restricting* neighborhood functions denoted by alw are used to forbid a partial solution to have some arcs in $\vec{A}(\vec{G}_\varphi)$. We use this kind of neighborhood functions to apply the inclusion-exclusion principle in order to kill some of the renegade partial solutions mentioned above.

³Here, we use *sheaf* to denote an edge subset of G to prepare the reader that this definition will be applied for the sheaf of some node of the given elimination tree of G .

Given a guessing neighborhood function $\mathbb{E}$ and a mod-degree $\mathbf{pd}$, we know the set of vertices from $\mathbf{dom}$ that is expected to belong in a neighborhood term $N^\sigma(X, Y)$ occurring in φ . In addition to a vertex interpretation and an edge interpretation of φ , we can interpret all the terms and equalities of φ as follows.

Definition 5.5. *Given a vertex interpretation $\mathbf{f}$ of φ on $\mathbf{dom}$, an edge interpretation $\mathbf{g}$ of φ on $\mathbf{sheaf}$, a neighborhood function $\mathbb{E}$ on $\mathbf{dom}$, and a mod-degree $\mathbf{pd}$ on $\mathbf{dom}$, we define*

- $\llbracket t \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})} := \llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \cap \mathbf{sheaf}$ for every edge set term t ,
- $\llbracket t \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})} := \llbracket t \rrbracket^{(G, \mathbf{f}, \mathbf{g})} \cap \mathbf{dom}$ for every neighborhood term t that is not a neighborhood operator $N^\sigma(C)$ with $C \in \mathfrak{X}$.
- For every $C \in \mathfrak{X}$ and every $\sigma \subseteq \mathbb{N}$ that occurs in φ ,

$$\begin{aligned} \llbracket N^\sigma(C) \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})} := & \left\{ v \in \mathbf{dom} : \mathbb{E}(v, C) < d_\varphi \text{ and } \mathbb{E}(v, C) \in \sigma \right\} \cup \\ & \left\{ v \in \mathbf{dom} : \mathbb{E}(v, C) = d_\varphi \text{ and } \mathbf{pd}(v, C) + (d_\varphi \cdot p_\varphi) \in \sigma \right\}. \end{aligned}$$

Also, we set

- $\perp(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd}) := \{t_1 = t_2 \in \mathbf{EQ}(\varphi) : \llbracket t_1 \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})} \neq \llbracket t_2 \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})}\}.$
- $\perp^E(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd}) = \perp(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd}) \cap \mathbf{EQ}^E(\varphi)$ and $\perp^V(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd}) = \perp(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd}) \cap \mathbf{EQ}^V(\varphi).$

Observe that since φ is a $\mathbf{coreNEO}_2$ formula, Definition 5.5 provide an exhaustive way of interpreting the edge terms, neighborhood terms, and the equalities of φ .

Remark 5.6. *Observe that $\llbracket t \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})}$ does not depend on $\mathbf{g}$ when t is a neighborhood term, and it is also the case for $\perp^V(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})$. Thus, we sometimes simply write $\llbracket t \rrbracket^{(\mathbf{f}, \mathbb{E}, \mathbf{pd})}$ and $\perp^V(\mathbf{f}, \mathbb{E}, \mathbf{pd})$. For the same reasons, we sometimes write $\llbracket X \rrbracket^{(\mathbf{f})}$ and $\llbracket t \rrbracket^{(\mathbf{g})}$ when X is a vertex set variable and t is an edge term—sometimes, we add the superscript G to emphasize what the underlying graph is. Moreover, we sometimes use $\perp^E(\mathbf{g})$ as a shortcut for $\perp^E(\mathbf{f}, \mathbf{g}, \mathbb{E}, \mathbf{pd})$.*

The rest of this subsection is dedicated to introducing some notations and lemmas that will be used later on. Since some of the equalities in $\mathbf{EQ}(\varphi)$ might be negated or joined by disjunction in φ , a model of φ might, in general, falsify some of those equalities. For this reason, we will be interested in the set of equalities from $\mathbf{EQ}(\varphi)$ falsified by a certain partial solution (which will be formally defined in the next subsection). The following definition and lemma below will later be useful to argue about the equalities falsified by a restriction of a partial solution:

Definition 5.7. *Let $\mathbf{dom}' \subseteq \mathbf{dom} \subseteq V(G)$. Given a vertex interpretation $\mathbf{f}$ on $\mathbf{dom}$, we define the restriction $\mathbf{f}|_{\mathbf{dom}'}$ of $\mathbf{f}$ on $\mathbf{dom}'$ as the vertex interpretation on $\mathbf{dom}'$ such that $\mathbf{f}|_{\mathbf{dom}'}(X) = \mathbf{f}(X) \cap \mathbf{dom}'$ for every $X \in \mathbf{var}^V(\varphi)$. Moreover, for every neighborhood function or mod-degree $\mathbf{h}$ on $\mathbf{dom}$, we define the restriction $\mathbf{h}|_{\mathbf{dom}'}$ of $\mathbf{h}$ on $\mathbf{dom}'$ as the neighborhood function or mod-degree, respectively, such that $\mathbf{h}|_{\mathbf{dom}'}(u, C) = \mathbf{h}(u, C)$ for every $u \in \mathbf{dom}'$ and every $C \in \mathfrak{X}$.*

Similarly, let $\mathbf{sheaf}' \subseteq \mathbf{sheaf} \subseteq E(G)$. Given an edge interpretation $\mathbf{g}$ on $\mathbf{sheaf}$, we define the restriction $\mathbf{g}|_{\mathbf{sheaf}'}$ of $\mathbf{g}$ on $\mathbf{sheaf}'$ as the edge interpretation on $\mathbf{sheaf}'$ such that $\mathbf{g}|_{\mathbf{sheaf}'}(X) = \mathbf{g}(X) \cap \mathbf{sheaf}'$ for every $X \in \mathbf{var}^E(\varphi)$. And for a mapping $\vec{W}: \mathfrak{X} \rightarrow 2^{\vec{\mathbf{sheaf}}}$, we define the restriction $\vec{W}|_{\mathbf{sheaf}'}: \mathfrak{X} \rightarrow 2^{\vec{\mathbf{sheaf}'}}$ on $\mathbf{sheaf}'$ such that $\vec{W}|_{\mathbf{sheaf}'}(C) = \vec{W}(C) \cap \vec{\mathbf{sheaf}'}$ for all $C \in \mathfrak{X}$.

Lemma 5.8. *Let $\text{dom}' \subseteq \text{dom} \subseteq V(G)$ and let $\text{sheaf}' \subseteq \text{sheaf} \subseteq E(G)$. Further let $\mathbf{f}$ be a vertex interpretation of φ on dom , let $\mathbf{g}$ be an edge interpretation of φ on sheaf , let $\mathbb{E}$ be a neighborhood function on dom , and let pd be a mod-degree on dom . Then it holds that*

$$\perp(\mathbf{f}|_{\text{dom}'}, \mathbf{g}|_{\text{sheaf}'}, \mathbb{E}|_{\text{dom}'}, \text{pd}|_{\text{dom}'}) \subseteq \perp(\mathbf{f}, \mathbf{g}, \mathbb{E}, \text{pd}).$$

Proof. It suffices to observe that, by Definition 4.1, for every neighborhood term t , we have

$$\llbracket t \rrbracket^{(\mathbf{f}|_{\text{dom}'}, \mathbf{g}|_{\text{sheaf}'}, \mathbb{E}|_{\text{dom}'}, \text{pd}|_{\text{dom}'})} = \llbracket t \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \text{pd})} \cap \text{dom}'$$

and that for every edge set term t , we have

$$\llbracket t \rrbracket^{(\mathbf{f}|_{\text{dom}'}, \mathbf{g}|_{\text{sheaf}'}, \mathbb{E}|_{\text{dom}'}, \text{pd}|_{\text{dom}'})} = \llbracket t \rrbracket^{(\mathbf{f}, \mathbf{g}, \mathbb{E}, \text{pd})} \cap \text{sheaf}'.$$

□

To argue about the equalities falsified by a union of partial solutions, we will also use the following definition and the lemma after it.

Definition 5.9. *Let $t \in \mathbb{N}$, let $\text{dom} \subseteq V(G)$, and let $\text{sheaf} \subseteq E(G)$. Further, let $\text{dom}_1, \dots, \text{dom}_t$ be a partition of dom , let $\text{sheaf}_1, \dots, \text{sheaf}_t$ be a partition of sheaf , and for every $i \in [t]$ let $(\mathbf{f}_i, \mathbf{g}_i, \mathbb{E}_i, \text{pd}_i)$ be such that $\mathbf{f}_i$ is a vertex interpretation of φ on dom_i , $\mathbf{g}_i$ is an edge interpretation of φ on sheaf_i , $\mathbb{E}_i$ be a neighborhood function on dom_i , and pd_i be a mod-degree on dom_i . Then we define*

- the vertex interpretation $\mathbf{f}_{[t]}$ on dom via $\mathbf{f}_{[t]}(X) = \bigcup_{i \in [t]} \mathbf{f}_i(X)$ for all $X \in \text{var}^V(\varphi)$,
- the vertex interpretation $\mathbf{g}_{[t]}$ on sheaf via $\mathbf{g}_{[t]}(Y) = \bigcup_{i \in [t]} \mathbf{g}_i(Y)$ for all $Y \in \text{var}^E(\varphi)$,
- the neighborhood function $\mathbb{E}_{[t]}$ on dom via $\mathbb{E}_{[t]}(v, C) = \mathbb{E}_i(v, C)$ (for every $v \in \text{dom}$ and $C \in \mathfrak{X}$) where $i \in [t]$ is the unique value such that $v \in \text{dom}_i$,
- the mod-degree $\text{pd}_{[t]}$ on dom via $\text{pd}_{[t]}(v, C) = \text{pd}_i(v, C)$ (for every $v \in \text{dom}$ and $C \in \mathfrak{X}$) where $i \in [t]$ is the unique value such that $v \in \text{dom}_i$.

Lemma 5.10. *For objects from the above definition we have*

$$\perp(\mathbf{f}_{[t]}, \mathbf{g}_{[t]}, \mathbb{E}_{[t]}, \text{pd}_{[t]}) = \bigcup_{i \in [t]} \perp(\mathbf{f}_i, \mathbf{g}_i, \mathbb{E}_i, \text{pd}_i).$$

Proof. It suffices to observe that, by Definition 5.5, for every neighborhood term or edge term t , we have

$$\llbracket t \rrbracket^{(\mathbf{f}_{[t]}, \mathbf{g}_{[t]}, \mathbb{E}_{[t]}, \text{pd}_{[t]})} = \bigcup_{i \in [t]} \llbracket t \rrbracket^{(\mathbf{f}_i, \mathbf{g}_i, \mathbb{E}_i, \text{pd}_i)}.$$

□

5.3 Partial solutions, indexes, and compatibility

To simplify the definitions and statements in this subsection, we fix a triad $\text{tss} = (\text{tail}, \text{subtree}, \text{sheaf})$ of G . To get some intuition, one should think of subtree as the set of “finished” vertices, while the vertices in tail are still “active” and may, for example, get further neighbors. Moreover, we denote by broom the broom of tss , i.e., $\text{broom} := \text{tail} \cup \text{subtree}$. Now we are ready to present the notion of a partial solution used by our algorithm.

Definition 5.11. A partial solution for the triad $\mathbf{tss}$ is a quadruple $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ such that the following conditions are satisfied:

(Domain-respecting) We have:

- $\hat{\mathbf{f}}$ is a vertex interpretation of φ on **broom**,
- $\hat{\mathbf{g}}$ is an edge interpretation of φ on **sheaf**,
- $\hat{\mathbb{E}}$ is a neighborhood function on **broom**,
- $\vec{\mathcal{W}}$ maps each element $(X, Y) \in \mathfrak{X}$ to a **simple** subset of arcs from $\vec{\mathbf{sheaf}}$.

(Edge-consistent) For every $(X, Y) \in \mathfrak{X}$ and every edge $uv \in \mathbf{sheaf}$ the following equivalence holds: there exists $j \in [d_\varphi]$ with $(u^1, v^j) \in \vec{\mathcal{W}}(X, Y)$ if and only if $u \in \hat{\mathbf{f}}(X)$ and $uv \in \hat{\mathbf{g}}(Y)$.

(Respect-subtree- $\hat{\mathbb{E}}$) For every $(X, Y) \in \mathfrak{X}$, every $v \in \mathbf{subtree}$ and every $j \in [\hat{\mathbb{E}}(v, (X, Y))]$, the set $\vec{\mathcal{W}}(X, Y)$ contains at least one arc incoming into v^j .

(Respect-broom- $\hat{\mathbb{E}}$) For every $(X, Y) \in \mathfrak{X}$, every $v \in \mathbf{broom}$ and every $j \in [d_\varphi]$ with $j > \hat{\mathbb{E}}(v, (X, Y))$, the set $\vec{\mathcal{W}}(X, Y)$ contains no arcs incoming into v^j .

We denote by $\mathbf{Part}(\mathbf{tss})$ the set of all partial solutions for $\mathbf{tss}$.

It is not difficult to see (we will formally prove it later though) that the properties **(Edge-consistent)** and **(Respect-subtree- $\hat{\mathbb{E}}$)** imply that for every vertex $v \in \mathbf{subtree}$ and every pair $(X, Y) \in \mathfrak{X}$, the vertex v has at least $\hat{\mathbb{E}}(v, (X, Y))$ neighbors in $w \in \hat{\mathbf{f}}(X)$ such that $vw \in \hat{\mathbf{g}}(Y)$. Later in the proofs of Lemma 5.20 and Lemma 5.22 we will see how to ensure that (unless the special case $\hat{\mathbb{E}}(v, (X, Y)) = d_\varphi$ occurs), that $\hat{\mathbb{E}}(v, (X, Y))$ is also an upper bound on the number of such neighbors, i.e., this is precisely this number.

Definition 5.12. Given $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ satisfying **(Domain-respecting)** for $\mathbf{tss}$, we define the mod-degree $\mathbf{pd}_{\vec{\mathcal{W}}}$ on **broom** such that for every $w \in \mathbf{broom}$ and $C \in \mathfrak{X}$, we have:

$$\mathbf{pd}_{\vec{\mathcal{W}}}(w, C) := |\{(z^1, w^i) \in \vec{\mathcal{W}}(C) : z \in \mathbf{broom}\}| \pmod{p_\varphi}.$$

For a partial solution $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$, in particular satisfying **(Edge-consistent)**, we will be able to use, for a vertex $v \in \mathbf{subtree}$ and $(X, Y) \in \mathfrak{X}$, the value $\mathbf{pd}_{\vec{\mathcal{W}}}(v, (X, Y))$ as the number of neighbors $w \in \hat{\mathbf{f}}(X)$ of v such that $vw \in \hat{\mathbf{g}}(Y)$. Altogether, assuming that $\hat{\mathbb{E}}(v, (X, Y))$ is indeed also an upper bound on the number of such neighbors (which will be ensured later), we will be able to decide whether $v \in \mathbf{subtree}$ belongs to $N^\sigma(\hat{\mathbf{f}}(X), \hat{\mathbf{g}}(Y))$ based on the values $\hat{\mathbb{E}}(v, (X, Y))$ and $\mathbf{pd}_{\vec{\mathcal{W}}}(v, (X, Y))$.

This motivates the following definition of the indexes used by our recursive algorithm. Given two neighborhood functions $\mathbf{alw}$ and $\mathbb{E}$ on the same domain $\mathbf{dom}$, we write $\mathbf{alw} \leq \mathbb{E}$ if for every $v \in \mathbf{dom}$ and $C \in \mathfrak{X}$, we have $\mathbf{alw}(v, C) \leq \mathbb{E}(v, C)$.

Definition 5.13. An index for $\mathbf{tss}$ is a tuple $(\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}, \mathbf{F})$ where:

- $\mathbf{f}$ a vertex interpretation of φ on **tail**.
- $\mathbb{E}$ and $\mathbf{alw}$ are neighborhood functions on **tail** such that $\mathbf{alw} \leq \mathbb{E}$.
- $\mathbf{pd}$ is a mod-degree on **tail**,
- $\mathbf{F} \subseteq \mathbf{EQ}(\varphi)$.

We denote by $\text{Index}(\text{tss})$ the set of all indexes for tss . We say that a partial solution $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}(\text{tss})$ is compatible with an index $(\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F}) \in \text{Index}(\text{tss})$ when the following conditions are satisfied:

(f-consistent) For every $X \in \text{var}^V(\varphi)$, we have $\mathbf{f}(X) = \hat{\mathbf{f}}(X) \cap \text{tail}$.

(E-consistent) For every $v \in \text{tail}$ and every $C \in \mathfrak{X}$, we have $\mathbb{E}(v, C) = \hat{\mathbb{E}}(v, C)$.

(pd-consistent) For every $v \in \text{tail}$ and every $C \in \mathfrak{X}$, we have $\text{pd}(v, C) = \text{pd}_{\vec{\mathcal{W}}}(v, C)$.

(Respect-alw) For every $v \in \text{tail}$, every $C \in \mathfrak{X}$, and every arc $(u^1, v^j) \in \vec{\mathcal{W}}(C)$, we have $j \leq \text{alw}(v, C)$.

(Respect-F) We have $\perp(\hat{\mathbf{f}}|_{\text{subtree}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}|_{\text{subtree}}, (\text{pd}_{\vec{\mathcal{W}}})|_{\text{subtree}}) \subseteq \mathbf{F}$.

For every index $I \in \text{Index}(\text{tss})$, we denote the set of all partial solutions in $\text{Part}(\text{tss})$ compatible with I by $\text{Part}(\text{tss}, I)$. Moreover, we define $\text{Part}^=(\text{tss}, I)$ as the subset of $\text{Part}(\text{tss}, I)$ that contains all the partial solutions $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ such that $\perp(\hat{\mathbf{f}}|_{\text{subtree}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}|_{\text{subtree}}, (\text{pd}_{\vec{\mathcal{W}}})|_{\text{subtree}}) = \mathbf{F}$.

The neighborhood function alw will be used in the so-called forget-computation (see Subsection 6.4) to ensure that when a vertex is moved to the “finished” set subtree , it satisfies **(Respect-subtree- $\hat{\mathbb{E}}$)**, i.e., certain copies of this vertex in $\vec{G}_\varphi$ have incoming arcs. This, in turn, ensures that the function $\hat{\mathbb{E}}$ indeed provides lower bounds on the number of certain sets of relevant neighbors of this vertex. The inclusion-exclusion (see Theorem 3.5) allows to replace the task of ensuring the existence of incoming arcs by the task of forbidding such arcs. The neighborhood function alw is then used to ensure that some copies of the vertex are forbidden to have incoming arcs. We will provide all technical details in Subsection 6.4.

Since alw and $\mathbf{F}$ are involved only in **(Respect-alw)** and **(Respect-F)**, respectively, we deduce the following inclusions.

Observation 5.14. Let $(\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F}) \in \text{Index}(\text{tss})$. For every neighborhood function alw' on tail such that $\text{alw}' \leq \text{alw}$, we have:

$$\text{Part}(\text{tss}, (\mathbf{f}, \mathbb{E}, \text{alw}', \text{pd}, \mathbf{F})) \subseteq \text{Part}(\text{tss}, (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F})).$$

Moreover, for every $\mathbf{F}' \subseteq \mathbf{F}$, we have:

$$\text{Part}(\text{tss}, (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F}')) \subseteq \text{Part}(\text{tss}, (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F})).$$

We stratify the sets $\text{Part}(\text{tss}, I)$ as follows. The values $c_1, \dots, c_\ell$ simply keep track of the cardinalities of sets involved in size measurements in φ , while a and b will later be used to ensure that $\hat{\mathbb{E}}$ provides correct estimates on the cardinalities of certain neighborhoods.

Definition 5.15. Given a partial solution $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}(\text{tss})$, we define the following constants:

$$a_{\text{tss}}(P) := \sum_{C \in \mathfrak{X}} \sum_{\substack{v \in \text{subtree}: \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C),$$

$$b_{\text{tss}}(P) := \sum_{C \in \mathfrak{X}} \sum_{\substack{v \in \text{broom}: \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} |\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}|,$$

and for every $i \in [\ell]$,

$$c_{\text{tss}}^i(P) := |\llbracket X_i \rrbracket^{(G, \hat{\mathbf{f}}_{\text{subtree}}, \hat{\mathbf{g}})}|.$$

Moreover, we use $\bar{c}_{\text{tss}}(P)$ to denote the tuple $(c_{\text{tss}}^1(P), \dots, c_{\text{tss}}^\ell(P))$. Given an index $I \in \text{Index}(\text{tss})$, $a, b \in \mathbb{N}$ and $\bar{c} \in \mathbb{N}^\ell$ we define $\text{Part}(\text{tss}, I, a, b, \bar{c})$ as the set of all partial solutions $P \in \text{Part}(\text{tss}, I)$ such that $a_{\text{tss}}(P) = a$, $b_{\text{tss}}(P) = b$ and $\bar{c}_{\text{tss}}(P) = \bar{c}$. We define $\text{Part}^=(\text{tss}, I, a, b, \bar{c})$ from $\text{Part}(\text{tss}, I)$ similarly.

The following lemma will be useful to bound the range of values we are interested in. One should think of td as the depth of an elimination tree of G which we will fix soon:

Lemma 5.16. *Let G be a graph with at most $n \cdot \text{td}$ edges for some value td . Then for every $P \in \text{Part}(\text{tss})$, the value $a_{\text{tss}}(P)$ is at most $|\varphi| \cdot d_\varphi \cdot n$ and the value $b_{\text{tss}}(P)$ is at most $2 \cdot |\varphi| \cdot n \cdot \text{td}$. Further, for every $i \in [\ell]$, we have $c_{\text{tss}}^i(P) \leq n \cdot \text{td}$.*

Proof. The size of $\mathbb{X}$ is upper-bounded by $|\varphi|$, the sizes of subtree and broom are upper-bounded by n , and the size of sheaf is upper-bounded by $n \cdot \text{td}$. From this the bound on $a_{\text{tss}}(P)$ follows immediately. For every $C \in \mathbb{X}$, every edge uv of G contributes at most two arcs, namely one of form (u^1, v^i) and the other of form (v^1, u^j) (for some integers i, j), to the simple set $\vec{W}(C)$ so the bound on $b_{\text{tss}}(P)$ follows as well. And for $i \in [\ell]$, $c_{\text{tss}}^i(P)$ is the size of a set of vertices or edges, and thus it is upper-bounded by $n \cdot \text{td}$. $\square$

5.4 Relations between partial solutions and the formula satisfiability

Let root denote the triad defined as $\text{root} := (\emptyset, V(G), E(G))$. Informally speaking, this corresponds to the case where all vertices and edges are “processed” already. In this subsection, we will see how we can determine whether $G \models \varphi$ holds from the cardinality $|\text{Part}^=(\text{root}, I, a, b, \bar{c})|$ for some values of $I \in \text{Index}(\text{root})$, $a, b \in \mathbb{N}$ and $\bar{c} \in \mathbb{N}^\ell$.

Given $\mathbf{F} \subseteq \text{EQ}(\varphi)$ and $\bar{c} = (c^1, \dots, c^\ell) \in \mathbb{N}^\ell$, we denote by $\llbracket \varphi \rrbracket^{(\mathbf{F}, \bar{c})}$ the logical value of the formula obtained from φ by (1) replacing every subformula $t_1 = t_2 \in \text{EQ}(\varphi)$ by 0 if it belongs to $\mathbf{F}$, and by 1 otherwise; and (2) replacing every subformula $|X_i| \prec m$ with $\prec \in \{\leq, \geq, =\}$ by 1 if $c^i \prec m$, and by 0 otherwise.

This subsection is dedicated to prove the following characterization for $G \models \varphi$ based on our notion of partial solutions. For simplicity, if the domain of some function is empty, we denote this function by $\emptyset$. Similarly, we denote the unique vertex interpretation on the empty set by $\emptyset$.

Lemma 5.17. *We have $G \models \varphi$ if and only if there exist $\mathbf{F} \subseteq \text{EQ}(\varphi)$, $i \in \mathbb{N}$, and $\bar{c} \in \mathbb{N}^\ell$ such that $\llbracket \varphi \rrbracket^{(\mathbf{F}, \bar{c})} = 1$ and $\text{Part}^=(\text{root}, (\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}), i, i, \bar{c}) \neq \emptyset$.*

Observe that in this lemma we restrict ourselves to partial solutions P satisfying $a_{\text{root}}(P) = b_{\text{root}}(P) = i$ for some $i \in \mathbb{N}$ (cf. Definition 5.15). As we will see in the proof of this lemma, this restriction is a crucial trick to ensure that the neighborhood function $\hat{\mathbb{E}}$ used in P provides correct estimates on the cardinalities of relevant neighborhoods.

In order to prove Lemma 5.17, we first give an additional definition.

Definition 5.18. *Given $\mathbf{F} \subseteq \text{EQ}(\varphi)$ and $\bar{c} = (c^1, \dots, c^\ell) \in \mathbb{N}^\ell$, we say that an interpretation $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ of φ realizes $(\mathbf{F}, \bar{c})$ if the following conditions are satisfied:*

$$(\text{Realize-}\perp) \quad \perp(G, \hat{\mathbf{f}}, \hat{\mathbf{g}}) = \mathbf{F},$$

$$(\text{Realize-}\bar{c}) \quad \llbracket X_i \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})} = c^i \text{ for every } i \in [\ell].$$

Since in the above definition $\mathbf{F}$ is the set of equalities falsified by the interpretation $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$, and $\bar{c}$ determines the cardinalities of all sets involved in size measurements, the following observation immediately follows from the semantics of the logic:

Observation 5.19. *We have $G \models \varphi$ if and only if there exist $\mathbf{F} \subseteq \text{EQ}(\varphi)$ and $\bar{c} \in \mathbb{N}^\ell$ with $\llbracket \varphi \rrbracket^{(\mathbf{F}, \bar{c})} = 1$ such that there exists an interpretation $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ of φ realizing $(\mathbf{F}, \bar{c})$.*

With this observation in hand, we can split the proof of the equivalence from Lemma 5.17 into two lemmas below.

Lemma 5.20. *Let $I = (\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}) \in \text{Index}(\text{root})$ and $\bar{c} \in \mathbb{N}^\ell$. If an interpretation $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ of φ realizes $(\mathbf{F}, \bar{c})$, then there exist functions $\hat{\mathbb{E}}$ and $\vec{\mathcal{W}}$ with the following properties. Let $P := (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$. Then we have $P \in \text{Part}^-(\text{root}, I, a_{\text{root}}(P), b_{\text{root}}(P), \bar{c})$ and $a_{\text{root}}(P) = b_{\text{root}}(P)$.*

Proof. Let $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ be an interpretation of φ that realizes the pair $(\mathbf{F}, \bar{c})$. Also, let $c^1, \dots, c^\ell \in \mathbb{N}$ such that $\bar{c} = (c^1, \dots, c^\ell) \in \mathbb{N}^\ell$. In the following, we construct a partial solution $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ that belongs to $\text{Part}^-(\text{root}, I, a_{\text{root}}(P), b_{\text{root}}(P), \bar{c})$ with $a_{\text{root}}(P) = b_{\text{root}}(P)$. We construct $\hat{\mathbb{E}}$ and $\vec{\mathcal{W}}$ by doing the following for every $v \in V(G)$ and $C = (X, Y) \in \mathfrak{X}$. Let $U(v, C)$ be the set of all the vertices u in $\hat{\mathbf{f}}(X)$ such that uv is an edge in $\hat{\mathbf{g}}(Y)$. We set $\hat{\mathbb{E}}(v, C) = \min(d_\varphi, |U(v, C)|)$. For each $u \in U(v, C)$, we add one arc (u^1, v^j) to $\vec{\mathcal{W}}(C)$ in such a way that for every j between 1 and $\hat{\mathbb{E}}(v, C)$, there is at least one arc (u^1, v^j) in $\vec{\mathcal{W}}(C)$, and there is no such arc for any $j > \hat{\mathbb{E}}(v, C)$. Since $\text{pd}_{\vec{\mathcal{W}}}(v, C) = |\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}| \bmod p_\varphi$, we deduce from Observation 5.2 and the construction of $\hat{\mathbb{E}}$ and $\vec{\mathcal{W}}$ the following.

Observation 5.21. *For every neighborhood term $N^\sigma(X, Y)$ occurring in φ and $v \in V(G)$, we have $v \in \llbracket N^\sigma(X, Y) \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})}$ if and only if either:*

- $\mathbb{E}(v, (X, Y)) < d_\varphi$ and $\mathbb{E}(v, (X, Y)) \in \sigma$, or
- $\mathbb{E}(v, (X, Y)) = d_\varphi$ and $\text{pd}_{\vec{\mathcal{W}}}(v, C) + (d_\varphi \cdot d_\varphi) \in \sigma$.

Now, we prove that $P \in \text{Part}(\text{root})$, that is, P satisfies all the conditions of Definition 5.11. By definition, $\hat{\mathbf{f}}$ is a vertex interpretation on $V(G)$ and $\hat{\mathbf{g}}$ an edge interpretation on $E(G)$. By construction, $\hat{\mathbb{E}}$ is a neighborhood function on $V(G)$ and $\vec{\mathcal{W}}$ maps each $(X, Y) \in \mathfrak{X}$ to a simple subset of arcs from $E(G)$ so (Domain-respecting) is satisfied. Now, for each $C = (X, Y) \in \mathfrak{X}$, observe that:

- By construction, for every $v \in V(G)$ and $i \in [\hat{\mathbb{E}}(v, (X, Y))]$, there is at least one arc in $\vec{\mathcal{W}}(X, Y)$ incoming into v^i ; and for every $j \in [d_\varphi]$ with $j > \hat{\mathbb{E}}(v, (X, Y))$, there is no arc in $\vec{\mathcal{W}}(X, Y)$ incoming into v^j . Hence, (Respect-subtree- $\hat{\mathbb{E}}$) and (Respect-broom- $\hat{\mathbb{E}}$) are satisfied.
- For every $uv \in E(G)$ and every $X \in \mathfrak{X}$, there exists $i \in [d_\varphi]$ such that $(u^1, v^i) \in \vec{\mathcal{W}}(X, Y)$ if and only if $u \in U(v, C)$. Since $U(v, C)$ is the set of all vertices $u \in \hat{\mathbf{f}}(X)$ such that $uv \in \hat{\mathbf{g}}(Y)$, it follows that (Edge-consistent) is satisfied.

It follows that $P \in \text{Part}(\text{root})$. We need to prove that P belongs to $\text{Part}(\text{root}, I)$. Since the tail of root is empty, (f-consistent), ($\mathbb{E}$ -consistent), (pd-consistent) and (Respect- alw) are trivially satisfied by P . We need now to prove that (Respect-F) is satisfied. From Observation 5.21, we deduce that $\llbracket N^\sigma(X, Y) \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})} = \llbracket N^\sigma(X, Y) \rrbracket^{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})}$ for every neighborhood operator $N^\sigma(X, Y)$ occurring in φ . And for any neighborhood term t that is not a neighborhood operator, we have $\llbracket t \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})} =$

$\llbracket t \rrbracket^{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})}$ simply by definition, and the same holds for every edge set term t . Consequently, we have $\perp(G, \hat{\mathbf{f}}, \hat{\mathbf{g}}) = \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})$. From (Realize- $\perp$), it follows that

$$\perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}}) = \mathbf{F} \quad (2)$$

and thus P satisfies (Respect-F). This proves that P belongs to $\text{Part}(\text{root}, I)$. Thanks to (2), we also know that P satisfies the equality requirement to be in $\text{Part}^=(\text{root}, I)$. Since $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ satisfies (Realize- $\bar{c}$), for each $i \in [\ell]$, we have that $\llbracket |X_i| = c^i \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})} = 1$, so consequently, we have $c_{\text{root}}^i(P) = c^i$. It follows that P belongs to $\text{Part}^=(\text{root}, I, a_{\text{root}}(P), b_{\text{root}}(P), \bar{c})$.

It remains to prove $a_{\text{root}}(P) = b_{\text{root}}(P)$. For every vertex $v \in V(G)$ and every $C \in \mathfrak{X}$ with $\hat{\mathbb{E}}(v, C) < d_\varphi$, we have $\hat{\mathbb{E}}(v, C) = |U(v, C)|$ and by construction, there are exactly $|U(v, C)|$ arcs in $\vec{\mathcal{W}}(C)$ incoming into the vertices $v^1, \dots, v^{d_\varphi}$. We conclude that

$$a_{\text{root}}(P) = \sum_{C \in \mathfrak{X}} \sum_{\substack{v \in V(G) \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) = \sum_{C \in \mathfrak{X}} \sum_{\substack{v \in V(G) \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} |\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}| = b_{\text{root}}(P).$$

□

We now show that the converse of the above lemma holds as well, namely:

Lemma 5.22. *Let $I = (\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}) \in \text{Index}(\text{root})$ and let $\bar{c} \in \mathbb{N}^\ell$. Then, for every $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}^=(\text{tss}(\text{root}), I, a_{\text{root}}(P), b_{\text{root}}(P), \bar{c})$ with $a_{\text{root}}(P) = b_{\text{root}}(P)$, the interpretation $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ of φ realizes $(\mathbf{F}, \bar{c})$.*

Proof. First, let $c^1, \dots, c^\ell \in \mathbb{N}$ such that $\bar{c} = (c^1, \dots, c^\ell) \in \mathbb{N}^\ell$. Also, let $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ be a partial solution in $\text{Part}^=(\text{root}, I, a_{\text{root}}(P), b_{\text{root}}(P), \bar{c})$ such that $a_{\text{root}}(P) = b_{\text{root}}(P)$. By Definition 5.15, for every $i \in [\ell]$, we have $c_{\text{tss}}^i(P) := |\llbracket |X_i| \rrbracket^{(G, \hat{\mathbf{f}}|_{V(G)}, \hat{\mathbf{g}})}|$. As $\hat{\mathbf{f}}|_{V(G)} = \hat{\mathbf{f}}$, it follows that $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ satisfies (Realize- $\bar{c}$).

It remains to prove that $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ satisfies (Realize- $\perp$), i.e., $\perp(G, \hat{\mathbf{f}}, \hat{\mathbf{g}}) = \mathbf{F}$. As $P \in \text{Part}^=(\text{root}, I)$, we have $\perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}}) = \mathbf{F}$, thus it is enough to prove that $\perp(G, \hat{\mathbf{f}}, \hat{\mathbf{g}}) = \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})$. To prove this, we now check that $\llbracket t \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})} = \llbracket t \rrbracket^{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})}$ holds for every neighborhood or edge term t occurring in φ . By Definition 5.5 this is trivially satisfied when t is not a neighborhood operator. So we remain with the case that $t = N^\sigma(X, Y)$ for some set σ occurring in φ and some $C = (X, Y) \in \mathfrak{X}$, and this is why we need the following claim.

Claim 5.23. *For each $v \in V(G)$ and $C = (X, Y) \in \mathfrak{X}$, the number of edges $uv \in \hat{\mathbf{g}}(Y)$ with $u \in \hat{\mathbf{f}}(X)$ is:*

- exactly $|\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}|$,
- exactly $\hat{\mathbb{E}}(v, C)$ when $\hat{\mathbb{E}}(v, C) < d_\varphi$, and
- at least $\hat{\mathbb{E}}(v, C)$ when $\hat{\mathbb{E}}(v, C) = d_\varphi$.

Proof of the claim. Let $v \in V(G)$ and $(X, Y) \in \mathfrak{X}$. By (Edge-consistent), for every edge $uv \in E(G)$, there exists an arc $(u^1, v^j) \in \vec{\mathcal{W}}(X, Y)$ if and only if $u \in \hat{\mathbf{f}}(X)$ and $uv \in \hat{\mathbf{g}}(Y)$. As $\vec{\mathcal{W}}(X, Y)$ is simple, we deduce that $|\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}|$ is exactly the number of edges $uv \in \hat{\mathbf{g}}(Y)$ with $u \in \hat{\mathbf{f}}(X)$.

Thanks to (Respect-subtree- $\mathbb{E}$), we know that for each vertex v^j with $j \leq \hat{\mathbb{E}}(v, (X, Y))$, there exists an arc $(u^1, v^j) \in \vec{\mathcal{W}}(X, Y)$. It follows that there are at least $\hat{\mathbb{E}}(v, (X, Y))$ edges $uv \in \hat{\mathbf{g}}(Y)$ with $u \in \hat{\mathbf{f}}(X)$.

It remains to prove that there are exactly $\hat{\mathbb{E}}(v, (X, Y))$ of these edges when $\hat{\mathbb{E}}(v, (X, Y)) < d_\varphi$. This is where we use the fact that

$$a_{\text{root}}(P) = \sum_{C \in \mathfrak{X}} \sum_{\substack{v \in V(G) \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) = \sum_{C \in \mathfrak{X}} \sum_{\substack{v \in V(G) \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} |\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}| = b_{\text{root}}(P).$$

We already know that $|\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}| \geq \hat{\mathbb{E}}(v, (X, Y))$ holds for every $v \in V(G)$ and $(X, Y) \in \mathfrak{X}$. So the above equality can only hold when $|\{(u^1, v^i) \in \vec{\mathcal{W}}(C)\}| = \hat{\mathbb{E}}(v, (X, Y))$ for every $v \in V(G)$ and $(X, Y) \in \mathfrak{X}$ satisfying $\hat{\mathbb{E}}(v, (X, Y)) < d_\varphi$. And this concludes the proof of the claim. $\triangle$

By Definition 5.5, a vertex v belongs to $\llbracket N^\sigma(X, Y) \rrbracket^{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})}$ iff one of the following holds:

- $\hat{\mathbb{E}}(v, (X, Y)) < d_\varphi$ and $\hat{\mathbb{E}}(v, (X, Y)) \in \sigma$, or
- $\hat{\mathbb{E}}(v, (X, Y)) = d_\varphi$ and $\text{pd}_{\vec{\mathcal{W}}}(v, (X, Y)) + (d_\varphi \cdot p_\varphi) \in \sigma$.

When $\hat{\mathbb{E}}(v, (X, Y)) < d_\varphi$, by Claim 5.23, we get that $\hat{\mathbb{E}}(v, (X, Y))$ is exactly the number of edges $uv \in \hat{\mathbf{g}}(Y)$ with $u \in \hat{\mathbf{f}}(X)$. Consequently, when $\hat{\mathbb{E}}(v, (X, Y)) < d_\varphi$, we have v belongs to $\llbracket N^\sigma(X, Y) \rrbracket^{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})}$ if and only if v belongs to $\llbracket N^\sigma(X, Y) \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})}$.

Now recall that by definition we have that $\text{pd}_{\vec{\mathcal{W}}}(v, (X, Y)) \equiv_{p_\varphi} |\{(u^1, v^i) \in \vec{\mathcal{W}}(X, Y)\}|$. By Claim 5.23, $|\{(u^1, v^i) \in \vec{\mathcal{W}}(X, Y)\}|$ is exactly the number of edges $uv \in \hat{\mathbf{g}}(Y)$ with $u \in \hat{\mathbf{f}}(X)$. Moreover, when $\hat{\mathbb{E}}(v, (X, Y)) = d_\varphi$, this number is at least d_φ . From Observation 5.2, we deduce that when $\hat{\mathbb{E}}(v, (X, Y)) = d_\varphi$, then v belongs to $\llbracket N^\sigma(X, Y) \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})}$ if and only if $\text{pd}_{\vec{\mathcal{W}}}(v, (X, Y)) + (d_\varphi \cdot p_\varphi) \in \sigma$. We conclude that $\llbracket N^\sigma(X, Y) \rrbracket^{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})} = \llbracket N^\sigma(X, Y) \rrbracket^{(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})}$. As discussed above, this implies that $(G, \hat{\mathbf{f}}, \hat{\mathbf{g}})$ realizes $(\mathbf{F}, \bar{c})$ and concludes this proof. $\square$

We conclude this subsection by observing that Lemmas 5.20 and 5.22 imply Lemma 5.17.

6 Recursive Equalities

In this section, we formally define the polynomials that our algorithm will evaluate as well as the recursive equalities on which our algorithm relies on: the root/leaf/forget/join equalities. This is the technical core of our work.

6.1 The definitions of polynomials

Similarly as in [27, 32, 37], we encode the different choices of $a, b \in \mathbb{N}$ and $\bar{c} \in \mathbb{N}^\ell$ as degrees of formal variables $\alpha, \beta, \gamma_1, \dots, \gamma_\ell$, so that all the relevant values $|\text{Part}(\text{tss}, I, a, b, \bar{c})|$ can be stored as coefficients of one polynomial from $\mathbb{Z}[\alpha, \beta, \gamma_1, \dots, \gamma_\ell]$.

Definition 6.1. For every triad tss of G and every index $I = (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F}) \in \text{Index}(\text{tss})$, we define the polynomial $\mathcal{P}_{\text{tss}}(I)$ as follows:

$$\mathcal{P}_{\text{tss}}(I) = \sum_{a, b, c^1, \dots, c^\ell \in \mathbb{N}} |\text{Part}(\text{tss}, I, a, b, (c^1, \dots, c^\ell))| \cdot \alpha^a \beta^b \gamma_1^{c^1} \dots \gamma_\ell^{c^\ell}.$$

Moreover, we define its equality variant $\mathcal{P}_{\text{tss}}^=(I)$ as follows:

$$\mathcal{P}_{\text{tss}}^=(I) = \sum_{a, b, c^1, \dots, c^\ell \in \mathbb{N}} |\text{Part}^=(\text{tss}, I, a, b, (c^1, \dots, c^\ell))| \cdot \alpha^a \beta^b \gamma_1^{c^1} \dots \gamma_\ell^{c^\ell}.$$

Given $\bar{c} = (c^1, \dots, c^\ell) \in \mathbb{N}^\ell$, we denote by $\gamma^{\bar{c}}$ the monomial $\prod_{i \in [\ell]} \gamma_i^{c^i}$.

The following lemma bounds the coefficients of these polynomials. Recall that we assumed that $d_\varphi \geq 1$, this was done in order to simplify the following upper bounds.

Lemma 6.2. *Let G be a graph with at most $n \cdot \text{td}$ edges for some value td . For every triad $\text{tss} = (\text{tail}, \text{subtree}, \text{sheaf})$ of G and every index $I \in \text{Index}(\text{tss})$, we have that every coefficient of the polynomial $\mathcal{P}_{\text{tss}}(I)$ is non-negative and upper-bounded by $(d_\varphi + 1)^{5n \cdot \text{td} \cdot |\varphi|} = 2^{5n \cdot \text{td} \cdot |\varphi| \cdot \log(d_\varphi + 1)}$.*

Proof. Every partial solution is a quadruple so we bound the number of options for each of the four elements. For this note that the cardinalities of the sets $\text{var}^V(\varphi)$, $\text{var}^E(\varphi)$, and $\mathfrak{B}$ are upper-bounded by $|\varphi|$. The size of broom of tss is upper-bounded by n and the size of sheaf is upper-bounded by $|E(G)| \leq n \cdot \text{td}$ by Observation 3.1. The first element is a vertex interpretation on broom so there are at most $(2^n)^{|\varphi|} = 2^{n \cdot |\varphi|}$ options for it. Similarly, the second element is an edge interpretation on sheaf so there are at most $(2^{n \cdot \text{td}})^{|\varphi|} = 2^{n \cdot \text{td} \cdot |\varphi|}$ options. The third component is a neighborhood function on broom so there are at most $((d_\varphi + 1)^n)^{|\varphi|} = (d_\varphi + 1)^{n \cdot |\varphi|}$ possibilities. Finally, the fourth component maps every element of $\mathfrak{B}$ to a simple subset of sheaf , i.e., for every edge $uv \in \text{sheaf}$, at most one arc of form $u^1 v^i$ and at most one arc of form $v^1 u^j$ belongs to this simple subset for some $i, j \in [d_\varphi]$. Hence, there are at most $((d_\varphi + 1)^2)^{n \cdot \text{td} \cdot |\varphi|} = (d_\varphi + 1)^{2n \cdot \text{td} \cdot |\varphi|}$ options for the fourth element and the claim follows. $\square$

For our algorithm we will strongly rely on the following polynomials:

Definition 6.3. *Let p be a prime number such that $\mathbb{F}_p^*$ admits the p_φ -th root of unity and let $\omega(p)$ denote the smallest (with respect to the natural ordering of integers) such root of unity. For every triad tss of G and every index $I = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}, \mathbf{F}) \in \text{Index}(\text{tss})$, we define the following polynomial*

$$\mathcal{Q}_{\text{tss}}^p(I) \equiv_p \sum_{q \in \text{PD}(\text{tail})} \omega(p)^{q \cdot \mathbf{pd}} \cdot \mathcal{P}_{\text{tss}}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, q, \mathbf{F}) \quad (3)$$

from $\mathbb{F}_p[\alpha, \beta, \gamma_1, \dots, \gamma_\ell]$.

Note that if we fix all elements of the index apart from $\mathbf{pd}$, then this definition captures the discrete Fourier transform where the domain is $\text{PD}(\text{tail}) = \text{tail} \times \mathfrak{B}$.

In the remainder of this section we provide the recursive equalities on which our algorithm relies on. To this end, we fix an elimination tree T of G and let td denote its depth. The direct recursion relying on these equalities would result in an algorithm whose space complexity is polynomial in n . In the next section we will show how, relying on the Chinese remainder theorem, to improve this dependence to be only logarithmic in n . In particular, we will describe which (sufficiently short) prime numbers p are used. Let us recall that for a node u of T the triads $\text{tss}[u]$ and $\text{tss}(u)$ are $(\text{tail}[u], \text{subtree}(u), \text{sheaf}(u))$ and $(\text{tail}(u), \text{subtree}[u], \text{sheaf}(u))$, respectively. We define the following shortcuts:

Definition 6.4. *For every node u of T , we denote by $\text{Index}(u)$ and $\text{Index}[u]$ the sets $\text{Index}(\text{tss}(u))$ and $\text{Index}(\text{tss}[u])$, respectively. Every notation in the form $\spadesuit_{\text{tss}[u]}(\dots)$ is denoted by $\spadesuit_u[\dots]$. For example, we have $\mathcal{P}_u[I] := \mathcal{P}_{\text{tss}[u]}(I)$, $\mathcal{Q}_u^p[I] := \mathcal{Q}_{\text{tss}[u]}^p(I)$, $\text{Part}_u(I) := \text{Part}_{\text{tss}[u]}(I)$, and $a_u[P] := a_{\text{tss}[u]}(P)$.*

Moreover, every notation in the form $\spadesuit_{\text{tss}(u)}(\dots)$ is denoted by $\spadesuit_u(\dots)$. For example, we have $\mathcal{P}_u^-(I) := \mathcal{P}_{\text{tss}(u)}^-(I)$, $\mathcal{Q}_u^p(I) := \mathcal{Q}_{\text{tss}(u)}^p(I)$, and $b_u(P) := b_{\text{tss}(u)}(P)$.

6.2 Root equalities

Recall that if the domain of some function is empty, we denote this function by $\emptyset$. Similarly, we denote the unique vertex interpretation on the empty set by $\emptyset$. Now for simplicity, by $\emptyset^4$ we denote the tuple $(\emptyset, \emptyset, \emptyset, \emptyset)$. The following lemma reflects, for a set $F \subseteq \text{EQ}(\varphi)$, the relation between the partial solutions where only the equalities from F are possibly “falsified” on the one hand, and the partial solutions where precisely the equalities from F are falsified on the other hand. In fact, the analogue of this lemma holds for any node of the elimination tree T but in our applications we will only use it for the root.

Lemma 6.5. *For the root r of the elimination tree T and for every $F \subseteq \text{EQ}(\varphi)$, it holds that*

$$\mathcal{P}_r^=(\emptyset^4, F) = \sum_{F' \subseteq F} (-1)^{|F'|} \cdot \mathcal{P}_r(\emptyset^4, F \setminus F').$$

Proof. The proof is almost a direct application of the inclusion-exclusion principle. To see this formally, consider arbitrary but fixed $a, b \in \mathbb{N}$ and $\bar{c} \in \mathbb{N}^\ell$. For an equality $x \in F$ we define the set

$$\text{Part}_r(\emptyset^4, F, a, b, \bar{c}, x) = \{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_r(\emptyset^4, F, a, b, \bar{c}) : x \in \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})\}.$$

First, it holds that

$$\begin{aligned} \bigcap_{x \in F} \text{Part}_r(\emptyset^4, F, a, b, \bar{c}, x) &= \{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_r(\emptyset^4, F, a, b, \bar{c}) : F \subseteq \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})\} \\ &= \{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_r(\emptyset^4, F, a, b, \bar{c}) : F = \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}})\} = \text{Part}_r^=(\emptyset^4, F, a, b, \bar{c}) \end{aligned}$$

where the second equality holds since we have $\perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}}) \subseteq F$ for every $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ from $\text{Part}_r(\emptyset^4, F, a, b, \bar{c})$ due to (Respect-F) together with the fact that both broom and subtree of $\text{tss}(r)$ are equal to $V(G)$. Then by inclusion-exclusion principle (Theorem 3.5) we get

$$\begin{aligned} |\text{Part}_r^=(\emptyset^4, F, a, b, \bar{c})| &= \sum_{F' \subseteq F} (-1)^{|F'|} \cdot |\text{Part}_r(\emptyset^4, F, a, b, \bar{c}) \setminus \left(\bigcup_{x \in F'} \text{Part}_r(\emptyset^4, F, a, b, \bar{c}, x) \right)| \\ &= \sum_{F' \subseteq F} (-1)^{|F'|} \cdot |\text{Part}_r(\emptyset^4, F \setminus F', a, b, \bar{c})| \end{aligned} \tag{4}$$

where the last equality holds due to

$$\begin{aligned} &\text{Part}_r(\emptyset^4, F, a, b, \bar{c}) \setminus \left(\bigcup_{x \in F'} \text{Part}_r(\emptyset^4, F, a, b, \bar{c}, x) \right) \\ &= \{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_r(\emptyset^4, F, a, b, \bar{c}) : x \notin \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}}) \text{ for all } x \in F'\} \\ &= \{(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_r(\emptyset^4, F, a, b, \bar{c}) : \perp(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \text{pd}_{\vec{\mathcal{W}}}) \subseteq F \setminus F'\} = \text{Part}_r(\emptyset^4, F \setminus F', a, b, \bar{c}) \end{aligned}$$

for every $F' \subseteq F$. Since (4) holds for arbitrary $a, b, \bar{c}$, the lemma follows. $\square$

Now that we know how to compute the polynomials $\mathcal{P}_r^=(I)$ from the polynomials $\mathcal{P}_r(I)$. Observe that given the polynomials $\mathcal{P}_r^=(I)$, we can decide whether $G \models \varphi$ thanks to Lemma 5.17. In the remainder of the section, we prove that $\mathcal{Q}_r^p(I) \equiv_p \mathcal{P}_r(I)$ for every index I . This follows from the fact that $\text{tail}(r)$ is empty and thus the DFT-transformation is the identity function (modulo p).

Lemma 6.6. *For the root r of the elimination tree T and for every prime number p such that $\mathbb{F}_p^*$ admits the p_φ -th root of unity and every index $I = (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, F) \in \text{Index}(r)$, it holds that $\mathcal{Q}_r^p(I) \equiv_p \mathcal{P}_r(I)$.*

Proof. It holds that $\text{tail}(r) = \emptyset$ and thus also $\text{tail}(r) \times \mathfrak{X}$ is empty. So the set $\text{PD}(\text{tail}(r)) = (\mathbb{Z}_{p_\varphi})^{\text{tail}(r) \times \mathfrak{X}} = \mathbb{Z}_{p_\varphi}^\emptyset$ consists of precisely one mapping whose domain is empty. Since pd is such a mapping, we have $\text{PD}(\text{tail}(r)) = \{\text{pd}\}$. Thus we have

$$\begin{aligned} \mathcal{Q}_r^p(I) &\equiv_p \mathcal{Q}_{\text{tss}(r)}^p(I) \stackrel{(3)}{\equiv_p} \sum_{q \in \text{PD}(\text{tail}(r))} \left(\prod_{(v,C) \in \text{tail}(r) \times \mathfrak{X}} \omega(p)^{q(v,C) \cdot \text{pd}(v,C)} \right) \cdot \mathcal{P}_{\text{tss}(r)}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, q, \mathbf{F}) \\ &\equiv_p \sum_{q \in \text{PD}(\text{tail}(r))} 1 \cdot \mathcal{P}_{\text{tss}(r)}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, q, \mathbf{F}) \equiv_p \sum_{q \in \{\text{pd}\}} 1 \cdot \mathcal{P}_{\text{tss}(r)}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, q, \mathbf{F}) \\ &\equiv_p \mathcal{P}_{\text{tss}(r)}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, \text{pd}, \mathbf{F}) \equiv_p \mathcal{P}_{\text{tss}(r)}(I) \equiv_p \mathcal{P}_r(I). \end{aligned}$$

□

So now we may restrict our considerations to the computation DFT-transformed polynomials $\mathcal{Q}_u(I)$ and $\mathcal{Q}_u[I]$ for all the nodes u of T . To this end, we will sometime prove some equalities about the polynomials $\mathcal{P}_u[I]$ and $\mathcal{P}_u(I)$ as intermediary steps in our proofs.

6.3 Leaf equality

We start with the equality for the non-recursive computation at the leaves of T :

Lemma 6.7. *Let u be a leaf of T and let $I = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \text{pd}, \mathbf{F}) \in \text{Index}[u]$. Further, let p be a prime number such that $\mathbb{F}_p^*$ admits the p_φ -th root of unity. Then it holds that:*

$$\mathcal{Q}_u^p[I] \equiv_p \prod_{e=vw \in \text{sheaf}(u)} \sum_{\hat{\mathbf{g}}_e \in \mathbb{G}(e, \mathbf{F})} \gamma^{\bar{c}_{\hat{\mathbf{g}}_e}} \cdot \prod_{(X,Y) \in \mathfrak{X}} L(v, w, \hat{\mathbf{g}}_e, (X, Y)) \cdot L(w, v, \hat{\mathbf{g}}_e, (X, Y))$$

where

$$\mathbb{G}(e, \mathbf{F}) = \{\hat{\mathbf{g}} : \hat{\mathbf{g}} \text{ is an edge interpretation on } \{e\} \text{ and } \perp^E(\hat{\mathbf{g}}) \subseteq \mathbf{F}\},$$

$\bar{c}_{\hat{\mathbf{g}}_e} = (c_{\hat{\mathbf{g}}_e}^1, \dots, c_{\hat{\mathbf{g}}_e}^\ell) \in \{0, 1\}^\ell$ with $c_{\hat{\mathbf{g}}_e}^i = |\hat{\mathbf{g}}_e(X_i)|$ if $X_i \in \text{var}^E(\varphi)$ and $c_{\hat{\mathbf{g}}_e}^i = 0$ if $X \in \text{var}^V(\varphi)$ for all $i \in [\ell]$, and

$$L(v, w, \hat{\mathbf{g}}_e, (X, Y)) = \begin{cases} \mathbf{alw}(w, (X, Y)) \cdot \omega(p)^{\text{pd}(w, (X, Y))} \cdot \beta^{\lfloor \mathbb{E}(w, (X, Y)) \rfloor} & \text{if } v \in \mathbf{f}(X) \text{ and } e \in \hat{\mathbf{g}}_e(Y) \\ 1 & \text{otherwise.} \end{cases}$$

Proof. Since u is a leaf, we have $\text{tss}[u] = (\text{tail}[u], \emptyset, \text{sheaf}(u))$ and the broom of $\text{tss}[u]$ is equal to $\text{tail}[u]$, we will sometimes use implicitly these facts in this proof. Before we start to reformulate $\mathcal{Q}_u^p[I]$, we need the following notations and a claim.

Given $e = vw \in \text{sheaf}(u)$ and an edge interpretation $\hat{\mathbf{g}}$ of φ on $\text{sheaf}(u)$, we denote by $\hat{\mathbf{g}}_e$ the function $\hat{\mathbf{g}}|_{\{e\}}$. We define $\mathbb{G}(\mathbf{F})$ as the set of all edge interpretations $\hat{\mathbf{g}}$ of φ on $\text{sheaf}(u)$ such that $\hat{\mathbf{g}}_e \in \mathbb{G}(e, \mathbf{F})$ for every $e \in \text{sheaf}(u)$. Moreover, we denote by $\vec{\mathbb{W}}$ the set of all functions $\vec{\mathcal{W}}$ that maps each element from $\mathfrak{X}$ to a simple subset of arcs from $\text{sheaf}(u)$. Given $\vec{\mathcal{W}} \in \vec{\mathbb{W}}$ and $C \in \mathfrak{X}$, we define $\vec{\mathcal{W}}(C, v, w)$ as the intersection between $\vec{\mathcal{W}}(C)$ and the set of arcs $\{(v^1, w^i) : i \in [d_\varphi]\}$. Since $\vec{\mathcal{W}}(C)$ is a simple set of arcs, it follows that $\vec{\mathcal{W}}(C, v, w)$ contains at most one arc. Furthermore, we define

$$\vec{\mathbb{W}}(\hat{\mathbf{g}}_e, (X, Y), v, w) := \begin{cases} \{(v^1, w^i) : i \leq \mathbf{alw}(w, C)\} & \text{if } v \in \mathbf{f}(X) \text{ and } e \in \hat{\mathbf{g}}_e(Y) \\ \{\emptyset\} & \text{otherwise.} \end{cases}$$

Finally, we define $\vec{\mathbb{W}}(\hat{\mathbf{g}})$ as the set of all $\vec{\mathcal{W}} \in \vec{\mathbb{W}}$ such that $\vec{\mathcal{W}}(C, v, w) \in \vec{\mathbb{W}}(\hat{\mathbf{g}}_e, C, v, w)$ for every $C \in \mathfrak{X}$ and pair (v, w) such that $e = vw \in \text{sheaf}(u)$.

Claim 6.8. *For every $y \in \text{PD}(\text{tail}[u])$, we have $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_u[I[\text{pd} \rightarrow y]]$ if and only if the following conditions are satisfied:*

- $\hat{\mathbf{f}} = \mathbf{f}$,
- $\hat{\mathbb{E}} = \mathbb{E}$,
- $y = \text{pd}_{\vec{\mathcal{W}}}$,
- $\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})$,
- $\vec{\mathcal{W}} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})$.

Proof of the claim. Let $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_u[I[\text{pd} \rightarrow y]]$ for some $y \in \text{PD}(\text{tail}[u])$. Since the broom of $\text{tss}[u]$ is $\text{tail}[u]$, it holds that $\hat{\mathbf{f}} = \mathbf{f}$, $\hat{\mathbb{E}} = \mathbb{E}$ and $y = \text{pd}_{\vec{\mathcal{W}}}$ by **(f-consistent)**, **(E-consistent)** and **(pd-consistent)**. Consider an arbitrary edge $e = vw \in \text{sheaf}(u)$. Then by Lemma 5.8, we have $\perp^E(\hat{\mathbf{g}}_e) \subseteq \perp(\mathbf{f}_{|\emptyset}, \hat{\mathbf{g}}, \mathbb{E}_{|\emptyset}, (\text{pd}_{\vec{\mathcal{W}}})_{|\emptyset}) \subseteq \mathbf{F}$ where the last inclusion holds due to **(Respect-F)**. So we have $\hat{\mathbf{g}}_e \in \mathbb{G}(e, \mathbf{F})$. As it holds for any $e \in \text{sheaf}(u)$, we deduce that $\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})$. Let $C = (X, Y) \in \mathfrak{X}$. Due to **(Edge-consistent)**, if $v \notin \mathbf{f}(X)$ or $e \notin \hat{\mathbf{g}}(Y)$, then we have $\vec{\mathcal{W}}(C, v, w) = \emptyset$ and otherwise, there exists an arc $(v^1, w^j) \in \vec{\mathcal{W}}(C)$. In the latter case, since $\vec{\mathcal{W}}(C)$ is simple and due to **(Respect- alw)**, such a value j is unique and satisfies $1 \leq j \leq \text{alw}(w, (X, Y))$. Now observe that $e \notin \hat{\mathbf{g}}(Y)$ is equivalent to $e \notin \hat{\mathbf{g}}_e(Y)$. Altogether we deduce that $\vec{\mathcal{W}}(C, v, w) \in \vec{\mathbb{W}}(\hat{\mathbf{g}}_e, C, v, w)$ and we conclude that $\vec{\mathcal{W}} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})$.

For the other direction let $\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})$ and $\vec{\mathcal{W}} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})$ with $y = \text{pd}_{\vec{\mathcal{W}}}$. First, we argue that $P = (\mathbf{f}, \hat{\mathbf{g}}, \mathbb{E}, \vec{\mathcal{W}})$ is a partial solution for $\text{tss}[u]$. Observe that P trivially satisfies **(Domain-respecting)**. Further, by definition of $\vec{\mathbb{W}}(\hat{\mathbf{g}})$, for every $(X, Y) \in \mathfrak{X}$ and edge $e = vw \in \text{sheaf}(u)$, we have $\vec{\mathcal{W}}(C, v, w) = \{(v^1, w^i)\}$ (for some $i \in [d_\varphi]$) if and only if $v \in \mathbf{f}(X)$ and $e \in \hat{\mathbf{g}}_e(e)$. Since $e \in \hat{\mathbf{g}}_e(e)$ is equivalent to $e \in \hat{\mathbf{g}}(e)$, we deduce that P satisfies **(Edge-consistent)**. Furthermore, for every $w \in \text{tail}[u]$ and $C \in \mathfrak{X}$, if $\vec{\mathcal{W}}(C)$ contains an arc (v^1, w^i) , then by definition of $\vec{\mathbb{W}}(\hat{\mathbf{g}}_e, (X, Y), v, w)$, we have $i \leq \text{alw}(w, C)$. Since I is an index, we have $\text{alw} \leq \mathbb{E}$, and thus we have $i \leq \mathbb{E}(w, C)$ as required by **(Respect-broom-E)**. Finally, P trivially satisfies **(Respect-subtree-E)** because $\text{subtree}(u)$ is empty. It follows that P is indeed a partial solution for $\text{tss}[u]$.

Now it remains to prove that $P \in \text{Part}_u[I[\text{pd} \mapsto \text{pd}_{\vec{\mathcal{W}}}]$. The conditions **(f-consistent)**, **(E-consistent)**, and **(pd-consistent)** are trivially satisfied. As mentioned above, for every $C \in \mathfrak{X}$ and arc $(v^1, w^j) \in \vec{\mathcal{W}}(C)$, we have $j \leq \text{alw}(w, C)$, so **(Respect- alw)** is satisfied. The set $\text{subtree}(u)$ is empty so $\perp^V(\mathbf{f}_{|\text{subtree}(u)}, \hat{\mathbf{g}}, \mathbb{E}_{|\text{subtree}(u)}, (\text{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}(u)}) = \emptyset \subseteq \mathbf{F}$. And finally we have

$$\perp^E(\mathbf{f}_{|\emptyset}, \hat{\mathbf{g}}, \mathbb{E}_{|\emptyset}, (\text{pd}_{\vec{\mathcal{W}}})_{|\emptyset}) = \perp^E(\hat{\mathbf{g}}) \stackrel{\text{Lemma 5.10}}{=} \bigcup_{e \in \text{sheaf}(u)} \perp(\hat{\mathbf{g}}_e) \subseteq \mathbf{F}$$

where the last inclusion holds due to $\hat{\mathbf{g}}_e \in \mathbb{G}(e, \mathbf{F})$ for all $e \in \text{sheaf}(u)$ so **(Respect-F)** holds. Hence, P is indeed compatible with $I[\text{pd} \mapsto \text{pd}_{\vec{\mathcal{W}}}]$. $\triangle$

We are now ready to start our reformulation of $\mathcal{Q}_u^p[I]$. In the rest of the proof ω denotes $\omega(p)$.

$$\begin{aligned} \mathcal{Q}_u^p[I] &\stackrel{(3)}{\equiv_p} \sum_{y \in \text{PD}(\text{tail}[u])} \omega^{\text{pd} \cdot y} \cdot \mathcal{P}_u[I[\text{pd} \mapsto y]] \\ &\equiv_p \sum_{\substack{y \in \text{PD}(\text{tail}[u]) \\ P \in \text{Part}_u[I[\text{pd} \mapsto y]]}} \omega^{\text{pd} \cdot y} \cdot \alpha^{a_u[P]} \cdot \beta^{b_u[P]} \cdot \gamma^{\bar{c}_u[P]} \\ &\stackrel{6.8}{\equiv_p} \sum_{\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})} \sum_{\vec{\mathcal{W}} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})} \omega^{\text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}}} \cdot \alpha^{a_u[(\mathbf{f}, \hat{\mathbf{g}}, \mathbb{E}, \vec{\mathcal{W}})]} \cdot \beta^{b_u[(\mathbf{f}, \hat{\mathbf{g}}, \mathbb{E}, \vec{\mathcal{W}})]} \cdot \gamma^{\bar{c}_u[(\mathbf{f}, \hat{\mathbf{g}}, \mathbb{E}, \vec{\mathcal{W}})]}. \end{aligned} \quad (5)$$

The following claim will help us to reformulate the inner product in (5).

Claim 6.9. Let $P = (\mathbf{f}, \hat{\mathbf{g}}, \mathbb{E}, \vec{\mathcal{W}})$ with $\hat{\mathbf{g}} \in \mathbb{G}(\mathbb{F})$ and $\vec{\mathcal{W}} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})$. Then, we have $a_u[P] = 0$,

$$\begin{aligned}\bar{c}_u[P] &= \sum_{e \in \text{sheaf}(u)} \bar{c}_{\hat{\mathbf{g}}_e}, \\ b_u[P] &= \sum_{\substack{(v,w): \\ vw \in \text{sheaf}(u)}} \sum_{(X,Y) \in \mathfrak{K}} [v \in \mathbf{f}(X)] \cdot [e \in \hat{\mathbf{g}}_e(Y)] \cdot [\mathbb{E}(w, (X, Y)) < d_\varphi], \text{ and} \\ \text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}} &\equiv_{p_\varphi} \sum_{\substack{(v,w): \\ vw \in \text{sheaf}(u)}} \sum_{(X,Y) \in \mathfrak{K}} [v \in \mathbf{f}(X)] \cdot [e \in \hat{\mathbf{g}}_e(Y)] \cdot \text{pd}(w, C).\end{aligned}$$

Proof of the claim. Since $\text{subtree}(u) = \emptyset$, by definition, we have $a_u[P] = 0$. Let $i \in [\ell]$. If $X_i \in \text{var}^V(\varphi)$, then we have $c_u^i[P] = |\mathbf{f}(X_i) \cap \text{subtree}(u)| = 0$ because $\text{subtree}(u) = \emptyset$. On the other hand, if $X_i \in \text{var}^E(\varphi)$, then $c_u^i[P] = |\hat{\mathbf{g}}(X_i)|$. Since $|\hat{\mathbf{g}}(X_i)| = \sum_{e \in \text{sheaf}(u)} |\hat{\mathbf{g}}_e(X_i)|$, we deduce that $\bar{c}_u[P] = \sum_{e \in \text{sheaf}(u)} \bar{c}_{\hat{\mathbf{g}}_e}$. Now, notice that $b_u[P]$ can be reformulated as follows:

$$b_u[P] := \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u]: \\ \mathbb{E}(w, C) < d_\varphi}} |\{(v^1, w^i) \in \vec{\mathcal{W}}(C)\}| = \sum_{C \in \mathfrak{K}} \sum_{(v^1, w^i) \in \vec{\mathcal{W}}(C)} [\mathbb{E}(w, C) < d_\varphi].$$

Now, observe that by definition, for every $C \in \mathfrak{K}$, we have

$$\vec{\mathcal{W}}(C) = \bigcup_{\substack{(v,w): \\ vw \in \text{sheaf}(u)}} \vec{\mathcal{W}}(C, v, w).$$

Moreover, by Claim 6.8 and the definition of $\vec{\mathbb{W}}(\hat{\mathbf{g}}_e, (X, Y), v, w)$, we have

$$|\vec{\mathcal{W}}(C, v, w)| = [v \in \mathbf{f}(X)] \cdot [e \in \hat{\mathbf{g}}_e(Y)].$$

We deduce that the equality with $b_u[P]$ holds. It remains to prove the equality with $\text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}}$. Remember that by definition, we have

$$\text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}} := \sum_{(w,C) \in \text{tail}[u] \times \mathfrak{K}} \text{pd}(w, C) \cdot \text{pd}_{\vec{\mathcal{W}}}(w, C) \pmod{p_\varphi}.$$

Moreover, we have $\text{pd}_{\vec{\mathcal{W}}}(w, C) := |\{(z^1, w^j) \in \vec{\mathcal{W}}(C) : z \in \text{broom}[u]\}| \pmod{p_\varphi}$. We deduce that

$$\text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}} \equiv_{p_\varphi} \sum_{C \in \mathfrak{K}} \sum_{(v^1, w^j) \in \vec{\mathcal{W}}(C)} \text{pd}(w, C).$$

With the same arguments used for the equality with $b_u[P]$, we conclude that the equality with $\text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}}$ holds as well. $\triangle$

Consider again $P = (\mathbf{f}, \hat{\mathbf{g}}, \mathbb{E}, \vec{\mathcal{W}})$ with $\hat{\mathbf{g}} \in \mathbb{G}(\mathbb{F})$ and $\vec{\mathcal{W}} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})$. Recall that ω is the p_φ -th root of unity in $\mathbb{F}_p^*$ so in particular, we have $\omega^{p_\varphi} \equiv_p 1$, we deduce from Claim 6.9 that

$$\omega^{\text{pd} \cdot \text{pd}_{\vec{\mathcal{W}}}} \equiv_p \prod_{\substack{(v,w): \\ vw \in \text{sheaf}(u)}} \prod_{C \in \mathfrak{K}} \omega^{\text{pd}(w, C) \cdot [v \in \mathbf{f}(X)] \cdot [e \in \hat{\mathbf{g}}_e(Y)]}.$$

Again from Claim 6.9, it follows that

$$\beta^{b_u[P]} = \prod_{\substack{(v,w): \\ vw \in \text{sheaf}(u)}} \prod_{C \in \mathfrak{X}} \beta^{[v \in \mathbf{f}(X)] \cdot [e \in \hat{\mathbf{g}}_e(Y)] \cdot [\mathbb{E}(w,C) < d_\varphi]} \text{ and } \gamma^{\bar{c}_u[P]} = \prod_{e \in \text{sheaf}(u)} \bar{c}_{\hat{\mathbf{g}}_e}.$$

Hence, we have

$$\omega^{\text{pd} \cdot \text{pd}_{\vec{W}}} \cdot \alpha^{a_u[P]} \cdot \beta^{b_u[P]} \cdot \gamma^{\bar{c}_u[P]} = \prod_{e=vw \in \text{sheaf}(u)} \gamma^{\bar{c}_{\hat{\mathbf{g}}_e}} \prod_{(X,Y) \in \mathfrak{X}} L'(v, w, \hat{\mathbf{g}}_e, (X, Y)) \cdot L'(w, v, \hat{\mathbf{g}}_e, (X, Y)),$$

where $L'(v, w, \hat{\mathbf{g}}_e, (X, Y)) = \begin{cases} \omega^{\text{pd}(w, (X, Y))} \cdot \beta^{[\mathbb{E}(w, (X, Y)) < d_\varphi]} & \text{if } v \in \mathbf{f}(X) \text{ and } e \in \hat{\mathbf{g}}_e(Y) \\ 1 & \text{otherwise.} \end{cases}$

By plugging the above in (5), we obtain:

$$\begin{aligned} \mathcal{Q}_u^p[I] &\equiv_p \sum_{\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})} \sum_{\vec{W} \in \vec{\mathbb{W}}(\hat{\mathbf{g}})} \prod_{e=vw \in \text{sheaf}(u)} \gamma^{\bar{c}_{\hat{\mathbf{g}}_e}} \prod_{(X,Y) \in \mathfrak{X}} L'(v, w, \hat{\mathbf{g}}_e, (X, Y)) \cdot L'(w, v, \hat{\mathbf{g}}_e, (X, Y)) \\ &\equiv_p \sum_{\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})} |\vec{\mathbb{W}}(\hat{\mathbf{g}})| \cdot \prod_{e=vw \in \text{sheaf}(u)} \gamma^{\bar{c}_{\hat{\mathbf{g}}_e}} \prod_{(X,Y) \in \mathfrak{X}} L'(v, w, \hat{\mathbf{g}}_e, (X, Y)) \cdot L'(w, v, \hat{\mathbf{g}}_e, (X, Y)). \end{aligned}$$

By definition of $\vec{\mathbb{W}}(\hat{\mathbf{g}})$, we have

$$|\vec{\mathbb{W}}(\hat{\mathbf{g}})| = \prod_{\substack{e=vw \in \text{sheaf}(u) \\ C \in \mathfrak{X}}} |\vec{\mathbb{W}}(\hat{\mathbf{g}}_e, C, v, w)| \cdot |\vec{\mathbb{W}}(\hat{\mathbf{g}}_e, C, w, v)|.$$

Since $|\vec{\mathbb{W}}(\hat{\mathbf{g}}_e, C, v, w)| \cdot L'(v, w, \hat{\mathbf{g}}_e, (X, Y)) = L(v, w, \hat{\mathbf{g}}_e, (X, Y))$, it follows that

$$\mathcal{Q}_u^p[I] \equiv_p \sum_{\hat{\mathbf{g}} \in \mathbb{G}(\mathbf{F})} \prod_{e=vw \in \text{sheaf}(u)} \gamma^{\bar{c}_{\hat{\mathbf{g}}_e}} \prod_{(X,Y) \in \mathfrak{X}} L(v, w, \hat{\mathbf{g}}_e, (X, Y)) \cdot L(w, v, \hat{\mathbf{g}}_e, (X, Y)).$$

From the definition of $\mathbb{G}(\mathbf{F})$, we conclude that

$$\mathcal{Q}_u^p[I] \equiv_p \prod_{e=vw \in \text{sheaf}(u)} \sum_{\hat{\mathbf{g}}_e \in \mathbb{G}(e, \mathbf{F})} \gamma^{\bar{c}_{\hat{\mathbf{g}}_e}} \prod_{(X,Y) \in \mathfrak{X}} L(v, w, \hat{\mathbf{g}}_e, (X, Y)) \cdot L(w, v, \hat{\mathbf{g}}_e, (X, Y)).$$

□

6.4 Forget equality

Next we prove an equality reflecting the relation between the partial solutions of $\text{tss}(u)$ and of $\text{tss}[u]$ for a node u :

Lemma 6.10. *Let u be a node of T and let $I = (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F}) \in \text{Index}(u)$. Let $\mathcal{J}$ denote the set of all tuples $(\mathbf{f}_u, \mathbb{E}_u, \text{pd}_u)$ such that the following hold:*

- $\mathbf{f}_u$ is a vertex interpretation on $\{u\}$,
- $\mathbf{pd}_u$ is a mod-degree on $\{u\}$, and
- $\mathbb{E}_u$ is a neighborhood function on $\{u\}$,
- $\perp^V(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \subseteq \mathbf{F}$.

And let $\mathcal{J}'$ be the set of all tuples $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u)$ such that $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$ and $\mathbf{alw}_u$ is a neighborhood function on $\{u\}$ such that $\mathbf{alw}_u \leq \mathbb{E}_u$. Given, $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u) \in \mathcal{J}'$, we denote by $\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+$ and $\mathbf{pd}_u^+$ the functions $\mathbf{f} \cup \mathbf{f}_u, \mathbb{E} \cup \mathbb{E}_u, \mathbf{alw} \cup \mathbf{alw}_u$, and $\mathbf{pd} \cup \mathbf{pd}_u$, respectively. Then we have:

$$\mathcal{P}_u(I) = \sum_{(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u) \in \mathcal{J}'} S_{\pm}(\mathbb{E}_u, \mathbf{alw}_u) \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, \mathbf{pd}_u^+, \mathbf{F}] \cdot \alpha^{a_{\mathbb{E}_u}} \cdot \gamma^{\bar{c}_{\mathbf{f}_u}}, \quad (6)$$

with

$$S_{\pm}(\mathbb{E}_u, \mathbf{alw}_u) = \prod_{C \in \mathfrak{K}} \left((-1)^{\mathbb{E}_u(u, C) - \mathbf{alw}_u(u, C)} \cdot \binom{\mathbb{E}_u(u, C)}{\mathbf{alw}_u(u, C)} \right), \quad a_{\mathbb{E}_u} = \sum_{\substack{C \in \mathfrak{K}: \\ \mathbb{E}_u(u, C) < d_{\varphi}}} \mathbb{E}_u(u, C), \text{ and}$$

$\bar{c}_{\mathbf{f}_u} = (c_{\mathbf{f}_u}^1, \dots, c_{\mathbf{f}_u}^{\ell}) \in \{0, 1\}^{\ell}$ such that $c_{\mathbf{f}_u}^i = |\mathbf{f}_u(X_i)|$ if $X_i \in \text{var}^V(\varphi)$, and $c_{\mathbf{f}_u}^i = 0$ if $X_i \in \text{var}^E(\varphi)$ for every $i \in [\ell]$. Further, for every prime number p such that $\mathbb{F}_p^*$ admits a p_{φ} -th root of unity it holds that:

$$\begin{aligned} \mathcal{Q}_u^p(I) \equiv_p & \sum_{(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u) \in \mathcal{J}'} S_{\pm}(\mathbb{E}_u, \mathbf{alw}_u) \cdot \alpha^{a_{\mathbb{E}_u}} \cdot \gamma^{\bar{c}_{\mathbf{f}_u}} \cdot \\ & \frac{1}{p_{\varphi}^{|\mathfrak{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega(p)^{-\hat{q} \cdot \mathbf{pd}_u} \cdot \mathcal{Q}_u^p[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, \mathbf{pd} \cup \hat{q}, \mathbf{F}] \end{aligned} \quad (7)$$

where $\frac{1}{p_{\varphi}^{|\mathfrak{K}|}}$ is the multiplicative inverse of $p_{\varphi}^{|\mathfrak{K}|}$ in $\mathbb{F}_p$.

Proof. We start by proving the “non-DFT-transformed” equality (6). Before we proceed with a formal proof, we provide an intuition behind it. In essence, this equality is the application of the inclusion-exclusion principle. Recall that we have $\text{subtree}[u] = \text{subtree}(u) \dot{\cup} \{u\}$ whilst the triads $\text{tss}(u)$ and $\text{tss}[u]$ share the same broom the same sheaf. Now the following observations are crucial. Roughly, the partial solutions of $\text{tss}[u]$ and $\text{tss}(u)$ differ “only” on two aspects: in the partial solutions of $\text{tss}(u)$, the equalities “falsified by u ” have to belong to $\mathbf{F}$ as well (due to [\(Respect-F\)](#)), and second, due to [\(Respect-subtree-Ê\)](#), the sets in $\vec{\mathcal{W}}$ have to contain arcs incoming into certain vertices from $u^1, \dots, u^{d_{\varphi}}$. The definition of $\mathcal{J}$ ensures that [\(Respect-F\)](#) is satisfied on the vertex u . Thus, when we fix the tuple $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$ the computation of the desired cardinality of the set of partial solutions of $\text{tss}(u)$ boils down to ensuring that [\(Respect-subtree-Ê\)](#) holds for u .

To simplify the explanation in this sketch, assume that $\mathfrak{K}$ contains only one element C . Thus, we have to ensure that for every $j \in [\mathbb{E}_u(u, C)]$, the set $\vec{\mathcal{W}}(C)$ contains an arc incoming into u^j . So if we denote by P_j the set of partial solutions of $\text{tss}[u]$ where $\vec{\mathcal{W}}(C)$ contains an arc incoming into u^j , then what we seek is the value $|\bigcap_{j \in [\mathbb{E}_u(u, C)]} P_j|$. By inclusion-exclusion principle (see Theorem 3.5), we can compute this cardinality by considering all possible $S \subseteq [\mathbb{E}_u(u, C)]$ and determining the cardinality of the set of partial solutions where we *forbid* the edges incoming into u^j for all $j \in S$. Finally, we observe that for fixed S , due to symmetries, in terms of counting, this is the same as to forbid the edges incoming into $u^{\mathbb{E}_u(u, C) - |S| + 1}, \dots, u^{\mathbb{E}_u(u, C)}$, i.e., we can use the function $\mathbf{alw}_u$ with $\mathbf{alw}_u(u, C) = \mathbb{E}_u(u, C) - |S|$. Now we formalize this idea.

Given a neighborhood function $\mathbb{E}_u$ on $\{u\}$, we define $\mathcal{T}_{\mathbb{E}_u} := \{(C, j) : C \in \mathfrak{X}, j \in [\mathbb{E}_u(u, C)]\}$. For every $\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}$ and $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$, we define the set $\text{Part}_u[J, \mathcal{T}]$ of all $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbb{E}_u, \mathbf{pd}_u^+, \mathbf{F}]$ such that for every $(C, j) \in \mathcal{T}$, there is at least one arc in $\vec{\mathcal{W}}(C)$ incoming into u^j . Notice that we use $\mathbf{alw} \cup \mathbb{E}_u$ here, because it gives us the least restrictive index of $\text{tss}[u]$ in the sense that for every neighborhood function $\mathbf{alw}_u$ on $\{u\}$ such that $\mathbf{alw}_u \leq \mathbb{E}_u$, we have by Observation 5.14 the following inclusion

$$\text{Part}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, \mathbf{pd}_u^+, \mathbf{F}] \subseteq \text{Part}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbb{E}_u, \mathbf{pd}_u^+, \mathbf{F}].$$

Given $a, b \in \mathbb{N}$ and $\bar{c} \in \mathbb{N}^\ell$, we also define the set $\text{Part}_u[J, \mathcal{T}, a, b, \bar{c}]$ of all $P \in \text{Part}_u[J, \mathcal{T}]$ such that $a_u[P] = a$, $b_u[P] = b$ and $\bar{c}_u[P] = \bar{c}$. We aim to prove with the following claims that we have

$$\mathcal{P}_u(I) = \sum_{\substack{(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J} \\ a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell}} |\text{Part}_u[\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]| \cdot \alpha^{a+a_{\mathbb{E}_u}} \cdot \beta^b \cdot \gamma^{\bar{c}+\bar{c}_{\mathbf{f}_u}}. \quad (8)$$

We will see later how to prove (6) from (8) with the inclusion-exclusion principle.

Claim 6.11. *For every $P \in \text{Part}_u(I)$, there is a unique $J \in \mathcal{J}$ such that $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$.*

Proof of the claim. Let $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_u(I)$. Let $\mathbf{f}_u, \mathbb{E}_u$ and $\mathbf{pd}_u$ be the restrictions of $\hat{\mathbf{f}}, \hat{\mathbb{E}}$ and $\mathbf{pd}_{\vec{\mathcal{W}}}$ on $\{u\}$. First, observe that $P \in \text{Part}_u(I)$ and (Respect-F) imply the inclusion

$$\perp(\hat{\mathbf{f}}_{|\text{subtree}[u]}, \hat{\mathbf{g}}, \hat{\mathbb{E}}_{|\text{subtree}[u]}, (\mathbf{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}[u]}) \subseteq \mathbf{F}. \quad (9)$$

By applying Lemma 5.8 on (9) with $\text{dom}' = \{u\}$ and $\text{dom} = \text{subtree}[u]$, we deduce the inclusion

$$\perp^V(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \subseteq \perp^V(\hat{\mathbf{f}}_{|\text{subtree}[u]}, \hat{\mathbb{E}}_{|\text{subtree}[u]}, (\mathbf{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}[u]}) \subseteq \mathbf{F}.$$

Thus, we have $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$. Next, we argue that $P \in \text{Part}(\text{tss}[u])$ holds because we have $P \in \text{Part}(\text{tss}(u))$. Indeed, since each of the broom and sheaf is the same for $\text{tss}(u)$ and $\text{tss}[u]$, the quadruple P satisfies (Domain-respecting), (Edge-consistent), and (Respect-broom- $\hat{\mathbb{E}}$) for $\text{tss}[u]$ too. Moreover, the subtree of $\text{tss}[u]$ (i.e., $\text{subtree}(u)$) is a subset of the subtree of $\text{tss}(u)$ (i.e., $\text{subtree}[u]$), consequently P satisfies (Respect-subtree- $\hat{\mathbb{E}}$) for $\text{tss}[u]$. Thus, we have $P \in \text{Part}(\text{tss}[u])$.

Now, we show that P is compatible with $J' = (\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbb{E}_u, \mathbf{pd}_u^+, \mathbf{F})$.

- Recall that $\mathbf{f}_u^+ = \mathbf{f} \cup \mathbf{f}_u$ and $\mathbf{f}_u$ is the restriction of $\hat{\mathbf{f}}$ on $\{u\}$. Moreover, $\mathbf{f}(X) = \hat{\mathbf{f}}(X) \cap \text{tail}(u)$ for every $X \in \text{var}^V(\varphi)$ because P is compatible with I . We deduce that P satisfies (f-consistent) for J' because for every $X \in \text{var}^V(\varphi)$, we have:

$$\hat{\mathbf{f}}(X) \cap \text{tail}[u] = (\hat{\mathbf{f}}(X) \cap \text{tail}(u)) \cup (\hat{\mathbf{f}}(X) \cap \{u\}) = \mathbf{f}(X) \cup \mathbf{f}_u(X) = \mathbf{f}_u^+(X).$$

- Recall that $\mathbb{E}_u^+ = \mathbb{E} \cup \mathbb{E}_u$. Since $P \in \text{Part}_u(I)$, we have $\mathbb{E}_u^+(v, C) = \mathbb{E}(v, C) = \hat{\mathbb{E}}(v, C)$ for every $v \in \text{tail}(u)$ and $C \in \mathfrak{X}$. As $\mathbb{E}_u$ is the restriction of $\hat{\mathbb{E}}$ on $\{u\}$, we have $\mathbb{E}_u^+(u, C) = \hat{\mathbb{E}}(u, C)$ for every $C \in \mathfrak{X}$. Hence, P satisfies ($\mathbb{E}$ -consistent) for J' . With the same arguments, we can prove that P satisfies (pd-consistent) for J' .
- Since $P \in \text{Part}_u(I)$, we know that it satisfies (Respect-alw) for J' and every vertex in $\text{tail}(u)$. Moreover, P satisfying (Respect-broom- $\hat{\mathbb{E}}$) (for $\text{tss}(u)$ and $\text{tss}[u]$) implies that for every $C \in \mathfrak{X}$ and every arc $(v^1, u^j) \in \vec{\mathcal{W}}(C)$, we have $j \leq \hat{\mathbb{E}}(u, C)$. As $\hat{\mathbb{E}}(u, C) = \mathbb{E}_u(u, C) = (\mathbf{alw} \cup \mathbb{E}_u)(u, C)$, we deduce that P satisfies (Respect-alw) for J' and the vertex u as well.

- By applying Lemma 5.8 on (9) with $\text{dom}' = \text{subtree}(u)$ and $\text{dom} = \text{subtree}[u]$, we obtain

$$\perp(\hat{\mathbf{f}}_{|\text{subtree}(u)}, \hat{\mathbf{g}}, \hat{\mathbb{E}}_{|\text{subtree}(u)}, (\mathbf{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}(u)}) \subseteq \mathbf{F}.$$

It follows that P satisfies (Respect-F) for J' .

Hence, we have $P \in \text{Part}_u[J']$.

Now, we need to prove that $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$, i.e., for each $(C, j) \in \mathcal{T}_{\mathbb{E}_u}$, there is at least one arc in $\vec{\mathcal{W}}(C)$ incoming into u^j . Since $P \in \text{Part}(\text{tss}(u))$, from (Respect-subtree- $\hat{\mathbb{E}}$), it follows that for every $w \in \text{subtree}[u]$, every $C \in \mathfrak{K}$, and every $j \in [\hat{\mathbb{E}}(w, C)]$, there is at least one arc in $\vec{\mathcal{W}}(C)$ incoming into w^j . As $u \in \text{subtree}[u]$, we thus have $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$.

Finally, it is easy to see that J is the unique element of $\mathcal{J}$ such that $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$ because of the conditions (f-consistent), ($\mathbb{E}$ -consistent) and (pd-consistent). Δ

Claim 6.12. *The collection $(\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}])_{J \in \mathcal{J}}$ is a partition of $\text{Part}_u(I)$.*

Proof of the claim. Thanks to Claim 6.11, to prove this statement, it is sufficient to prove that for every $J \in \mathcal{J}$ and $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$, we have $P \in \text{Part}_u(I)$. Let $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$ and $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$.

We start by proving that $P \in \text{Part}(\text{tss}(u))$. Observe that P satisfies (Domain-respecting) and (Edge-consistent), (Respect-broom- $\hat{\mathbb{E}}$) for $\text{tss}(u)$ because P satisfies these properties for $\text{tss}[u]$ and $\text{tss}(u)$ and $\text{tss}[u]$ have the same broom and sheaf. It remains to check that P satisfies (Respect-subtree- $\hat{\mathbb{E}}$) for $\text{tss}(u)$. Observe that P satisfies this condition for $\text{tss}[u]$ and we have $\text{subtree}(u) = \text{subtree}[u] \setminus \{u\}$. So it remains to prove that, for every $C \in \mathfrak{K}$ and every $j \in [\hat{\mathbb{E}}(u, C)]$, there is at least one arc incoming into u^j in $\vec{\mathcal{W}}(C)$. Since $\hat{\mathbb{E}}(u, C) = \mathbb{E}_u(u, C)$, this follows from the definition of $\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$.

It remains to prove that P is compatible with I . Let $J' = (\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbb{E}_u, \mathbf{pd}_u^+, \mathbf{F})$.

- Observe that $\mathbf{f}$, $\mathbb{E}$, $\mathbf{alw}$ and $\mathbf{pd}$ are the restrictions of $\mathbf{f}_u^+$, $\mathbb{E}_u^+$, $\mathbf{alw} \cup \mathbb{E}_u$ and $\mathbf{pd}_u^+$ on $\text{tail}(u)$. Since $\text{tail}(u) \subseteq \text{tail}[u]$, we deduce that P satisfies (f-consistent), ($\mathbb{E}$ -consistent), (pd-consistent), and (Respect-alw) for I because it satisfies these properties for J' .
- It remains to deal with (Respect-F). For $S \in \{\text{subtree}(u), \text{subtree}[u]\}$, we denote by $P|_S$ the tuple $(\hat{\mathbf{f}}|_S, \hat{\mathbf{g}}, \hat{\mathbb{E}}|_S, (\mathbf{pd}_{\vec{\mathcal{W}}})|_S)$. As P satisfies this condition for J' , we know that

$$\perp(P|_{\text{subtree}(u)}) \subseteq \mathbf{F}. \quad (10)$$

We want to prove that $\perp(P|_{\text{subtree}[u]}) \subseteq \mathbf{F}$. Observe that, for each $S \in \{\text{subtree}(u), \text{subtree}[u]\}$, the set $\perp(P|_S)$ is a disjoint union of $\perp^V(P|_S)$ and $\perp^E(\hat{\mathbf{g}})$. From (10), we know that $\perp^E(\hat{\mathbf{g}}) \subseteq \mathbf{F}$. Thus, it remains to prove that $\perp^V(P|_{\text{subtree}[u]}) \subseteq \mathbf{F}$. As $\text{subtree}[u] = \text{subtree}(u) \cup \{u\}$, by Lemma 5.10, we have

$$\perp^V(P|_{\text{subtree}[u]}) = \perp^V(P|_{\text{subtree}(u)}) \cup \perp^V(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \quad (11)$$

because $\mathbf{f}_u$, $\mathbb{E}_u$ and $\mathbf{pd}_u$ are the restrictions of $\hat{\mathbf{f}}$, $\hat{\mathbb{E}}$ and $\mathbf{pd}_{\vec{\mathcal{W}}}$ on $\{u\}$. The property $J \in \mathcal{J}$ implies the inclusion $\perp^V(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \subseteq \mathbf{F}$. Moreover, (10) implies that $\perp^V(P|_{\text{subtree}(u)}) \subseteq \mathbf{F}$ holds. Hence, we deduce from (11) that $\perp^V(P|_{\text{subtree}[u]}) \subseteq \mathbf{F}$. Thus, P satisfies (Respect-F) for I .

We conclude that $P \in \text{Part}_u(I)$ and this proves the present claim. Δ

Claim 6.13. For every $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$ and every $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$, we have $a_u(P) = a_u[P] + a_{\mathbb{E}_u}$, $b_u(P) = b_u[P]$, and $\bar{c}_u(P) = \bar{c}_u[P] + \bar{c}_{\mathbf{f}_u}$.

Proof of the claim. Let $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$ and let $P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]$. We start by proving $a_u(P) = a_u[P] + a_{\mathbb{E}_u}$. By definition, we have

$$a_u(P) = \sum_{C \in \mathbb{X}} \sum_{\substack{v \in \text{subtree}[u] \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C).$$

Recall that $a_{\mathbb{E}_u}$ is the sum of $\mathbb{E}_u(u, C)$ over all $C \in \mathbb{X}$ with $\mathbb{E}_u(u, C) < d_\varphi$. Moreover, $a_u[P]$ is the sum of $\hat{\mathbb{E}}(v, C)$ over all $v \in \text{subtree}(u)$ and $C \in \mathbb{X}$ with $\hat{\mathbb{E}}(v, C) < d_\varphi$. As $\text{subtree}[u] = \text{subtree}(u) \cup \{u\}$ and $\mathbb{E}_u$ is the restriction of $\hat{\mathbb{E}}$ on $\{u\}$, we have $a_u(P) = a_u[P] + a_{\mathbb{E}_u}$.

Now, observe that $b_u(P) = b_u[P]$ because $\text{tss}(u)$ and $\text{tss}[u]$ have the same broom.

It remains to prove that $\bar{c}_u(P) = \bar{c}_u[P] + \bar{c}_{\mathbf{f}_u}$, i.e., for every $i \in [\ell]$, we have $c_u^i(P) = c_u^i[P] + c_{\mathbf{f}_u}^i$. Let $i \in [\ell]$. From the definition of $c_u^i(P)$ and $c_u^i[P]$, we have

$$c_u^i(P) = |\llbracket X_i \rrbracket^{(G, \hat{\mathbf{f}}_{\text{subtree}[u]}, \hat{\mathbf{g}})}| \text{ and } c_u^i[P] = |\llbracket X_i \rrbracket^{(G, \hat{\mathbf{f}}_{\text{subtree}(u)}, \hat{\mathbf{g}})}|.$$

If X_i is an edge set variable, then $c_u^i(P) = c_u^i[P] = |\hat{\mathbf{g}}(X_i)|$ while $c_{\mathbf{f}_u}^i = 0$ holds so the claim follows. Assume now that X_i is a vertex set variable. Then, we have

$$c_u^i(P) = |\hat{\mathbf{f}}(X_i) \cap \text{subtree}[u]| = |\hat{\mathbf{f}}(X_i) \cap \text{subtree}(u)| + |\hat{\mathbf{f}}(X_i) \cap \{u\}| = c_u^i[P] + |\hat{\mathbf{f}}(X_i) \cap \{u\}|.$$

Since $P \in \text{Part}[J, \mathcal{T}_{\mathbb{E}_u}]$, we know that $\mathbf{f}_u$ is the restriction of $\hat{\mathbf{f}}$ on $\{u\}$. Thus, we have $\hat{\mathbf{f}}(X_i) \cap \{u\} = \mathbf{f}_u(X_i)$. By definition, we have $c_{\mathbf{f}_u}^i = |\mathbf{f}_u(X_i)|$. We conclude that $c_u^i(P) = c_u^i[P] + c_{\mathbf{f}_u}^i$. Δ

We are ready to prove (8). By definition, we have

$$\begin{aligned} \mathcal{P}_u(I) &:= \sum_{a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell} |\text{Part}_u(I, a, b, \bar{c})| \cdot \alpha^a \cdot \beta^b \cdot \gamma^{\bar{c}} = \sum_{P \in \text{Part}_u(I)} \alpha^{a_u(P)} \cdot \beta^{b_u(P)} \cdot \gamma^{\bar{c}_u(P)} \\ &\stackrel{\text{Claim 6.12}}{=} \sum_{\substack{J \in \mathcal{J} \\ P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]}} \alpha^{a_u(P)} \cdot \beta^{b_u(P)} \cdot \gamma^{\bar{c}_u(P)} \stackrel{\text{Claim 6.13}}{=} \sum_{\substack{(\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J} \\ P \in \text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}]}} \alpha^{a_u[P] + a_{\mathbb{E}_u}} \cdot \beta^{b_u[P]} \cdot \gamma^{\bar{c}_u[P] + \bar{c}_{\mathbf{f}_u}} \\ &= \sum_{\substack{J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J} \\ a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell}} |\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]| \cdot \alpha^{a + a_{\mathbb{E}_u}} \cdot \beta^b \cdot \gamma^{\bar{c} + \bar{c}_{\mathbf{f}_u}}. \end{aligned}$$

For $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$, from the definition of $\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]$, we have:

$$\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}] = \bigcap_{T \in \mathcal{T}_{\mathbb{E}_u}} \text{Part}_u[J, \{T\}, a, b, \bar{c}].$$

For every $\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}$, we denote by $\overline{\text{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$ the set of all $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \text{Part}_u[J, \emptyset, a, b, \bar{c}]$ such that for every $(C, i) \in \mathcal{T}$, there is no arc in $\vec{\mathcal{W}}(C)$ incoming into u^i . Observe that

$$\overline{\text{Part}}_u[J, \mathcal{T}, a, b, \bar{c}] = \text{Part}_u[J, \emptyset, a, b, \bar{c}] \setminus \left(\bigcup_{T \in \mathcal{T}} \text{Part}_u[J, \{T\}, a, b, \bar{c}] \right).$$

By the inclusion-exclusion principle (Theorem 3.5), we deduce that

$$|\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]| = \sum_{\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}} (-1)^{|\mathcal{T}|} |\overline{\text{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]|. \quad (12)$$

We will prove the “non-DFT-transformed” equality (6) with (8), (12) and the following claim.

Claim 6.14. Let $J = (\mathbf{f}_u, \mathbb{E}_u, \mathbf{pd}_u) \in \mathcal{J}$, $a, b \in \mathbb{N}$, $\bar{c} \in \mathbb{N}^\ell$. For every $\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}$, there exists a neighborhood function $\mathbf{alw}_{\mathcal{T}}$ on $\{u\}$ with $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_{\mathcal{T}}, \mathbf{pd}_u) \in \mathcal{J}'$ such that

$$|\overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]| = |\mathbf{Part}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbf{alw}_{\mathcal{T}}, \mathbf{pd}_u^+, \mathbf{F}, a, b, \bar{c}]|, \text{ and} \quad (13)$$

$$(-1)^{|\mathcal{T}|} = \prod_{C \in \mathfrak{X}} (-1)^{\mathbb{E}_u(u, C) - \mathbf{alw}_{\mathcal{T}}(u, C)}. \quad (14)$$

Furthermore, for every $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u) \in \mathcal{J}'$, we have

$$|\{\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u} : \mathbf{alw}_u = \mathbf{alw}_{\mathcal{T}}\}| = \prod_{C \in \mathfrak{X}} \binom{\mathbb{E}_u(u, C)}{\mathbf{alw}_u(u, C)}. \quad (15)$$

Proof of the claim. First, we fix an arbitrary $\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}$ and construct $\mathbf{alw}_{\mathcal{T}}$ as follows. For every $C \in \mathfrak{X}$, let $\mathbf{Alw}(C)$ be the set of all $i \in [\mathbb{E}_u(u, C)]$ and such that $(C, i) \notin \mathcal{T}$ and we let h_C be a bijection between $\mathbf{Alw}(C)$ and $\{1, 2, \dots, |\mathbf{Alw}(C)|\}$. We define $\mathbf{alw}_{\mathcal{T}}$ as the neighborhood function on $\{u\}$ such that for every $C \in \mathfrak{X}$, we have $\mathbf{alw}_{\mathcal{T}}(u, C) = |\mathbf{Alw}(C)|$. Observe that $\mathbf{alw}_{\mathcal{T}}$ is indeed a neighborhood function because by construction we have $\mathbf{alw}_{\mathcal{T}}(u, C) \in [\mathbb{E}_u(u, C)] \subseteq [d_\varphi]$.

Let $J_{\mathcal{T}}^+ := (\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbf{alw}_{\mathcal{T}}, \mathbf{pd}_u^+, \mathbf{F})$. It is easy to see that $J_{\mathcal{T}}^+ \in \mathcal{J}'$ because by construction we have $\mathbf{alw}_{\mathcal{T}} \leq \mathbb{E}_u$. Next, we provide a bijection between $\overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$ and $\mathbf{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$. For this, for every $C \in \mathfrak{X}$ we define the bijection g_C between the set

$$\{(z^1, w^j) \in \vec{A}(\vec{G}_\varphi) : w \neq z\} \cup \{(z^1, u^j) \in \vec{A}(\vec{G}_\varphi) : j \in \mathbf{Alw}(C)\} \quad (16)$$

and the set

$$\{(z^1, w^j) \in \vec{A}(\vec{G}_\varphi) : w \neq z\} \cup \{(z^1, u^j) \in \vec{A}(\vec{G}_\varphi) : j \in [\mathbf{Alw}(C)]\} \quad (17)$$

via

$$g_C((z^1, w^j)) = \begin{cases} (z^1, w^j) & \text{if } w \neq u \\ (z^1, u^{h_C(j)}) & \text{otherwise.} \end{cases}$$

Let h be the function that maps every $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$ to $h(P) := (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{Z}})$ where

$$\vec{\mathcal{Z}}(C) := \{g_C(a) : a \in \vec{\mathcal{W}}(C)\}$$

for every $C \in \mathfrak{X}$. This is well-defined because for every $C \in \mathfrak{X}$ and every arc $(v^1, u^j) \in \vec{\mathcal{W}}(C)$, we have $j \in \mathbf{Alw}(C)$. Indeed, by definition of $\overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$, we know that $(C, j) \notin \mathcal{T}$ and by (Respect-broom- $\hat{\mathbb{E}}$) and ($\mathbb{E}$ -consistent) we have $j \in [\hat{\mathbb{E}}(u, C)] = [\mathbb{E}_u(u, C)]$. In simple words, for every $C \in \mathfrak{X}$, h “reorders” the arcs in $\vec{\mathcal{W}}(C)$ incoming into $\{u^j : j \in \mathbf{Alw}(C)\}$ in such a way that they keep their tail and their heads are now in $\{u^j : j \leq |\mathbf{Alw}(C)| = \mathbf{alw}_{\mathcal{T}}(u, C)\}$. In particular, g_C is a bijection between $\vec{\mathcal{W}}(C)$ and $\vec{\mathcal{Z}}(C)$.

We first prove that (1) h maps every element of $\overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$ to an element in $\mathbf{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$, and (2) every element in $\mathbf{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$ has a preimage under h in $\overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$. Later we show that the bijectivity easily follows from these two properties.

Let $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$ and let $h(P) = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{Z}})$. We need to check that $h(P) \in \mathbf{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$. First observe that $h(P)$ is a partial solution for $\text{tss}[u]$. This is clear from the construction of $\vec{\mathcal{Z}}$ and the fact that g_C is a bijection between $\vec{\mathcal{W}}(C)$ and $\vec{\mathcal{Z}}(C)$ for every $C \in \mathfrak{X}$. In particular, (Respect-broom- $\hat{\mathbb{E}}$) holds because $\{1, 2, \dots, |\mathbf{Alw}(C)|\} \subseteq [\mathbb{E}_u(u, C)] = [\hat{\mathbb{E}}(u, C)]$ for every $C \in \mathfrak{X}$. Now, we show that $h(P)$ is compatible with $J_{\mathcal{T}}^+$. Since $P \in \overline{\mathbf{Part}}_u[J, \mathcal{T}, a, b, \bar{c}] \subseteq \mathbf{Part}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw} \cup \mathbb{E}_u, \mathbf{pd}_u^+, \mathbf{F}, a, b, \bar{c}]$, we only need to check the conditions involving $\vec{\mathcal{Z}}$ and $\mathbf{alw}_{\mathcal{T}}$,

that is: **(pd-consistent)** and **(Respect- alw)**. By construction of $\vec{\mathcal{Z}}$, for every $v \in \text{tail}[u]$ and $C \in \mathfrak{K}$, we have

$$\text{pd}(v, C) = \text{pd}_{\vec{\mathcal{W}}}(v, C) = |\{(w^1, v^j) \in \vec{\mathcal{W}}(C)\}| = |\{(w^1, v^j) \in \vec{\mathcal{Z}}(C)\}| = \text{pd}_{\vec{\mathcal{Z}}}(v, C),$$

where the second equality holds because g_C is a bijection between $\vec{\mathcal{W}}(C)$ and $\vec{\mathcal{Z}}(C)$ that maps, for every vertex w of G , all arcs incoming into $w^1, \dots, w^{d_\varphi}$ to arcs incoming into vertices $w^1, \dots, w^{d_\varphi}$. Thus, $h(P)$ satisfies **(pd-consistent)**. By construction, for every $C \in \mathfrak{K}$ and arc $(v^1, w^j) \in \vec{\mathcal{W}}(C)$ with $w \neq u$, we have $(v^1, w^j) \in \vec{\mathcal{Z}}(C)$. The tuple P satisfies **(Respect- alw)** on $\text{tail}(u)$ with the function $\text{alw} \cup \mathbb{E}_u$ so $h(P)$ satisfies **(Respect- alw)** on $\text{tail}(u)$ with the function $\text{alw} \cup \text{alw}_{\mathcal{T}}$. Moreover, for every $C \in \mathfrak{K}$ and every arc $(v^1, u^j) \in \vec{\mathcal{Z}}(C)$, we have $j \leq [\text{Alw}(C)] = \text{alw}_{\mathcal{T}}(u, C)$ by construction. Clearly, $a_u[P] = a_u[h(P)]$ and $\bar{c}_u[P] = \bar{c}_u[h(P)]$. Moreover, since the number of arcs incoming into $v^1, \dots, v^{d_\varphi}$ is the same in $\vec{\mathcal{W}}(C)$ and $\vec{\mathcal{Z}}(C)$ for every $C \in \mathfrak{K}$, we have $b_u[P] = b_u[h(P)]$. We deduce that $h(P) \in \text{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$.

Let $P' = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{Z}}) \in \text{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$. Let $\vec{\mathcal{W}}$ be the function that maps every $C \in \mathfrak{K}$ to the set $\{g_C^{-1}(a) : a \in \vec{\mathcal{Z}}(C)\}$. This is well-defined because for every $(v^1, u^j) \in \vec{\mathcal{Z}}(C)$, we have $j \leq \text{alw}_{\mathcal{T}}(u, C) = [\text{Alw}(C)]$ due to **(Respect- alw)**, so $g_C^{-1}((v^1, u^j)) = (v^1, u^{h_C^{-1}(j)})$ exists and g_C is a bijection between $\vec{\mathcal{W}}(C)$ and $\vec{\mathcal{Z}}(C)$. We claim that $P := (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \overline{\text{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$. With the same arguments used above, we can prove that P is a partial solution for $\text{tss}[u]$, we have $a_u[P'] = a_u[P]$, $b_u[P'] = b_u[P]$, $\bar{c}_u[P'] = \bar{c}_u[P]$ and P satisfies all the conditions to be compatible with $(\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw} \cup \mathbb{E}_u, \text{pd}_u^+, \mathbf{F})$ at the exception of **(Respect- alw)**. This latter condition is satisfied by P for $\text{tail}(u)$ because P' satisfies **(Respect- alw)** for $\text{tail}(u)$ and the function $\text{alw} \cup \text{alw}_{\mathcal{T}}$, and for $\{u\}$ it is satisfied because P' satisfies **(Respect-broom- $\hat{\mathbb{E}}$)** and it satisfies **($\mathbb{E}$ -consistent)** for $\mathbb{E}_u$. We deduce that $P \in \text{Part}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw} \cup \mathbb{E}_u, \text{pd}_u^+, \mathbf{F}, a, b, \bar{c}]$. Furthermore, observe that by construction, for every $C \in \mathfrak{K}$ and arc $(v^1, u^i) \in \vec{\mathcal{W}}(C)$, we have $i \in \text{Alw}(C)$, i.e., $(C, i) \notin \mathcal{T}$. It follows that $P \in \overline{\text{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$, and, by the construction of $\vec{\mathcal{W}}$, it holds that $h(P) = P'$.

Finally, recall that g_C is a bijection between (16) and (17) for every $C \in \mathfrak{K}$. Thus h is injective and every partial solution in $\text{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$ has at most one preimage in h . Altogether, we obtain that h is a bijection between $\overline{\text{Part}}_u[J, \mathcal{T}, a, b, \bar{c}]$ and $\text{Part}_u[J_{\mathcal{T}}^+, a, b, \bar{c}]$. We conclude that (13) is true.

Observe that (14) simply follows from the following equation:

$$|\mathcal{T}| = \sum_{C \in \mathfrak{K}} |\{(C, i) \in \mathcal{T} : i \in [\mathbb{E}_u(u, C)]\}| = \sum_{C \in \mathfrak{K}} \mathbb{E}_u(u, C) - |\text{Alw}(C)| = \sum_{C \in \mathfrak{K}} \mathbb{E}_u(u, C) - \text{alw}_{\mathcal{T}}(u, C).$$

Finally, to prove (8), observe that for every neighborhood function alw_u on $\{u\}$ such that $\text{alw}_u \leq \mathbb{E}_u$ and every $\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}$, we have $\text{alw}_u = \text{alw}_{\mathcal{T}}$ if and only if $|\mathcal{T}(C)| = \mathbb{E}_u(u, C) - \text{alw}_u(u, C)$ for every $C \in \mathfrak{K}$. Since $\mathcal{T}(C) \subseteq [\mathbb{E}_u(u, C)]$, there are $\binom{\mathbb{E}_u(u, C)}{\mathbb{E}_u(u, C) - \text{alw}_u(u, C)} = \binom{\mathbb{E}_u(u, C)}{\text{alw}_u(u, C)}$ possibilities for each $\mathcal{T}(C)$ so that $\text{alw}_u = \text{alw}_{\mathcal{T}}$ holds. Δ

In the following, given $J' = (\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'$, we define the tuple

$$J'^+ := (\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw}_u^+, \text{pd}_u^+, \mathbf{F}).$$

Thanks to Claim 6.14, we can reformulate $|\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]|$ as follows.

$$\begin{aligned}
|\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]| &\stackrel{(12)}{=} \sum_{\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u}} (-1)^{|\mathcal{T}|} |\overline{\text{Part}_u}[J, \mathcal{T}, a, b, \bar{c}]| \\
&\stackrel{6.14}{=} \sum_{\substack{\text{alw}_u: \\ J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'}} \sum_{\substack{\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u} \\ \text{alw}_u = \text{alw}_{\mathcal{T}}}} (-1)^{|\mathcal{T}|} |\text{Part}_u[J'^+, a, b, \bar{c}]| \\
&\stackrel{(14)}{=} \sum_{\substack{\text{alw}_u: \\ J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'}} \sum_{\substack{\mathcal{T} \subseteq \mathcal{T}_{\mathbb{E}_u} \\ \text{alw}_u = \text{alw}_{\mathcal{T}}}} \left(\prod_{C \in \mathbb{X}} (-1)^{\mathbb{E}_u(u, C) - \text{alw}_u(u, C)} \right) \cdot |\text{Part}_u[J'^+, a, b, \bar{c}]| \\
&\stackrel{(15)}{=} \sum_{\substack{\text{alw}_u: \\ J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'}} \left(\prod_{C \in \mathbb{X}} (-1)^{\mathbb{E}_u(u, C) - \text{alw}_u(u, C)} \cdot \binom{\mathbb{E}_u(u, C)}{\text{alw}_u(u, C)} \right) \cdot |\text{Part}_u[J'^+, a, b, \bar{c}]| \\
&= \sum_{\substack{\text{alw}_u: \\ J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'}} S_{\pm}(\mathbb{E}_u, \text{alw}_u) \cdot |\text{Part}_u[J'^+, a, b, \bar{c}]| \tag{19}
\end{aligned}$$

We are now ready to prove (6):

$$\begin{aligned}
\mathcal{P}_u(I) &\stackrel{(8)}{=} \sum_{\substack{J=(\mathbf{f}_u, \mathbb{E}_u, \text{pd}_u) \in \mathcal{J} \\ a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell}} |\text{Part}_u[J, \mathcal{T}_{\mathbb{E}_u}, a, b, \bar{c}]| \cdot \alpha^{a+a_{\mathbb{E}_u}} \cdot \beta^b \cdot \gamma^{\bar{c}+\bar{c}_{\mathbf{f}_u}} \\
&\stackrel{(19)}{=} \sum_{\substack{(\mathbf{f}_u, \mathbb{E}_u, \text{pd}_u) \in \mathcal{J} \\ a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell}} \sum_{\substack{\text{alw}_u: \\ J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'}} S_{\pm}(\mathbb{E}_u, \text{alw}_u) \cdot |\text{Part}_u[J'^+, a, b, \bar{c}]| \cdot \alpha^{a+a_{\mathbb{E}_u}} \cdot \beta^b \cdot \gamma^{\bar{c}+\bar{c}_{\mathbf{f}_u}} \\
&= \sum_{(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'} S_{\pm}(\mathbb{E}_u, \text{alw}_u) \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw}_u^+, \text{pd}_u^+, \mathbf{F}] \cdot \alpha^{a_{\mathbb{E}_u}} \cdot \gamma^{\bar{c}_{\mathbf{f}_u}}.
\end{aligned}$$

Now we will use this result to prove the “DFT-transformed” equality (7). In the following, given $J' = (\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'$, we denote by $M(J')$ the product $S_{\pm}(\mathbb{E}_u, \text{alw}_u) \cdot \alpha^{a_{\mathbb{E}_u}} \cdot \gamma^{\bar{c}_{\mathbf{f}_u}}$. We remind the reader that we aim at proving that

$$\mathcal{Q}_u^p(I) \equiv_p \sum_{J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'} M(J') \cdot \frac{1}{p_{\varphi}^{|\mathbb{X}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega(p)^{-\hat{q} \cdot \text{pd}_u} \cdot \mathcal{Q}_u^p[\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw}_u^+, \text{pd} \cup \hat{q}, \mathbf{F}]$$

for every prime number p such that $\mathbb{F}_p^*$ admits a p_{φ} -th root of unity. For the remainder of the proof, we fix such an arbitrary prime number p . In the following $\equiv$ denotes $\equiv_p$ and ω denotes $\omega(p)$. We have:

$$\begin{aligned}
\mathcal{Q}_u^p(I) &\stackrel{(3)}{=} \sum_{q \in \text{PD}(\text{tail}(u))} \omega^{q \cdot \text{pd}} \cdot \mathcal{P}_u(\mathbf{f}, \mathbb{E}, \text{alw}, q, \mathbf{F}) \\
&\stackrel{(6)}{=} \sum_{q \in \text{PD}(\text{tail}(u))} \omega^{q \cdot \text{pd}} \cdot \left(\sum_{J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'} M(J') \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw}_u^+, q \cup \text{pd}_u, \mathbf{F}] \right) \\
&\equiv \sum_{J'=(\mathbf{f}_u, \mathbb{E}_u, \text{alw}_u, \text{pd}_u) \in \mathcal{J}'} M(J') \cdot \sum_{q \in \text{PD}(\text{tail}(u))} \left(\omega^{q \cdot \text{pd}} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \text{alw}_u^+, q \cup \text{pd}_u, \mathbf{F}] \right)
\end{aligned}$$

Now to prove (7), it suffices to show that for each $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u) \in \mathcal{J}'$, we have

$$\sum_{q \in \text{PD}(\text{tail}(u))} \omega^{q \cdot \mathbf{pd}} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, q \cup \mathbf{pd}_u, \mathbf{F}] \equiv \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \cdot \mathcal{Q}_u^p[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, \mathbf{pd} \cup \hat{q}, \mathbf{F}]. \quad (20)$$

Let $(\mathbf{f}_u, \mathbb{E}_u, \mathbf{alw}_u, \mathbf{pd}_u) \in \mathcal{J}'$. To show this, we start by rewriting the right-hand side of this equality which we denote by RHS. We rely on the fact that $\text{tail}[u] = \text{tail}(u) \cup \{u\}$ implies that

$$\text{PD}(\text{tail}[u]) = \{\tilde{q} : \tilde{q} : \text{tail}[u] \times \mathbb{K} \rightarrow \mathbb{Z}_{p_\varphi}\} = \{\tilde{q}_u \cup q : \tilde{q}_u \in \text{PD}(\{u\}), q \in \text{PD}(\text{tail}(u))\}. \quad (21)$$

Then we have:

$$\begin{aligned} \text{RHS} &\stackrel{\text{Def.}}{\equiv_p} \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \sum_{\tilde{q} \in \text{PD}(\text{tail}[u])} \omega^{\tilde{q} \cdot (\mathbf{pd} \cup \hat{q})} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, \tilde{q}, \mathbf{F}] \\ &\stackrel{(21)}{\equiv_p} \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \sum_{\substack{\tilde{q}_u \in \text{PD}(\{u\}), \\ q \in \text{PD}(\text{tail}(u))}} \omega^{(q \cup \tilde{q}_u) \cdot (\mathbf{pd} \cup \hat{q})} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, q \cup \tilde{q}_u, \mathbf{F}] \\ &\equiv_p \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \sum_{\substack{\tilde{q}_u \in \text{PD}(\{u\}), \\ q \in \text{PD}(\text{tail}(u))}} \omega^{q \cdot \mathbf{pd}} \cdot \omega^{\tilde{q}_u \cdot \hat{q}} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, q \cup \tilde{q}_u, \mathbf{F}] \\ &\equiv_p \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \sum_{\tilde{q}_u \in \text{PD}(\{u\})} \omega^{\tilde{q}_u \cdot \hat{q}} \sum_{q \in \text{PD}(\text{tail}(u))} \omega^{q \cdot \mathbf{pd}} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, q \cup \tilde{q}_u, \mathbf{F}] \end{aligned}$$

Now we define the function $\Delta : \text{PD}(\{u\}) \rightarrow \mathbb{F}_p[\alpha, \beta, \gamma_1, \dots, \gamma_\ell]$ via

$$\Delta(\tilde{q}_u) := \sum_{q \in \text{PD}(\text{tail}(u))} \omega^{q \cdot \mathbf{pd}} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, q \cup \tilde{q}_u, \mathbf{F}] \pmod{p}.$$

By inserting this definition in the above equality we obtain

$$\begin{aligned} \text{RHS} &\equiv_p \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \sum_{\tilde{q}_u \in \text{PD}(\{u\})} \omega^{\tilde{q}_u \cdot \hat{q}} \cdot \Delta(\tilde{q}_u) \\ &\equiv_p \frac{1}{p_\varphi^{|\mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \cdot (\text{DFT}_{\{u\} \times \mathbb{K}, p_\varphi, p}(\Delta))(\hat{q}) \\ &\equiv_p \frac{1}{p_\varphi^{|\{u\} \times \mathbb{K}|}} \sum_{\hat{q} \in \text{PD}(\{u\})} \omega^{-\hat{q} \cdot \mathbf{pd}_u} \cdot (\text{DFT}_{\{u\} \times \mathbb{K}, p_\varphi, p}(\Delta))(\hat{q}) \\ &\equiv_p \left(\text{DFT}_{\{u\} \times \mathbb{K}, p_\varphi, p}^{-1} \left(\text{DFT}_{\{u\} \times \mathbb{K}, p_\varphi, p}(\Delta) \right) \right)(\mathbf{pd}_u) \\ &\stackrel{A.3}{\equiv_p} \Delta(\mathbf{pd}_u) \equiv_p \sum_{q \in \text{PD}(\text{tail}(u))} \omega^{q \cdot \mathbf{pd}} \cdot \mathcal{P}_u[\mathbf{f}_u^+, \mathbb{E}_u^+, \mathbf{alw}_u^+, q \cup \mathbf{pd}_u, \mathbf{F}] \end{aligned}$$

proving the desired equality (20) and thus concluding the proof. $\square$

6.5 Join equality

And finally we prove the equality relating the partial solutions of $\text{tss}[u]$ and of $\text{tss}(v_1), \dots, \text{tss}(v_t)$ for a node u having children $v_1, \dots, v_t$. For an index $I = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}, \mathbf{F}) \in \text{Index}[u]$ and an assignment $\mathbf{pd}' \in \text{PD}(\text{tail}[u])$ we define the index $I[\mathbf{pd} \mapsto \mathbf{pd}'] = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}', \mathbf{F}) \in \text{Index}[u]$.

Lemma 6.15. *Let u be an internal node of T with children $v_1, \dots, v_t$. Then we have*

$$\mathcal{P}_u[I] = \sum_{\substack{\mathbf{pd}_1, \dots, \mathbf{pd}_t \in \text{PD}(\text{tail}[u]): \\ \mathbf{pd}_1 + \dots + \mathbf{pd}_t \equiv_{p\varphi} \mathbf{pd}}} \prod_{i=1}^t \mathcal{P}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i]) \quad (22)$$

and

$$\mathcal{Q}_u^p[I] \equiv_p \prod_{i=1}^t \mathcal{Q}_{v_i}^p(I) \quad (23)$$

for every $I = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}, \mathbf{F}) \in \text{Index}[u]$.

Proof. We start by proving the “non-transformed” equality (22). So let $I = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}, \mathbf{F}) \in \text{Index}[u]$ be an arbitrary but fixed index. First, we rewrite both sides of the desired equality. For the right-hand side we have:

$$\begin{aligned} & \sum_{\substack{\mathbf{pd}_1, \dots, \mathbf{pd}_t \in \text{PD}(\text{tail}[u]): \\ \mathbf{pd}_1 + \dots + \mathbf{pd}_t \equiv_{p\varphi} \mathbf{pd}}} \prod_{i=1}^t \mathcal{P}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i]) = \\ & \sum_{\substack{\mathbf{pd}_1, \dots, \mathbf{pd}_t \in \text{PD}(\text{tail}[u]): \\ \mathbf{pd}_1 + \dots + \mathbf{pd}_t \equiv_{p\varphi} \mathbf{pd}}} \prod_{i=1}^t \sum_{a, b, \bar{c} \in \mathbb{N}} |\text{Part}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i], a, b, \bar{c})| \alpha^a \beta^b \gamma^{\bar{c}} = \\ & \sum_{\substack{a_1, \dots, a_t \in \mathbb{N}, \\ b_1, \dots, b_t \in \mathbb{N}, \\ \bar{c}_1, \dots, \bar{c}_t \in \mathbb{N}^\ell, \\ \mathbf{pd}_1, \dots, \mathbf{pd}_t \in \text{PD}(\text{tail}[u]): \\ \mathbf{pd}_1 + \dots + \mathbf{pd}_t \equiv_{p\varphi} \mathbf{pd}}} \prod_{i=1}^t |\text{Part}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i], a_i, b_i, \bar{c}_i)| \alpha^{a_i} \beta^{b_i} \gamma^{\bar{c}_i} = \\ & \sum_{\substack{a_1, \dots, a_t \in \mathbb{N}, \\ b_1, \dots, b_t \in \mathbb{N}, \\ \bar{c}_1, \dots, \bar{c}_t \in \mathbb{N}^\ell, \\ \mathbf{pd}_1, \dots, \mathbf{pd}_t \in \text{PD}(\text{tail}[u]): \\ \mathbf{pd}_1 + \dots + \mathbf{pd}_t \equiv_{p\varphi} \mathbf{pd}}} \alpha^{\sum_{i=1}^t a_i} \beta^{\sum_{i=1}^t b_i} \gamma^{\sum_{i=1}^t \bar{c}_i} \prod_{i=1}^t |\text{Part}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i], a_i, b_i, \bar{c}_i)| = \\ & \sum_{a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell} \alpha^a \beta^b \gamma^{\bar{c}} \sum_{\substack{a_1 + \dots + a_t = a, \\ b_1 + \dots + b_t = b, \\ \bar{c}_1 + \dots + \bar{c}_t = \bar{c}, \\ \mathbf{pd}_1, \dots, \mathbf{pd}_t \in \text{PD}(\text{tail}[u]): \\ \mathbf{pd}_1 + \dots + \mathbf{pd}_t \equiv_{p\varphi} \mathbf{pd}}} \prod_{i=1}^t |\text{Part}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i], a_i, b_i, \bar{c}_i)| \end{aligned}$$

where $a_1 + \dots + a_t = a$ is a shortcut for $a_1, \dots, a_t \in \mathbb{N}: \sum_{i=1}^t a_i = a$, also $b_1 + \dots + b_t = b$ is a shortcut for $b_1, \dots, b_t \in \mathbb{N}: \sum_{i=1}^t b_i = b$, and $\bar{c}_1 + \dots + \bar{c}_t = \bar{c}$ is a shortcut for $\bar{c}_1, \dots, \bar{c}_t \in \mathbb{N}^\ell: \sum_{i=1}^t \bar{c}_i = \bar{c}$. And the left-hand side of (22) can be rewritten as follows:

$$\mathcal{P}_u[I] = \sum_{a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell} |\text{Part}_u[I, a, b, \bar{c}]| \cdot \alpha^a \beta^b \gamma^{\bar{c}}.$$

So to prove (22), it suffices to show that we have:

$$|\text{Part}_u[I, a, b, \bar{c}]| = \sum_{\substack{a_1 + \dots + a_t = a, \\ b_1 + \dots + b_t = b, \\ \bar{c}_1 + \dots + \bar{c}_t = \bar{c}, \\ \text{pd}_1, \dots, \text{pd}_t \in \text{PD}(\text{tail}[u]): \\ \text{pd}_1 + \dots + \text{pd}_t \equiv_{p_\varphi} \text{pd}}} \prod_{i=1}^t |\text{Part}_{v_i}(I[\text{pd} \mapsto \text{pd}_i], a_i, b_i, \bar{c}_i)| \quad (24)$$

for all $a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell$. To obtain the desired equality, we fix arbitrary $a, b \in \mathbb{N}, \bar{c} \in \mathbb{N}^\ell$, define the set

$$\begin{aligned} \text{RHS} := & \left\{ (P_1, \dots, P_t) : \right. \\ & \exists \text{pd}_1, \dots, \text{pd}_t \in \text{PD}(\text{tail}[u]), a_1, \dots, a_t, b_1, \dots, b_t \in \mathbb{N}, \bar{c}_1, \dots, \bar{c}_t \in \mathbb{N}^\ell : \\ & \sum_{i=1}^t \text{pd}_i \equiv_{p_\varphi} \text{pd}, \sum_{i=1}^t a_i = a, \sum_{i=1}^t b_i = b, \sum_{i=1}^t \bar{c}_i = \bar{c}, \\ & \left. P_i \in \text{Part}_{v_i}(I[\text{pd} \mapsto \text{pd}_i], a_i, b_i, \bar{c}) \text{ for every } i \in [t] \right\}. \end{aligned}$$

Observe that by Definition 5.13 and Definition 5.15 the sets $\text{Part}_{v_i}(I[\text{pd} \mapsto \text{pd}'], a', b', \bar{c}')$ and $\text{Part}_{v_i}(I[\text{pd} \mapsto \text{pd}''], a'', b'', \bar{c}'')$ are disjoint for every pair $(\text{pd}', a', b', \bar{c}') \neq (\text{pd}'', a'', b'', \bar{c}'') \in \text{PD}(\text{tail}[u]) \times \mathbb{N}^{\ell+2}$ and every $i \in [t]$. Therefore, the right-hand side of (24) is equal to the cardinality of RHS. So to prove (24) it remains to provide an injective mapping from $\text{Part}_u[I, a, b, \bar{c}]$ to RHS as well as an injective mapping from RHS to $\text{Part}_u[I, a, b, \bar{c}]$.

Observe that the tail of each triad in $\text{tss}[u], \text{tss}(v_1), \dots, \text{tss}(v_t)$ is equal to $\text{tail}[u]$, the sets $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ partition the set $\text{sheaf}(u)$, the sets $\text{subtree}(v_1), \dots, \text{subtree}(v_t)$ partition the set $\text{subtree}(u)$, and we have $\text{broom}[u] = \cup_{i \in [t]} \text{broom}[v_i]$. We will rely on these properties in the following.

Claim 6.16. *There exists an injective mapping h_1 from $\text{Part}_u[I, a, b, \bar{c}]$ to RHS.*

Proof of the claim. For the one direction, we map every partial solution $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ in $\text{Part}_u[I, a, b, \bar{c}]$ to the tuple $h_1(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) = (P_1, \dots, P_t)$ where for every $i \in [t]$, $P_i = (\hat{\mathbf{f}}_i, \hat{\mathbf{g}}_i, \hat{\mathbb{E}}_i, \vec{\mathcal{W}}_i)$ such that

- $\hat{\mathbf{f}}_i$ is the restriction $\hat{\mathbf{f}}|_{\text{broom}[v_i]}$ of $\hat{\mathbf{f}}$ on $\text{broom}[v_i]$,
- $\hat{\mathbf{g}}_i$ is the restriction $\hat{\mathbf{g}}|_{\text{sheaf}(v_i)}$ of $\hat{\mathbf{g}}$ on $\text{sheaf}(v_i)$,
- $\hat{\mathbb{E}}_i$ is the restriction $\hat{\mathbb{E}}|_{\text{broom}[v_i]}$ of $\hat{\mathbb{E}}$ on $\text{broom}[v_i]$,
- and $\vec{\mathcal{W}}_i$ is the restriction $\vec{\mathcal{W}}|_{\text{sheaf}(v_i)}$ of $\vec{\mathcal{W}}$ to $\text{sheaf}(v_i)$.

Before showing that this is indeed a mapping into RHS, let us start with injectivity since it is simpler. For this we consider two elements $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ and $Q = (\hat{\mathbf{f}}', \hat{\mathbf{g}}', \hat{\mathbb{E}}', \vec{\mathcal{W}}')$ in $\text{Part}_u[I, a, b, \bar{c}]$ such that $h_1(P) = h_1(Q)$. This implies the following:

- For every $X \in \text{var}^V(\varphi)$ we have

$$\begin{aligned}
\hat{\mathbf{f}}(X) &= \hat{\mathbf{f}}(X) \cap \text{broom}[u] = \hat{\mathbf{f}}(X) \cap (\text{broom}[v_1] \cup \dots \cup \text{broom}[v_t]) \\
&= (\hat{\mathbf{f}}(X) \cap \text{broom}[v_1]) \cup \dots \cup (\hat{\mathbf{f}}(X) \cap \text{broom}[v_t]) = \hat{\mathbf{f}}_1(X) \cup \dots \cup \hat{\mathbf{f}}_t(X) \\
h_1(P_1) &\stackrel{=}{=} h_1(P_2) \hat{\mathbf{f}}'_1(X) \cup \dots \cup \hat{\mathbf{f}}'_t(X) = (\hat{\mathbf{f}}'(X) \cap \text{broom}[v_1]) \cup \dots \cup (\hat{\mathbf{f}}'(X) \cap \text{broom}[v_t]) \\
&= \hat{\mathbf{f}}'(X) \cap (\text{broom}[v_1] \cup \dots \cup \text{broom}[v_t]) = \hat{\mathbf{f}}'(X) \cap \text{broom}[u] = \hat{\mathbf{f}}'(X).
\end{aligned}$$

Hence $\hat{\mathbf{f}} = \hat{\mathbf{f}}'$.

- For every $Y \in \text{var}^E(\varphi)$ we have

$$\begin{aligned}
\hat{\mathbf{g}}(Y) &= \hat{\mathbf{g}}(Y) \cap \text{sheaf}(u) = \hat{\mathbf{g}}(Y) \cap (\text{sheaf}(v_1) \cup \dots \cup \text{sheaf}(v_t)) \\
&= (\hat{\mathbf{g}}(Y) \cap \text{sheaf}(v_1)) \cup \dots \cup (\hat{\mathbf{g}}(Y) \cap \text{sheaf}(v_t)) = \hat{\mathbf{g}}_1(Y) \cup \dots \cup \hat{\mathbf{g}}_t(Y) \\
h_1(P_1) &\stackrel{=}{=} h_1(P_2) \hat{\mathbf{g}}'_2(Y) \cup \dots \cup \hat{\mathbf{g}}'_t(Y) = (\hat{\mathbf{g}}'(Y) \cap \text{sheaf}(v_1)) \cup \dots \cup (\hat{\mathbf{g}}'(Y) \cap \text{sheaf}(v_t)) \\
&= \hat{\mathbf{g}}'(Y) \cap (\text{sheaf}(v_1) \cup \dots \cup \text{sheaf}(v_t)) = \hat{\mathbf{g}}'(Y) \cap \text{sheaf}(u) = \hat{\mathbf{g}}'(Y).
\end{aligned}$$

Hence $\hat{\mathbf{g}} = \hat{\mathbf{g}}'$.

- For every vertex $z \in \text{subtree}(u)$ there exists a (unique) index $i \in [t]$ with $z \in \text{subtree}[v_i]$. Then for every $C \in \mathfrak{K}$ we have

$$\hat{\mathbb{E}}(z, C) = \hat{\mathbb{E}}_i(z, C) \stackrel{h_1(P_1) = h_1(P_2)}{=} \hat{\mathbb{E}}'_i(z, C) = \hat{\mathbb{E}}'(z, C).$$

And for every vertex $z \in \text{tail}[u]$ and every $C \in \mathfrak{K}$ we have $\hat{\mathbb{E}}(z, C) = \mathbb{E}(z, C) = \hat{\mathbb{E}}'(z, C)$ since both P_1 and P_2 satisfy **(E-consistent)** for I . So we get $\hat{\mathbb{E}}_1 = \hat{\mathbb{E}}_2$.

- And for every $C \in \mathfrak{K}$ we have

$$\begin{aligned}
\vec{\mathcal{W}}(C) &= \vec{\mathcal{W}}(C) \cap \vec{\text{sheaf}}(u) = \vec{\mathcal{W}}(C) \cap (\vec{\text{sheaf}}(v_1) \cup \dots \cup \vec{\text{sheaf}}(v_t)) \\
&= (\vec{\mathcal{W}}(C) \cap \vec{\text{sheaf}}(v_1)) \cup \dots \cup (\vec{\mathcal{W}}(C) \cap \vec{\text{sheaf}}(v_t)) = \vec{\mathcal{W}}_1(C) \cup \dots \cup \vec{\mathcal{W}}_t(C) \\
h_1(P_1) &\stackrel{=}{=} h_1(P_2) \vec{\mathcal{W}}'_1(C) \cup \dots \cup \vec{\mathcal{W}}'_t(C) = (\vec{\mathcal{W}}'(C) \cap \vec{\text{sheaf}}(v_1)) \cup \dots \cup (\vec{\mathcal{W}}'(C) \cap \vec{\text{sheaf}}(v_t)) \\
&= \vec{\mathcal{W}}'(C) \cap (\vec{\text{sheaf}}(v_1) \cup \dots \cup \vec{\text{sheaf}}(v_t)) = \vec{\mathcal{W}}'(C) \cap \vec{\text{sheaf}}(u) = \vec{\mathcal{W}}'(C)
\end{aligned}$$

Hence, $\vec{\mathcal{W}} = \vec{\mathcal{W}}'$.

Altogether, we obtain $P_1 = P_2$ implying the injectivity of h_1 .

It remains to show that for every $P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) \in \mathbf{Part}_u[I, a, b, \bar{c}]$, its image $h_1(P)$ belongs to RHS. For this, we first define the following values for every $i \in [t]$ and every $j \in [\ell]$:

$$\begin{aligned}
a_i &= \sum_{C \in \mathfrak{K}} \sum_{\substack{v \in \text{subtree}[v_i]: \\ \mathbb{E}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C), \\
b_i &= \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[v_i] \\ \hat{\mathbb{E}}(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\
c_i^j &= \begin{cases} |\hat{\mathbf{f}}(X_j) \cap \text{subtree}[v_i]| & \text{if } X_j \in \text{var}^V(\varphi) \\ |\hat{\mathbf{g}}(X_j) \cap \text{sheaf}(v_i)| & \text{if } X_j \in \text{var}^E(\varphi) \end{cases}.
\end{aligned}$$

For $i \in [t]$, we then define the vector $\bar{c}_i = (c_i^1, \dots, c_i^\ell) \in \mathbb{N}^\ell$. We then also have the following for every $i \in [t]$ and every $j \in [\ell]$:

$$a_i = \sum_{C \in \mathfrak{K}} \sum_{\substack{v \in \text{subtree}[v_i]: \\ \mathbb{E}_i(v, C) < d_\varphi}} \hat{\mathbb{E}}_i(v, C), \quad (25)$$

$$b_i = \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[v_i]: \\ \mathbb{E}_i(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \quad (26)$$

$$c_i^j = \begin{cases} |\hat{\mathbf{f}}_i(X_j) \cap \text{subtree}[v_i]| & \text{if } X_j \in \text{var}^V(\varphi) \\ |\hat{\mathbf{g}}_i(X_j)| & \text{if } X_j \in \text{var}^E(\varphi) \end{cases}. \quad (27)$$

Finally, for every $i \in [t]$ we define the mapping $\text{pd}_i: \text{tail}(v_i) \times \mathfrak{K} \rightarrow \mathbb{F}_{p_\varphi}$ to be equal to $\text{pd}_{\vec{\mathcal{W}}_i}$. We then have

$$\begin{aligned} \sum_{i=1}^t a_i &= \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{v \in \text{subtree}[v_i]: \\ \mathbb{E}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) = \sum_{C \in \mathfrak{K}} \sum_{i=1}^t \sum_{\substack{v \in \text{subtree}[v_i]: \\ \mathbb{E}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) \\ \text{subtree}(u) &= \text{subtree}[v_1] \dot{\cup} \dots \dot{\cup} \text{subtree}[v_t] \quad \sum_{C \in \mathfrak{K}} \sum_{\substack{v \in \text{subtree}(u): \\ \mathbb{E}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) = a_u[P]^{P \in \text{Part}_u[I, a, b, \bar{c}]} a \end{aligned}$$

and

$$\begin{aligned} \sum_{i=1}^t b_i &= \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[v_i]: \\ \mathbb{E}(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\ &\stackrel{\substack{\mathcal{W}_i(C) \subseteq \vec{\text{sheaf}}(v_i), \\ V(\text{sheaf}(v_i)) \subseteq \text{broom}[v_i] \subseteq \text{broom}[u]}}{=} \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u]: \\ \mathbb{E}(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\ &= \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u]: \\ \mathbb{E}(w, C) < d_\varphi}} \sum_{i=1}^t |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\ \text{sheaf}(u) &= \text{sheaf}(v_1) \dot{\cup} \dots \dot{\cup} \text{sheaf}(v_t) \quad \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u]: \\ \mathbb{E}(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}(C)\}| = b_u[P]^{P \in \text{Part}_u[I, a, b, \bar{c}]} b. \end{aligned}$$

Also for $j \in [\ell]$, if $X_j \in \text{var}^V(\varphi)$, we have:

$$\sum_{i=1}^t c_i^j = \sum_{i=1}^t |\hat{\mathbf{f}}(X_j) \cap \text{subtree}[v_i]| \stackrel{\text{subtree}(u) = \text{subtree}[v_1] \dot{\cup} \dots \dot{\cup} \text{subtree}[v_t]}{=} |\hat{\mathbf{f}}(X_j) \cap \text{subtree}(u)| = c_u^j[P]^{P \in \text{Part}_u[I, a, b, \bar{c}]} c^j$$

and if $X_j \in \text{var}^E(\varphi)$, we have

$$\sum_{i=1}^t c_i^j = \sum_{i=1}^t |\hat{\mathbf{g}}(X_j) \cap \text{sheaf}(v_i)| \stackrel{\text{sheaf}(u) = \text{sheaf}(v_1) \dot{\cup} \dots \dot{\cup} \text{sheaf}(v_t)}{=} |\hat{\mathbf{g}}(X_j)| = c_u^j[P]^{P \in \text{Part}_u[I, a, b, \bar{c}]} c^j.$$

We recall that we have $\text{tail}[u] = \text{tail}(v_i)$ for every $i \in [t]$. Then the following holds for every $w \in \text{tail}[u]$ and every $C \in \mathfrak{K}$:

$$\sum_{i=1}^t \text{pd}_i(w, C) = \sum_{i=1}^t \text{pd}_{\vec{\mathcal{W}}_i}(w, C) \equiv_{p_\varphi} \sum_{i=1}^t |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}|$$

$$\text{sheaf}(u) = \text{sheaf}(v_1) \dot{\cup} \dots \dot{\cup} \text{sheaf}(v_t) \Big| \{(z^1, w^j) \in \vec{\mathcal{W}}(C)\} \Big| \stackrel{P \in \text{Part}_u[I, a, b, \bar{c}]}{\equiv_{p_\varphi}} \text{pd}(w, C).$$

Therefore, we get $\text{pd} \equiv_{p_\varphi} \text{pd}_1 + \dots + \text{pd}_t$.

It thus remains to show that for every $i \in [t]$, the tuple $P_i = (\hat{\mathbf{f}}_i, \hat{\mathbf{g}}_i, \hat{\mathbb{E}}_i, \vec{\mathcal{W}}_i)$ belongs to the set $\text{Part}_{v_i}(I[\text{pd} \mapsto \text{pd}_i], a_i, b_i, \bar{c}_i)$. So let $i \in [t]$ be arbitrary but fixed. First, we show that P_i is a partial solution for $\text{tss}(v_i) = (\text{tail}(v_i), \text{subtree}[v_i], \text{sheaf}(v_i))$. For this we will strongly rely on the fact that P is a partial solution for $\text{tss}[u]$. By definition, P_i is **(Domain-respecting)**. Let now $C = (X, Y) \in \mathfrak{K}$ be arbitrary but fixed:

- **(Edge-consistent)**: Consider an arbitrary edge $zw \in \text{sheaf}(v_i)$. Note that we then have $w \in \text{broom}[v_i]$. Then there exists $j \in [d_\varphi]$ with $(z^1, w^j) \in \vec{\mathcal{W}}_i(X, Y)$ if and only if there exists $j \in [d_\varphi]$ with $(z^1, w^j) \in \vec{\mathcal{W}}(X, Y)$ —this is because $\text{sheaf}(v_q)$ is disjoint from $\text{sheaf}(v_i)$ for every $q \neq i \in [t]$. Since P is a partial solution for $\text{tss}[u]$ this is equivalent to $z \in \hat{\mathbf{f}}(X)$ and $zw \in \hat{\mathbf{g}}(Y)$. We have $z \in \text{broom}[v_i]$ and $zw \in \text{sheaf}(v_i)$ so this is equivalent to $z \in \hat{\mathbf{f}}_i(X)$ and $zw \in \hat{\mathbf{g}}_i(Y)$.
- **(Respect-subtree- $\hat{\mathbb{E}}$)**: Consider an arbitrary vertex $z \in \text{subtree}[v_i] \subseteq \text{subtree}(u)$ and arbitrary $j \in [\hat{\mathbb{E}}_i(v, C)] = [\hat{\mathbb{E}}(v, C)]$. Partial solution P having this property for $\text{tss}[u]$ implies that $\vec{\mathcal{W}}(C)$ contains an arc (w^1, z^j) for some vertex w . By construction of the digraph $\vec{G}_\varphi$, wz is then an edge of G . Every edge of G incident with z in G belongs to $\text{sheaf}(z) \subseteq \text{sheaf}(v_i)$. Therefore, the arc (w^1, z^j) is also contained in $\vec{\mathcal{W}}(C) \cap \text{sheaf}(v_i) = \vec{\mathcal{W}}_i(C)$.
- **(Respect-broom- $\hat{\mathbb{E}}$)**: Due to $P \in \text{Part}_u[I, a, b, \bar{c}]$, for every $v \in \text{broom}[u]$, the set $\vec{\mathcal{W}}(C)$ does not contain any arc incoming into v^j for any $j > \hat{\mathbb{E}}(v, C)$. Then this property holds for the subset $\text{broom}[v_i]$ of $\text{broom}[u]$, the subset $\vec{\mathcal{W}}_i(C)$ of $\vec{\mathcal{W}}(C)$, and the restriction $\hat{\mathbb{E}}_i$ of $\mathbb{E}$ to $\text{broom}[v_i]$ for every $i \in [t]$.

So P_i is indeed a partial solution for $\text{tss}(v_i)$.

Next, we show that P_i is compatible with the index $I[\text{pd} \mapsto \text{pd}_i] = (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}_i, \mathbf{F})$ for $\text{tss}(v_i)$. We will rely on the fact that P is compatible with I for $\text{tss}[u]$. The properties **(f-consistent)** and **($\mathbb{E}$ -consistent)** follow directly from $\text{tail}[u] = \text{tail}(v_i)$. The property **(pd-consistent)** holds by definition of pd_i . The property **(Respect-alw)** follows from $\text{tail}[u] = \text{tail}(v_i)$ as well as $\vec{\mathcal{W}}_i(C) \subseteq \vec{\mathcal{W}}(C)$ for every $C \in \mathfrak{K}$. It remains to show that **(Respect-F)** is satisfied. Recall that P is compatible with I so we have: $\perp(\hat{\mathbf{f}}|_{\text{subtree}(u)}, \hat{\mathbf{g}}, \hat{\mathbb{E}}|_{\text{subtree}(u)}, (\text{pd}_{\vec{\mathcal{W}}})|_{\text{subtree}(u)}) \subseteq \mathbf{F}$. Since $\text{subtree}[v_i] \subseteq \text{subtree}(u)$ and $\text{sheaf}(v_i) \subseteq \text{sheaf}[u]$ by Lemma 5.8 we get

$$\perp((\hat{\mathbf{f}}|_{\text{subtree}(u)})|_{\text{subtree}[v_i]}, \hat{\mathbf{g}}|_{\text{sheaf}(v_i)}, (\hat{\mathbb{E}}|_{\text{subtree}(u)})|_{\text{subtree}[v_i]}, ((\text{pd}_{\vec{\mathcal{W}}})|_{\text{subtree}(u)})|_{\text{subtree}[v_i]}) \subseteq$$

$$\perp(\hat{\mathbf{f}}|_{\text{subtree}(u)}, \hat{\mathbf{g}}, \hat{\mathbb{E}}|_{\text{subtree}(u)}, (\text{pd}_{\vec{\mathcal{W}}})|_{\text{subtree}(u)}) \subseteq \mathbf{F}.$$

Since $\text{subtree}[v_i] \subseteq \text{subtree}(u)$ and $\text{subtree}[v_i] \subseteq \text{broom}[v_i]$ by definition of P_i this implies

$$\perp((\hat{\mathbf{f}}_i)|_{\text{subtree}[v_i]}, \hat{\mathbf{g}}_i, (\hat{\mathbb{E}}_i)|_{\text{subtree}[v_i]}, (\text{pd}_{\vec{\mathcal{W}}_i})|_{\text{subtree}[v_i]}) \subseteq \mathbf{F}.$$

Now we show that $(\mathbf{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}[v_i]} = (\mathbf{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}$ holds. For this consider an arbitrary vertex $w \in \text{subtree}[v_i]$ and arbitrary $C \in \mathfrak{X}$. By Observation 3.2, w has no incident edges in $\text{sheaf}(v_j)$ for any $j \neq i \in [t]$. Then we have

$$\begin{aligned} (\mathbf{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}[v_i]}(w, C) &= \mathbf{pd}_{\vec{\mathcal{W}}}(w, C) = |\{(z^1, w^j) \in \vec{\mathcal{W}}(C)\}| = |\{(z^1, w^j) \in \vec{\mathcal{W}}(C) \cap \vec{\text{sheaf}}(v_i)\}| \\ &= |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| = \mathbf{pd}_{\vec{\mathcal{W}}_i}(w, C) = (\mathbf{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}(w, C) \end{aligned}$$

So indeed $(\mathbf{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}[v_i]} = (\mathbf{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}$ holds and we get

$$\perp((\hat{\mathbf{f}}_i)_{|\text{subtree}[v_i]}, \hat{\mathbf{g}}_i, (\hat{\mathbb{E}}_i)_{|\text{subtree}[v_i]}, (\mathbf{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}) \subseteq \mathbf{F}.$$

Altogether we obtain that P_i is compatible with $I[\mathbf{pd} \mapsto \mathbf{pd}_i]$ for $\text{tss}(v_i)$. Finally, by (25), (26), and (27), the partial solution P_i is contained in $\text{Part}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i], a_i, b_i, \bar{c}_i)$. Therefore, $h_1(P)$ indeed belongs to RHS, and φ is a well-defined mapping from $\text{Part}_u[I, a, b, \bar{c}]$ to RHS. Δ

Claim 6.17. *There exists an injective mapping h_2 from RHS to $\text{Part}_u[I, a, b, \bar{c}]$.*

Proof of the claim. Let $(P_1, \dots, P_t) \in \text{RHS}$ with $P_i = (\hat{\mathbf{f}}_i, \hat{\mathbf{g}}_i, \hat{\mathbb{E}}_i, \vec{\mathcal{W}}_i)$ for every $i \in [t]$. We define $h_2(P_1, \dots, P_t) := P = (\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ such that

- For every $X \in \text{var}^V(\varphi)$, we have $\hat{\mathbf{f}}(X) = \bigcup_{i \in [t]} \hat{\mathbf{f}}_i(X)$.
- For every $Y \in \text{var}^E(\varphi)$, we have $\hat{\mathbf{g}}(Y) = \bigcup_{i \in [t]} \hat{\mathbf{g}}_i(Y)$.
- For every $C \in \mathfrak{X}$ and every $z \in \text{subtree}(u)$, let $i \in [t]$ be such that $z \in \text{subtree}[v_i]$. Then we have $\hat{\mathbb{E}}(z, C) = \hat{\mathbb{E}}_i(z, C)$. And for every $C \in \mathfrak{X}$ and every $z \in \text{tail}[u]$ we let $\hat{\mathbb{E}}(z, C) = \mathbb{E}(z, C)$.
- For every $C \in \mathfrak{X}$ we have $\vec{\mathcal{W}}(C) = \bigcup_{i \in [t]} \vec{\mathcal{W}}_i(C)$.

By definition of RHS there exist $\mathbf{pd}_i \in \text{PD}(\text{tail}[u])$, $a_i, b_i \in \mathbb{N}$, and $\bar{c}_i \in \mathbb{N}^\ell$ with $P_i \in \text{Part}_{v_i}(I[\mathbf{pd} \mapsto \mathbf{pd}_i], a_i, b_i, \bar{c}_i)$ for every $i \in [t]$ and $\sum_{i \in [t]} \mathbf{pd}_i \equiv_{p_\varphi} \mathbf{pd}$, $\sum_{i \in [t]} a_i = a$, $\sum_{i \in [t]} b_i = b$, and $\sum_{i \in [t]} \bar{c}_i = \bar{c}$. We denote the index $I[\mathbf{pd} \mapsto \mathbf{pd}_i]$ by I_i .

First we show that P is a partial solution for $\text{tss}[u]$. Recall that $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ partition $\text{sheaf}(u)$ while $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$ partition $\text{subtree}(u)$. And we also have $\text{broom}[u] = \bigcup_{i \in [t]} \text{broom}[v_i]$. We will rely on the fact that P_i is a partial solution for $\text{tss}(v_i)$ for every $i \in [t]$. First, this implies that $\hat{\mathbf{f}}$ is a vertex interpretation on $\text{broom}[u]$ and $\hat{\mathbf{g}}$ is an edge interpretation on $\text{sheaf}(u)$. Further, $\hat{\mathbb{E}}$ is a well-defined neighborhood function on $\text{broom}[u]$. We now argue that for every $C \in \mathfrak{X}$, the set $\vec{\mathcal{W}}(C)$ is a simple subset of $\text{sheaf}(u)$. For every edge $zw \in E(G)$ there exists at most one index $i \in [t]$ with $zw \in \text{sheaf}(v_i)$. Since P_i satisfies (Domain-respecting), the set $\vec{\mathcal{W}}_i(C)$ is simple. Thus, $\vec{\mathcal{W}}(C)$ is also simple and P satisfies (Domain-respecting). Now we show that the remaining three conditions are satisfied as well:

- **(Edge-consistent):** Consider an arbitrary edge $zw \in \text{sheaf}(u)$ and let $i \in [t]$ be unique such that $zw \in \text{sheaf}(v_i)$. Then we have $z \in \text{broom}[v_i]$. Further, let $(X, Y) \in \mathfrak{X}$ be arbitrary. Then there exists $j \in [d_\varphi]$ with $(z^1, w^j) \in \vec{\mathcal{W}}(X, Y)$ if and only if there exists $j \in [d_\varphi]$ with $(z^1, w^j) \in \vec{\mathcal{W}}_i(X, Y)$. Since P_i is a partial solution of $\text{tss}(v_i)$ this is equivalent to $z \in \hat{\mathbf{f}}_i(X)$ and $zw \in \hat{\mathbf{g}}_i(Y)$.

First, we show that $z \in \hat{\mathbf{f}}_i(X)$ is equivalent to $z \in \hat{\mathbf{f}}(X)$. The property $z \in \hat{\mathbf{f}}_i(X)$ trivially implies $z \in \hat{\mathbf{f}}(X)$. On the other hand, let $z \in \hat{\mathbf{f}}(X)$. Then by definition there exists an

index $j \in [t]$ with $z \in \hat{\mathbf{f}}_j(X)$. Now if $z \in \text{tail}(v_j)$, since P_j satisfies **(f-consistent)** for $\mathbf{f}$, we have $z \in \mathbf{f}(X)$, and since P_i satisfies **(f-consistent)** for $\mathbf{f}$, we, in turn, get $z \in \hat{\mathbf{f}}_i(X)$. And if $z \in \text{subtree}[v_j]$, we have $i = j$ as otherwise $\text{subtree}[v_j]$ and $\text{broom}[v_i]$ are disjoint.

Since we have $zw \in \text{sheaf}(v_i)$ and the sets $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ are pairwise disjoint, the property $zw \in \hat{\mathbf{g}}_i(Y)$ is equivalent to $\hat{\mathbf{g}}(Y)$. So **(Edge-consistent)** is satisfied.

- **(Respect-subtree- $\hat{\mathbb{E}}$)**: Consider an arbitrary vertex $z \in \text{subtree}(u)$, an arbitrary $C \in \mathfrak{K}$, and an arbitrary index $j \in [\hat{\mathbb{E}}(z, C)]$. Let $i \in [t]$ be the unique index such that $z \in \text{subtree}[v_i]$. By definition of $\hat{\mathbb{E}}$ we then have $\hat{\mathbb{E}}_i(z, C) = \hat{\mathbb{E}}(z, C)$ and therefore, $j \in [\hat{\mathbb{E}}_i(z, C)]$. Since P_i is a partial solution of $\text{tss}(v_i)$, the set $\vec{\mathcal{W}}_i(C)$ contains an arc incoming into z^j so $\vec{\mathcal{W}}(C)$ does as well.
- **(Respect-broom- $\hat{\mathbb{E}}$)** Consider a vertex $w \in \text{broom}[u]$, $C \in \mathfrak{K}$, and an index $j \in [d_\varphi]$ with $j > \hat{\mathbb{E}}(w, C)$. If $w \in \text{tail}[u]$, then for every $i \in [t]$ we have $\text{tss}[u] = \text{tail}(v_i)$ and thus $\hat{\mathbb{E}}(u, C) = \mathbb{E}(u, C) = \hat{\mathbb{E}}_i(w)$ (since P_i satisfies **($\mathbb{E}$ -consistent)** for $I[\text{pd} \mapsto \text{pd}_i]$). As P_i also satisfies **(Respect-broom- $\hat{\mathbb{E}}$)**, the set $\vec{\mathcal{W}}_i(C)$ does not contain any arc incoming into w^j . And if $w \in \text{subtree}[v_i]$, then first, we have $w \notin \text{broom}[v_q]$ for every $q \neq i \in [t]$, and thus $\vec{\mathcal{W}}_q(C)$ contains no arcs incoming into w^j ; and second, we have $\hat{\mathbb{E}}(w, C) = \hat{\mathbb{E}}_i(w, C)$ and since P_i satisfies **(Respect-broom- $\hat{\mathbb{E}}$)** with respect to $I[\text{pd} \mapsto \text{pd}_i]$, the set $\vec{\mathcal{W}}_i(C)$ also contains no such arcs. So in both cases $\vec{\mathcal{W}}(C) = \bigcup_{i \in [t]} \vec{\mathcal{W}}_i(C)$ contains no arcs incoming into w^j .

Altogether, P is a partial solution for $\text{tss}[u]$.

Next we show that P is compatible with the index $I = (\mathbf{f}, \mathbb{E}, \text{alw}, \text{pd}, \mathbf{F})$ for $\text{tss}[u]$. First, for **(f-consistent)**, for every $X \in \text{var}^V(\varphi)$ we have

$$\hat{\mathbf{f}}(X) \cap \text{tail}[u] = \left(\bigcup_{i \in [t]} \hat{\mathbf{f}}_i(X) \right) \cap \text{tail}[u] = \bigcup_{i \in [t]} (\hat{\mathbf{f}}_i(X) \cap \text{tail}[u]) = \bigcup_{i \in [t]} (\hat{\mathbf{f}}_i(X) \cap \text{tail}(v_i)) = \bigcup_{i \in [t]} \mathbf{f}_i(X)$$

where the last equality holds since each P_i satisfies **(f-consistent)** for I_i . **($\mathbb{E}$ -consistent)** is satisfied by P by construction. Next, for every $z \in \text{tail}[u]$ and every $C \in \mathfrak{K}$ we have

$$\text{pd}(z, C) \equiv_{p_\varphi} \sum_{i \in [t]} \text{pd}_i(z, C) \equiv_{p_\varphi} \sum_{i \in [t]} \text{pd}_{\vec{\mathcal{W}}_i}(z, C) = \text{pd}_{\vec{\mathcal{W}}}(z, C)$$

where the second equality since P_i satisfies **(pd-consistent)** for I_i and the least equality holds since $\vec{\mathcal{W}}_1(C), \dots, \vec{\mathcal{W}}_t(C)$ partition $\vec{\mathcal{W}}(C)$. So **(pd-consistent)** is satisfied by P . Further, for every $z \in \text{tail}[u]$, every $C \in \mathfrak{K}$, and every edge $(w^1, z^j) \in \vec{\mathcal{W}}(C)$ the following holds. Let $i \in [t]$ be such that $(w^1, z^j) \in \vec{\mathcal{W}}_i(C)$. Since P_i is compatible with I_i and satisfies **(Respect-alw)**, we have $j \leq \text{alw}(z, C)$ —so this property also applies to P .

Finally, we show that the inclusion $\perp((\hat{\mathbf{f}})_{|\text{subtree}(u)}, \hat{\mathbf{g}}, \hat{\mathbb{E}}_{|\text{subtree}(u)}, (\text{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}(u)}) \subseteq \mathbf{F}$ holds. For every $i \in [t]$, since P_i satisfies **(Respect-F)**, for the index I_i , we have

$$\perp((\hat{\mathbf{f}}_i)_{|\text{subtree}[v_i]}, \hat{\mathbf{g}}_i, (\hat{\mathbb{E}}_i)_{|\text{subtree}[v_i]}, (\text{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}) \subseteq \mathbf{F}.$$

The following holds since $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$ partition $\text{subtree}(u)$ and $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ partition $\text{sheaf}(u)$:

- $\hat{\mathbf{f}}_{|\text{subtree}(u)} = \bigcup_{i \in [t]} (\hat{\mathbf{f}}_i)_{|\text{subtree}[v_i]}$, • $\hat{\mathbf{g}} = \bigcup_{i \in [t]} \hat{\mathbf{g}}_i$, • $\hat{\mathbb{E}}_{|\text{subtree}(u)} = \bigcup_{i \in [t]} (\hat{\mathbb{E}}_i)_{|\text{subtree}[v_i]}$,
- for every $i \neq j \in [t]$ and every vertex $z \in \text{subtree}[v_i]$, the vertex z has no incident edges in $\text{sheaf}(v_j)$ so for any $C \in \mathfrak{K}$ we have $\text{pd}_{\vec{\mathcal{W}}}(z, C) = \text{pd}_{\vec{\mathcal{W}}_i}(z, C)$. Hence, $(\text{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}(u)} = \bigcup_{i \in [t]} (\text{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}$.

So we get

$$\begin{aligned}
& \perp(\hat{\mathbf{f}}_{|\text{subtree}(u)}, \hat{\mathbf{g}}, \hat{\mathbb{E}}_{|\text{subtree}(u)}, (\text{pd}_{\vec{\mathcal{W}}})_{|\text{subtree}(u)}) = \\
& \perp(\bigcup_{i \in [t]} (\hat{\mathbf{f}}_i)_{|\text{subtree}[v_i]}, \bigcup_{i \in [t]} \hat{\mathbf{g}}_i, \bigcup_{i \in [t]} (\hat{\mathbb{E}}_i)_{|\text{subtree}[v_i]}, \bigcup_{i \in [t]} (\text{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}) \stackrel{\text{Lemma 5.10}}{=} \\
& \bigcup_{i \in [t]} \perp((\hat{\mathbf{f}}_i)_{|\text{subtree}[v_i]}, \hat{\mathbf{g}}_i, (\hat{\mathbb{E}}_i)_{|\text{subtree}[v_i]}, (\text{pd}_{\vec{\mathcal{W}}_i})_{|\text{subtree}[v_i]}) \subseteq \mathbf{F}.
\end{aligned}$$

Thus P satisfies **(Respect-F)** and therefore, we have $P \in \text{Part}_u[I]$.

Next, we show that we have $P \in \text{Part}_u[I, a, b, \bar{c}]$. First, it holds that

$$\begin{aligned}
a &= \sum_{i=1}^t a_i \stackrel{P_i \in \text{Part}_{v_i}(I_i, a_i, b_i, \bar{c}_i)}{=} \sum_{i=1}^t a_{v_i}(P_i) = \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{v \in \text{subtree}[v_i]: \\ \hat{\mathbb{E}}_i(v, C) < d_\varphi}} \hat{\mathbb{E}}_i(v, C) \\
&= \sum_{C \in \mathfrak{K}} \sum_{i=1}^t \sum_{\substack{v \in \text{subtree}[v_i]: \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) = \sum_{C \in \mathfrak{K}} \sum_{\substack{v \in \text{subtree}(u): \\ \hat{\mathbb{E}}(v, C) < d_\varphi}} \hat{\mathbb{E}}(v, C) = a_u[P].
\end{aligned}$$

Similarly, we have

$$\begin{aligned}
b &= \sum_{i=1}^t b_i \stackrel{P_i \in \text{Part}_{v_i}(I_i, a_i, b_i, \bar{c}_i)}{=} \sum_{i=1}^t b_{v_i}(P_i) = \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[v_i] \\ \hat{\mathbb{E}}_i(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\
V(\text{sheaf}(v_i)) \subseteq \text{broom}[v_i] & \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u] \\ \hat{\mathbb{E}}_i(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\
& \stackrel{\substack{\text{Def. of } \hat{\mathbb{E}}, \\ P_i \text{ compatible with } I_i}}{=} \sum_{i=1}^t \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u] \\ \hat{\mathbb{E}}(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\
&= \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u] \\ \hat{\mathbb{E}}(w, C) < d_\varphi}} \sum_{i=1}^t |\{(z^1, w^j) \in \vec{\mathcal{W}}_i(C)\}| \\
&= \sum_{C \in \mathfrak{K}} \sum_{\substack{w \in \text{broom}[u] \\ \hat{\mathbb{E}}(w, C) < d_\varphi}} |\{(z^1, w^j) \in \vec{\mathcal{W}}(C)\}| = b_u[P]
\end{aligned}$$

For every $j \in [\ell]$, if $X_j \in \text{var}^V(\varphi)$, then we have

$$c^j = \sum_{i=1}^t c_i^j \stackrel{P_i \in \text{Part}_{v_i}(I_i, a_i, b_i, \bar{c}_i)}{=} \sum_{i=1}^t c_{v_i}^j(P_i) = \sum_{i=1}^t |\hat{\mathbf{f}}_i(X_j) \cap \text{subtree}[v_i]|$$

$$\hat{\mathbf{f}}_q(X_j) \cap \text{subtree}[v_i] = \emptyset \text{ for } q \neq i \in [t] \sum_{i=1}^t |\hat{\mathbf{f}}(X_j) \cap \text{subtree}[v_i]| = |\hat{\mathbf{f}}(X_j) \cap \text{subtree}(u)| = c_u^j(P),$$

and if $X_j \in \text{var}^E(\varphi)$, then we have

$$c^j = \sum_{i=1}^t c_i^j \stackrel{P_i \in \text{Part}_{v_i}(I_i, a_i, b_i, \bar{c}_i)}{=} \sum_{i=1}^t c_{v_i}^j(P_i) = \sum_{i=1}^t \hat{\mathbf{g}}_i(X_j) = \hat{\mathbf{g}}(X_j) = c_u^j(P).$$

Altogether, we obtain that $P \in \text{Part}_u[I, a, b, \bar{c}] = \text{Part}_u[I, a, b, \bar{c}]$ holds and therefore, h_2 is a well-defined mapping from RHS to $\text{Part}_u[I, a, b, \bar{c}]$.

It remains to show that h_2 is injective. For this consider $(P_1^1, \dots, P_t^1), (P_1^2, \dots, P_t^2) \in \text{RHS}$ with $h_2(P_1^1, \dots, P_t^1) = h_2(P_1^2, \dots, P_t^2)$. And for $i \in [t]$ and $j \in [2]$, let $P_i^j = (\hat{\mathbf{f}}_i^j, \hat{\mathbf{g}}_i^j, \hat{\mathbb{E}}_i^j, \vec{\mathcal{W}}_i^j)$. Also let $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}}) = h_2(P_1^1, \dots, P_t^1) = h_2(P_1^2, \dots, P_t^2)$. Now fix an arbitrary $i \in [t]$. First, for every vertex $z \in \text{tail}[u]$, every $j \in [2]$, and every $X \in \text{var}^V(\varphi)$, since P_i^j is compatible with $I[\text{pd} \mapsto \text{pd}']$ for some $\text{pd}' \in \text{PD}(\text{tail}[u])$, we have $u \in \hat{\mathbf{f}}_i^j(X)$ if and only if $u \in \hat{\mathbf{f}}(X)$. Further, and every $X \in \text{var}^V(\varphi)$, since $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$ are pairwise disjoint, we have

$$\hat{\mathbf{f}}_i^1(X) \cap \text{subtree}[v_i] = \hat{\mathbf{f}}(X) \cap \text{subtree}[v_i] = \hat{\mathbf{f}}_i^2(X) \cap \text{subtree}[v_i].$$

Hence, $\hat{\mathbf{f}}_i^1 = \hat{\mathbf{f}}_i^2$. Further, for every $Y \in \text{var}^E(\varphi)$, since $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ are pairwise disjoint, we also have

$$\hat{\mathbf{g}}_i^1(Y) = \hat{\mathbf{g}}(Y) \cap \text{sheaf}(v_i) = \hat{\mathbf{g}}_i^2(Y),$$

i.e., $\hat{\mathbf{g}}_i^1 = \hat{\mathbf{g}}_i^2$. Similarly, for every vertex $z \in \text{tail}[u]$, every $j \in [2]$, and every $C \in \mathfrak{X}$, since P_i^j is compatible with $I[\text{pd} \mapsto \text{pd}']$ for some $\text{pd}' \in \text{PD}(\text{tail}[u])$, we have $\hat{\mathbb{E}}_i^j(z, C) = \mathbb{E}(z, C)$. And for every $j \in [2]$, $C \in \mathfrak{X}$, and $z \in \text{subtree}[v_i]$ we have

$$\hat{\mathbb{E}}_i^1(z, C) = \mathbb{E}(z, C) = \hat{\mathbb{E}}_i^2(z, C),$$

hence $\hat{\mathbb{E}}_i^1 = \hat{\mathbb{E}}_i^2$. And finally, for every $C \in \mathfrak{X}$, since $\text{sheaf}(v_1), \dots, \text{sheaf}(v_t)$ are pairwise disjoint, we have

$$\vec{\mathcal{W}}_i^1(C) = \vec{\mathcal{W}}(C) \cap \text{sheaf}(v_i) = \vec{\mathcal{W}}_i^2(C),$$

i.e., $\vec{\mathcal{W}}_i^1 = \vec{\mathcal{W}}_i^2$. Altogether, we get $(P_1^1, \dots, P_t^1) = (P_1^2, \dots, P_t^2)$ and therefore, h_2 is injective. $\triangle$

The existence of injective mappings $h_1: \text{Part}_u[I, a, b, \bar{c}] \rightarrow \text{RHS}$ and $h_2: \text{RHS} \rightarrow \text{Part}_u[I, a, b, \bar{c}]$ implies that we have $|\text{Part}_u[I, a, b, \bar{c}]| = |\text{RHS}|$ concluding the proof of (22).

Now given that this property holds, we will apply the convolution result for DFTs (cf. Theorem A.4) to prove the second part of the claim, namely that (23) holds for all $I \in \text{Index}[u]$. Let thus p be an arbitrary but fixed prime such that $\mathbb{F}_p^*$ admits a p_φ -th root of unity. Before we proceed with the formal proof, let us briefly sketch the idea behind it. If we fix $(\mathbf{f}, \mathbb{E}, \text{alw}, \mathbf{F})$ and only vary pd then $\mathcal{P}_u[\cdot]$ and $\mathcal{P}_{v_i}(\cdot)$ (modulo p) become functions from $(\text{tail}[u] \times \mathfrak{X}) \rightarrow \mathbb{Z}_{p_\varphi}$ to $\mathbb{F}_p[\alpha, \beta, \gamma_1, \dots, \gamma_\ell]$, and $\mathcal{Q}_u^p[\cdot]$ and $\mathcal{Q}_{v_i}^p(\cdot)$ become their discrete Fourier transforms, respectively. Given this, the convolution result (Theorem A.4) directly implies the desired equality (23). Now we provide the formal details.

In the remainder of the proof $\equiv$ without any subscript denotes $\equiv_p$ and ω denotes $\omega(p)$. Let $I = (\mathbf{f}, \mathbb{E}, \mathbf{alw}, \mathbf{pd}, \mathbf{F}) \in \text{Index}[u]$ be arbitrary but fixed. And let $D = \text{tail}[u] \times \mathbb{A}$. Recall that we have $\text{tail}[u] = \text{tail}(v_i)$ for every $i \in [t]$. We define the mappings $h, h^1, \dots, h^t, g, g^1, \dots, g^t: (D \rightarrow \mathbb{Z}_{p_\varphi}) \rightarrow \mathbb{F}_p[\alpha, \beta, \gamma_1, \dots, \gamma_\ell]$ as:

$$h(y) \equiv \mathcal{P}_u[\mathbf{f}, \mathbb{E}, \mathbf{alw}, y, \mathbf{F}] \text{ and } g(y) \equiv \mathcal{Q}_u^p[\mathbf{f}, \mathbb{E}, \mathbf{alw}, y, \mathbf{F}]$$

and

$$h^i(y) \equiv \mathcal{P}_{v_i}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, y, \mathbf{F}) \text{ and } g^i(y) \equiv \mathcal{Q}_{v_i}^p(\mathbf{f}, \mathbb{E}, \mathbf{alw}, y, \mathbf{F})$$

for every $i \in [t]$ and every $y: D \rightarrow \mathbb{Z}_{p_\varphi}$.

First, the equality (22) implies that we have

$$h(y) \equiv \sum_{\substack{y^1, \dots, y^t \in \text{PD}(\text{tail}[u]) \\ y^1 + \dots + y^t \equiv_{p_\varphi} y}} h^1(y^1) \cdot \dots \cdot h^t(y^t) \equiv \sum_{\substack{y^1, \dots, y^t: D \rightarrow \mathbb{Z}_{p_\varphi} \\ y^1 + \dots + y^t \equiv_{p_\varphi} y}} h^1(y^1) \cdot \dots \cdot h^t(y^t). \quad (28)$$

Second, by definitions of the polynomials $\mathcal{Q}_u^p[\cdot]$ and $\mathcal{Q}_{v_i}^p(\cdot)$ (cf. (3)) it holds that

$$g(y) \equiv \sum_{q \in \text{PD}(\text{tail}[u])} \prod_{(v, C) \in \text{tail}[u] \times \mathbb{A}} \omega^{q(v, C) \cdot y(v, C)} h(q) \equiv \sum_{q: D \rightarrow \mathbb{Z}_{p_\varphi}} \prod_{z \in D} \omega^{q(z) \cdot y(z)} h(q)$$

and

$$g^i(y) \equiv \sum_{q: \text{tail}(v_i) \times \mathbb{A} \rightarrow \mathbb{Z}_{p_\varphi}} \prod_{(v, C) \in \text{tail}(v_i) \times \mathbb{A}} \omega^{q(v, C) \cdot y(v, C)} h^i(q) \equiv \sum_{q: D \rightarrow \mathbb{Z}_{p_\varphi}} \prod_{z \in D} \omega^{q(z) \cdot y(z)} h^i(q)$$

for every $y: D \rightarrow \mathbb{Z}_{p_\varphi}$ and every $i \in [t]$, i.e., we have

$$g \equiv \text{DFT}_{D, p_\varphi, p}(h) \text{ and } g^i \equiv \text{DFT}_{D, p_\varphi, p}(h^i). \quad (29)$$

Then for every $y: D \rightarrow \mathbb{F}_p$ we have:

$$h(y) \stackrel{(28)}{\equiv} \sum_{\substack{y^1, \dots, y^t: D \rightarrow \mathbb{Z}_{p_\varphi} \\ y^1 + \dots + y^t \equiv_{p_\varphi} y}} h^1(y^1) \cdot \dots \cdot h^t(y^t)$$

$$\stackrel{\text{Theorem A.4}}{\equiv} (\text{DFT}_{D, p_\varphi, p}^{-1}(\text{DFT}_{D, p_\varphi, p}(h^1) \cdot \dots \cdot \text{DFT}_{D, p_\varphi, p}(h^t)))(y) \stackrel{(29)}{\equiv} (\text{DFT}_{D, p_\varphi, p}^{-1}(g^1 \cdot \dots \cdot g^t))(y),$$

i.e.,

$$h \equiv \text{DFT}_{D, p_\varphi, p}^{-1}(g^1 \cdot \dots \cdot g^t),$$

and hence

$$g \stackrel{(29)}{\equiv} \text{DFT}_{D, p_\varphi, p}(h) \equiv \text{DFT}_{D, p_\varphi, p}(\text{DFT}_{D, p_\varphi, p}^{-1}(g^1 \cdot \dots \cdot g^t)) \stackrel{\text{Theorem A.3}}{\equiv} g^1 \cdot \dots \cdot g^t. \quad (30)$$

Finally, by definitions of $h, h^1, \dots, h^t, g, g^1, \dots, g^t$ we obtain

$$\mathcal{Q}_u[\mathbf{f}, \mathbb{E}, \mathbf{alw}, y, \mathbf{F}] \equiv g(y) \stackrel{(30)}{\equiv} (g^1 \cdot \dots \cdot g^t)(y) \equiv \prod_{i=1}^t \mathcal{Q}_{v_i}(\mathbf{f}, \mathbb{E}, \mathbf{alw}, y, \mathbf{F})$$

for every $y: D \rightarrow \mathbb{Z}_{p_\varphi}$ concluding the proof of (23). $\square$

7 Model Checking

Here, we show how given the equalities from the previous section, we can obtain the desired space-efficient algorithm checking whether the input formula $\varphi \in \text{NEO}_2[\text{FRec}]$ is satisfied by the graph G when an elimination tree T of depth td is provided with the input. Directly using those equalities, one could obtain a recursive algorithm with running time single-exponential in td and polynomial space complexity: at each of $\mathcal{O}(\text{td})$ levels of the recursion, one would keep the current partial sums and products of polynomials. Since the number of non-zero monomials is upper-bounded by $\mathcal{O}(n^{\mathcal{O}(\ell)} \cdot d_\varphi \cdot |\varphi|)$ (cf. Lemma 5.16) and the length of each coefficient is upper-bounded by $\mathcal{O}(n \cdot \text{td} \cdot |\varphi| \cdot \log(d_\varphi))$ (cf. Lemma 6.2), the polynomial space complexity follows. In this section, we show how to improve the space complexity to have only logarithmic dependency on n and eventually prove our main result.

7.1 Getting to logarithmic space complexity

To improve the space complexity to have only logarithmic dependency on n , we will rely on the Chinese remainder theorem which allows us to use the evaluation of the polynomials from the previous subsection in a “small” set of “small” prime fields so that we can reconstruct the bits of the coefficients of the desired “root” polynomial in logarithmic in n space. This idea was first used on elimination trees by Pilipczuk and Wrochna [37] to solve the DOMINATING SET problem. In order to ensure that our equalities from the previous subsection hold, we will use the small primes p such that, additionally, $\mathbb{F}_p^*$ admits the p_φ -th root of unity. The idea of combining Chinese remainder theorem with discrete Fourier transforms was used by Focke et al. [20] to obtain fast convolutions for (σ, ρ) -DOMINATING SET on graphs of bounded treewidth. We remark that their algorithm has exponential space complexity. Now we formalize this sketch. We first set up the algebraic tools and after that, put everything together to obtain the desired algorithm.

First of all, the approach of Pilipczuk and Wrochna [37] only works for univariate polynomials so we start by providing a coefficient-preserving transformation of multivariate polynomials of bounded degree into univariate polynomials. Bergougnoux et al. [5] carried out this transformation for polynomials in two variables, and we now generalize their idea to our polynomials in $\ell + 2$ variables. For this let

$$A = 2 \cdot |\varphi| \cdot d_\varphi \cdot n \cdot \text{td} + 1,$$

the choice of A becomes clear in the proof of Lemma 7.6. We also define

$$A^* = A^{\ell+2} - 1.$$

Now consider an arbitrary polynomial $\mathcal{P}$ in $\ell + 2$ formal variables $\alpha, \beta, \gamma_1, \dots, \gamma_\ell$ of form

$$\mathcal{P} = \sum_{0 \leq a, b, c_1, \dots, c_\ell < A} q_{a, b, \bar{c}} \alpha^a \beta^b \gamma^{\bar{c}} \in \mathbb{Z}[\alpha, \beta, \gamma_1, \dots, \gamma_\ell],$$

i.e., the degree of every variable is upper-bounded by A (in the above, $\bar{c} = (c_1, \dots, c_\ell)$). Then we define the univariate polynomial $\mathcal{P}^\uparrow \in \mathbb{Z}[x]$ as

$$\mathcal{P}^\uparrow = \sum_{0 \leq a, b, c_1, \dots, c_\ell \leq A} q_{a, b, \bar{c}} x^a \left(\prod_{i \in [\ell]} x^{c_i \cdot A^i} \right) x^{b \cdot A^{\ell+1}} = \sum_{0 \leq a, b, c_1, \dots, c_\ell \leq A} q_{a, b, \bar{c}} x^{a + \sum_{i \in [\ell]} c_i \cdot A^i + b \cdot A^{\ell+1}}.$$

The following properties are crucial.

Observation 7.1. *The sum in exponent of each monomial in $\mathcal{P}^\uparrow$ is the representation of some number between 0 and A^* to the base A . Furthermore, every number has a unique base A representation. Thus the following hold:*

- *The degree of $\mathcal{P}^\uparrow$ is at most A^* .*
- *For every exponent $0 \leq j \leq A^*$, the coefficient at x^j in $\mathcal{P}^\uparrow$ is precisely $q_{a,b,\bar{c}}$ for the unique representation $j = a + \left(\sum_{i \in [\ell]} c_i \cdot A^i\right) + b \cdot A^{\ell+1}$ of j to the base A —hence, the task of reconstructing the coefficients of $\mathcal{P}$ is equivalent to the one of $\mathcal{P}^\uparrow$.*
- *For any $s \in \mathbb{Z}$ and for the evaluation of $\mathcal{P}^\uparrow$ at s , it holds that*

$$\mathcal{P}^\uparrow(s) = \mathcal{P}(s, s^{A^{\ell+1}}, (s^{A^i})_{i \in [\ell]}).$$

For univariate polynomials we will rely on the following implication of the Chinese remainder theorem.

Lemma 7.2 ([12, 37]). *Let n' be a non-negative integer and let $P(x) = \sum_{i=0}^{n'} q_i \cdot x^i$ be a polynomial in one variable x of degree at most n' with integer coefficients satisfying $0 \leq q_i \leq 2^{n'}$ for $i = 0, \dots, n'$. Further, let c be a non-negative integer. And let $p_1, \dots, p_{n'}$ be prime numbers satisfying $n' + 1 < p_1 < \dots < p_{n'} \leq c \cdot (n')^2$ and $\prod_{i=1}^{n'} p_i > 2^{n'}$ and such that for every $i \in [n']$, the value p_i can be computed in time $(c \cdot n')^{\mathcal{O}(1)}$ and space $\mathcal{O}(\log n' + \log c)$. Suppose that given an integer $i \in [n']$ and a value $s \in \mathbb{F}_{p_i}$, the evaluation $P(s) \pmod{p_i}$ can be computed in time **time** and space **space**. Then the coefficients $q_1, q_2, \dots, q_{n'}$ of P can be computed in time $\mathcal{O}(\text{time}) \cdot (c \cdot n')^{\mathcal{O}(1)}$ and space $\mathcal{O}(\text{space} + \log n' + \log c)$.*

Proof. This lemma is a slightly adapted version of Theorem 5.3 by Pilipczuk and Wrochna [37]. For the sake of completeness we provide the proof of this adapted theorem. It suffices to describe how to compute a single coefficient q_k as this can then be repeated for every $k \in [n']_0$ reusing the same space. Due to $\prod_{i=1}^{n'} p_i > 2^{n'} \geq q_k$, by the Chinese remainder theorem the coefficient q_k is uniquely determined by the rests $q_k \pmod{p_1}, q_k \pmod{p_2}, \dots, q_k \pmod{p_{n'}}$. Moreover, there is an effective version of this theorem [12] stating that from these rests, the bits of q_k can be recovered in time $(c \cdot n')^{\mathcal{O}(1)}$ and space $\mathcal{O}(\log c + \log n')$. For clarification let us remark that this procedure does not store all required rests simultaneously but rather uses the compositionality of logspace algorithms and then outputs the bits of q_k one after another.

So we remain with the task of computing $q_k \pmod{p_i}$ given some $i \in [n']$. First, the prime p_i can be computed in time $(c \cdot n')^{\mathcal{O}(1)}$ and space $\mathcal{O}(\log n' + \log c)$ by assumption. The bit length of p_i is also upper-bounded by $\mathcal{O}(\log n' + \log c)$. In the same time and space we can compute a generator α_i of the group $\mathbb{F}_{p_i}^*$ by trying every number $t \in \mathbb{F}_{p_i}^*$ and computing all of its exponents t^j (with $j \in [p_i - 1]$). In their proof Pilipczuk and Wrochna [37] show that $n' < p_i - 1$ implies the equality

$$q_k \equiv_{p_i} - \sum_{j=0}^{p_i-2} P(\alpha_i^j) \cdot \alpha_i^{-jk}.$$

And this value can now be computed by evaluating P at most $(c \cdot n')^{\mathcal{O}(1)}$ times and carrying out at most $(c \cdot n')^{\mathcal{O}(1)}$ arithmetic operations in $\mathbb{F}_{p_i}$. Altogether, the rest $q_k \pmod{p_i}$ is computed in time $\mathcal{O}(\text{time}) \cdot (c \cdot n')^{\mathcal{O}(1)}$ and space $\mathcal{O}(\text{space} + \log n' + \log c)$ so the lemma follows. $\square$

By Lemma 5.17, to check if G satisfies φ , it suffices to iterate over all $\mathbf{F} \subseteq \text{EQ}(\varphi)$, all integers $j \in [2 \cdot |\varphi| \cdot d_\varphi \cdot n \cdot \text{td}]_0$, and all $\bar{c} \in ([n \cdot \text{td}]_0)^\ell$ and check whether $\llbracket \varphi \rrbracket^{(\mathbf{F}, \bar{c})} = 1$ holds and the monomial $\alpha^j \beta^j \gamma^{\bar{c}}$ has a non-zero coefficient in $\mathcal{P}_r^-(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$ (here r denotes the root of the elimination tree T). To carry out this check in a space-efficient way we will make use of Lemma 7.2. We now define the set $p_1, \dots, p_m$ of primes that will be used in our algorithm. The choices of m and $\mathbf{M}$ in the definition below will become clear in the proof of Theorem 7.7.

Definition 7.3. Let $\mathbf{M} = 2^{5 \cdot n \cdot \text{td} \cdot |\varphi| \cdot \log(d_\varphi + 1)}$ and let $m = (5 \cdot n \cdot \text{td} \cdot |\varphi| \cdot d_\varphi)^{|\ell|+2}$. Consider the infinite set of the prime numbers having the rest 1 modulo p_φ . We define $p_1 < p_2 < \dots < p_m$ as the $(m+1)$ -st, $(m+2)$ -nd, $\dots$, and the $(2m)$ -th smallest number in this set, respectively.

Lemma 7.4. The following properties hold:

1. $m \geq \mathbf{A}^*$,
2. $2^m \geq \mathbf{M}$,
3. $\prod_{i \in [m]} p_i > 2^m$,
4. for $i \in [m]$, we have $m+1 < p_i \in \mathcal{O}(p_\varphi \cdot m \cdot \log m)$,
5. for $i \in [m]$, $\mathbb{F}_{p_i}^*$ admits the p_φ -th root of unity, i.e., $\omega(p_i)$ is well-defined,
6. for $i \in [m]$, the numbers p_i , $\omega(p_i)$, and the inverse of $p_\varphi^{|\boxtimes|}$ in $\mathbb{F}_{p_i}$ can be computed in $\mathcal{O}(\log p_\varphi + \log m)$ space and $(m \cdot p_\varphi \cdot |\varphi|)^{\mathcal{O}(1)}$ time.

Proof. The choice of m , $\mathbf{M}$, and $\mathbf{A}^*$ ensures that we have $m \geq \mathbf{A}^*$ and $2^m \geq \mathbf{M}$. Further, each of $p_1, \dots, p_m$ is larger than 2 so their product is larger than 2^m . Since p_1 is larger than the m -th prime number, we have $p_1 > m+1$. By the Prime Number Theorem for Arithmetic Progressions (Theorem 1.5 in [4]) it holds that $p_m \in \mathcal{O}(p_\varphi \cdot m \cdot \log m)$. For $i \in [m]$, p_φ divides the order $p_i - 1$ of $\mathbb{F}_{p_i}^*$ by definition so $\mathbb{F}_{p_i}^*$ admits the p_φ -th root of unity. The upper bound on p_i implies that the length of p_i is in $\mathcal{O}(\log p_\varphi + \log m)$. So to compute p_i we can iterate through positive integers $p_\varphi + 1, 2 \cdot p_\varphi + 1, 3 \cdot p_\varphi + 1, \dots$ and for every candidate, check whether it is prime. Once we find the $(m+i)$ -th candidate p_i satisfying this property, we stop. Now we can iterate through candidates $1, 2, \dots$ from $\mathbb{F}_{p_i}^*$ to check whether it is a p_φ -th root of unity: for this we compute (modulo p_i) all powers of this candidate up to p_φ -th and check whether p_φ -th power is the first power equal to 1. And similarly, we can find the inverse of $p_\varphi^{|\boxtimes|}$: First, compute the value $p_\varphi^{|\boxtimes|}$. Then try all candidates $1, 2, \dots$ from $\mathbb{F}_{p_i}^*$ and multiply each with $p_\varphi^{|\boxtimes|}$ (modulo p_i) until we find a value where this product becomes equal to 1. All of those computations can be carried out in space linear in the length of p_i and time polynomial in the sum of p_i and $|\boxtimes| \leq |\varphi|$ so the claim follows. $\square$

So now to apply Lemma 7.2 with $n' = m$, it remains to describe how, given $i \in [m]$ and a number $s \in \mathbb{F}_{p_i}$, to evaluate the polynomial $(\mathcal{P}_r^-(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))'$ at s modulo p_i . We will use the simple fact that the evaluation of a sum (resp. product) of two polynomials is equal to the sum (resp. product) of their evaluations. We start by proving the following intermediate lemma about the evaluation of $\mathcal{Q}_r^{p_i}(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$.

Lemma 7.5. Let r be the root of the elimination tree T and let $i \in [m]$. Given the prime number p_i , a vector $\bar{s} \in \mathbb{F}_{p_i}^{\ell+2}$ of length $\ell+2$ with entries from $\mathbb{F}_{p_i}$, and a set $\mathbf{F} \subseteq \text{EQ}(\varphi)$, we can evaluate the polynomial $\mathcal{Q}_r^{p_i}(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$ at the vector $\bar{s}$ (modulo p_i) in time $(2 \cdot (d_\varphi + 1)^2 \cdot p_\varphi)^{\ell \cdot (\text{td} + \mathcal{O}(1))} \cdot n^{\mathcal{O}(1)}$ and space $\mathcal{O}(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$.

Proof. We recall that by Lemma 7.4 we can compute $\omega(p_i)$ as well as the inverse of $p_\varphi^{\boxtimes |\bar{s}|}$ in $\mathbb{F}_{p_i}$ in $\mathcal{O}(\log p_\varphi + \log m)$ space and $(m \cdot |\varphi|)^{\mathcal{O}(1)}$ time. To compute the value $\mathcal{Q}_r^{p_i}(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})(\bar{s}) \pmod{p_i}$, we alternately use Lemma 6.10 (7) and Lemma 6.15 (23) to recurse until we reach a polynomial of form $\mathcal{Q}_u^{p_i}[I]$ for some leaf u and some index $I \in \text{Index}[u]$. Then we use Lemma 6.7 to evaluate this polynomial at $\bar{s}$. At each level of recursion we store the information about the current recursive call as well as $\mathcal{O}(1)$ partial sums and products in $\mathbb{F}_{p_i}$. Now we bound the space and time complexity of this procedure.

Time complexity. During the recursive evaluation at some node u via Lemma 6.10 (7), the restriction of the index on $\text{tail}(u)$ remains unchanged and we only branch on the restriction of the index restricted on $\{u\}$. Also the value $\mathbf{F}$ remains fixed. Further, during a recursive evaluation using Lemma 6.15 (23), the index remains the same in all recursive calls. Hence, for every node w and every index $I \in \text{Index}[w]$ (resp. $I \in \text{Index}(w)$), the evaluation of the polynomial $\mathcal{Q}_w^{p_i}[I]$ (resp. $\mathcal{Q}_w^{p_i}(I)$) at $\bar{s}$ is computed at most once during this process. We recall that td denotes the depth of T so $|\text{tail}[w]|, |\text{tail}(w)| \leq \text{td}$ holds. Hence, for $\text{tail} \in \{\text{tail}[w], \text{tail}(w)\}$ there are at most $2^{|\varphi| \cdot \text{td}}$ vertex interpretations on tail , at most $(d_\varphi + 1)^{|\varphi| \cdot \text{td}}$ neighborhood functions on tail , and at most $p_\varphi^{|\varphi| \cdot \text{td}}$ mod-degrees on tail . So each of the sets $\text{Index}[w]$ and $\text{Index}(w)$ contains at most $(2 \cdot (d_\varphi + 1)^2 \cdot p_\varphi)^{|\varphi| \cdot \text{td}}$ indexes in which the last coordinate is equal to $\mathbf{F}$. The tree T contains exactly n nodes. So altogether there are at most $2 \cdot n \cdot (2 \cdot (d_\varphi + 1)^2 \cdot p_\varphi)^{|\varphi| \cdot \text{td}}$ recursive calls. For the recursive equality Lemma 6.10 (7) the running time additional to recursive calls is bounded by $(2 \cdot (d_\varphi + 1) \cdot p_\varphi)^{\mathcal{O}(|\varphi|)} \cdot (\log p_i)^{\mathcal{O}(1)}$: this follows directly from the sums and products involved in that equality. And for Lemma 6.15 (23) this running time is bounded by $n \cdot (\log p_i)^{\mathcal{O}(1)}$ as we simply recurse with the same index to each child and then compute the product of at most n resulting values in $\mathbb{F}_{p_i}$. By Observation 3.1, the graph G has at most $n \cdot \text{td}$ edges. So for a (non-recursive) computation via Lemma 6.7 the running time is bounded by $(n \cdot \text{td} \cdot 2^{|\varphi|} \cdot |\varphi|) \cdot (p_\varphi + (\log p_i)^{\mathcal{O}(1)})$. By Lemma 7.4, we have $\log(p_i) \in \mathcal{O}(\log p_\varphi + \log m)$ and the definition of m implies $\log p_i \in \mathcal{O}(\log p_\varphi + |\varphi| \cdot (\log n + \log |\varphi| + \log d_\varphi))$. Altogether, we obtain that the running time of this process is at most $(2 \cdot (d_\varphi + 1)^2 \cdot p_\varphi)^{|\varphi| \cdot (\text{td} + \mathcal{O}(1))} \cdot n^{\mathcal{O}(1)}$.

Space complexity. At each of $\mathcal{O}(\text{td})$ levels of recursion we store the information about the current recursive call (i.e., the index queried) as well as a constant number of partial sums and products in $\mathbb{F}_{p_i}$. Since during the recursive evaluation at some node u , the index “restricted” to $\text{tail}(u)$ remains unchanged, at each level of recursion it suffices to store only the restriction of the index to $\{u\}$ instead of the whole index. Again, recall that $\log p_i \in \mathcal{O}(\log p_\varphi + |\varphi| \cdot (\log n + \log |\varphi| + \log d_\varphi))$ so every partial sum and product requires at most that many bits. Thus, the total space complexity is upper-bounded by $\mathcal{O}(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$. $\square$

From this we can derive the following:

Lemma 7.6. *Let r be the root of the elimination tree T and let $i \in [m]$. First, the polynomial $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ is well-defined. Second, given the prime number p_i , a value $s \in \mathbb{F}_{p_i}$, and a set $\mathbf{F} \subseteq \text{EQ}(\varphi)$, we can evaluate the polynomial $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ at the value s (modulo p_i) within the following time and space:*

- time $(2 \cdot (d_\varphi + 1)^2 \cdot p_\varphi)^{|\varphi| \cdot (\text{td} + \mathcal{O}(1))} \cdot n^{\mathcal{O}(\ell)}$ and
- space $\mathcal{O}(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$.

Proof. By Observation 3.1, the graph G has at most $n \cdot \text{td}$ edges. Then by Lemma 5.16, the degree of every variable in $\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$ is smaller than A so the polynomial $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$

is indeed well-defined. Then by Observation 7.1, the evaluation of $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ at s is equal to the evaluation of $\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$ at the vector $\bar{s}$ defined as $\bar{s} := (s, s^{A^{\ell+1}}, (s^{A^i})_{i \in [\ell]}) \in \mathbb{F}_{p_i}^{\ell+2}$. So as the first step we compute the vector $\bar{s}$ in time $(|\varphi| \cdot d_\varphi \cdot n \cdot \text{td})^{\ell+O(1)} \cdot (\log p_i)^{O(|\varphi|)}$ and space $(\log p_i)^{O(|\varphi|)}$ (we remark that this process can be accelerated using exponentiation by squaring but we do not use it since this is not going to be the bottleneck our main algorithm, presented in the next theorem). Now by Lemma 6.5 the evaluation of $\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$ at $\bar{s}$ boils down to the evaluation of $\mathcal{P}_r(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}')$ at $\bar{s}$ for every $\mathbf{F}' \subseteq \mathbf{F}$ and summing up the results with the correct sign. So fix $\mathbf{F}' \subseteq \mathbf{F}$. By Lemma 6.6, it holds that $\mathcal{P}_r(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}')(\bar{s}) \equiv_{p_i} \mathcal{Q}_r^{p_i}(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}')(\bar{s})$. By Lemma 7.5 we can compute this evaluation in time $(2 \cdot (d_\varphi + 1)^2 \cdot p_\varphi)^{|\varphi| \cdot (\text{td} + O(1))} \cdot n^{O(1)}$ and space $O(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$. This process is then repeated at most $2^{|\varphi|}$ times (for every $\mathbf{F}' \subseteq \mathbf{F}$) so the claim follows. $\square$

7.2 Wrapping up the model checking algorithm

Now we are ready to prove the main result of this section.

Theorem 7.7. *There is an algorithm that given a formula $\varphi \in \text{NEO}_2[\text{FRec}]$ with ℓ size measurements, a graph G and an elimination tree T of G of depth at most td , runs in time $((d_\varphi + 1) \cdot p_\varphi)^{O(|\varphi| \cdot \text{td})} \cdot n^{O(\ell)}$ and space $O(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$ and decides whether $G \models \varphi$ holds.*

Proof. First, by Lemma 4.2 we may assume that φ is in $\text{coreNEO}_2[\text{FRec}]$: otherwise, we can transform it with only a linear blowup in the length and without increasing the values d_φ and p_φ and by increasing the value ℓ by at most 2—we remark that this blowup is the reason for $O(|\varphi| \cdot \text{td})$ instead of $|\varphi| \cdot (\text{td} + O(1))$ in the exponent of the running time. Let now r denote the root of T . By Lemmas 5.16 and 5.17, to check if G satisfies φ , it suffices to iterate over all $\mathbf{F} \subseteq \text{EQ}(\varphi)$, all integers $j \in [2 \cdot |\varphi| \cdot d_\varphi \cdot n \cdot \text{td}]_0$, and all tuples $\bar{c} = (c_1, \dots, c_\ell) \in ([n \cdot \text{td}]_0)^\ell$ and check whether $\llbracket \varphi \rrbracket^{(\mathbf{F}, \bar{c})} = 1$ holds and the coefficient in $\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$ in front of $\alpha^j \beta^j \gamma^{\bar{c}}$ is non-zero. We recall that $\ell \leq |\varphi|$ holds. First of all, note that we can iterate over such candidates and check whether $\llbracket \varphi \rrbracket^{(\mathbf{F}, \bar{c})} = 1$ holds in time

$$O(2^{|\varphi|} \cdot |\varphi| \cdot d_\varphi \cdot n \cdot \text{td} \cdot (n \cdot \text{td} + 1)^\ell) = d_\varphi \cdot n^{O(\ell)} \quad (31)$$

and space $O(|\varphi| \cdot \log n + \log d_\varphi)$ (because td is upper-bounded by n).

So now we focus on checking whether the corresponding coefficient is non-zero for fixed values of j , $\bar{c}$, and $\mathbf{F}$. By Lemma 7.6, the polynomial $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ is well-defined and its degree is upper-bounded by A^* . By Observation 7.1, checking whether the desired coefficient is non-zero is equivalent to determining whether the coefficient at $x^{j + \left(\sum_{i \in [\ell]} c_i \cdot (A+1)^i\right) + j \cdot (A+1)^{\ell+1}}$ in $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ is non-zero. Note that to decide this, it suffices to compute the bits of this coefficient one by one and check if at least one is non-zero.

By Lemma 6.2, every coefficient of $\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F})$, and equivalently, of $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))'$, is non-negative and upper-bounded by M . Furthermore, the degree of $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ is less than A^* by construction. We now apply Lemma 7.2 to the polynomial $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$, the value $n' = m$, $c = p_\varphi$, and the primes $p_1, \dots, p_m$. By Lemma 7.4 together with observations we just made they satisfy the preconditions of the lemma. By Lemma 7.6, for any $i \in [m]$ and any $s \in \mathbb{F}_{p_i}$ we can compute the evaluation $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))'(s) \pmod{p_i}$ in time $((d_\varphi + 1) \cdot p_\varphi)^{O(|\varphi| \cdot \text{td})} \cdot n^{O(\ell)}$ and space $\text{space} = O(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$. Lemma 7.2 then implies that we can reconstruct

the bits of the coefficient of $(\mathcal{P}_r^=(\emptyset, \emptyset, \emptyset, \emptyset, \mathbf{F}))^\uparrow$ at the monomial $x^{j + \left(\sum_{i \in [\ell]} c_i \cdot (A+1)^i\right) + j \cdot (A+1)^{\ell+1}}$ by increasing the running time by a factor of $(p_\varphi \cdot m)^{O(1)}$ and using the additional space of $O(\text{space} +$

$\log n' + \log c$). Repeating this process for every candidate $j, \bar{c}, F$ increases the running time by a factor of (31) and remains the same space complexity so the claim follows. $\square$

8 Extensions of the Logic

8.1 Clique predicate

Here we sketch that it is not difficult to extend our algorithm to additionally handle a predicate of form $\text{clique}(X)$ which reflects for a vertex set variable X that the subgraph induced by X is a clique. The main tool for this is the *fast subset cover* by Björklund et al. [7] which we shortly present. Let U be a finite set and R be a ring. Let $g_1, \dots, g_t: 2^U \rightarrow R$ be set functions, for some integer t . For every $i \in [t]$, the *zeta-transform* $\zeta g_i: 2^U \rightarrow R$ of g_i is defined by

$$(\zeta g_i)(Y) = \sum_{X \subseteq Y} g_i(X),$$

and similarly, the *Möbius-transform* $\mu g_i: 2^U \rightarrow R$ of g_i is given by

$$(\mu g_i)(Y) = \sum_{X \subseteq Y} (-1)^{|Y \setminus X|} g_i(X).$$

The *cover product* $g_1 *_c g_2 *_c \dots *_c g_t: 2^U \rightarrow R$ of $g_1, \dots, g_t$ is defined as

$$(g_1 *_c g_2 *_c \dots *_c g_t)(Y) = \sum_{\substack{X_1, \dots, X_t \subseteq U: \\ X_1 \cup \dots \cup X_t = Y}} g_1(X_1) \cdot g_2(X_2) \cdot \dots \cdot g_t(X_t).$$

The following result of Björklund et al. [7] is crucial to handle the clique-predicates:

Lemma 8.1 ([7]). *Let U be a finite set, let R be a ring, and let $g_1, \dots, g_t: 2^U \rightarrow R$ be set functions for a positive integer t . Then for every $X \in 2^U$, it holds that*

$$(g_1 *_c g_2 *_c \dots *_c g_t)(X) = \mu((\zeta g_1)(X) \cdot (\zeta g_2)(X) \cdot \dots \cdot (\zeta g_t)(X)).$$

Informally speaking, this result allows to replace the time-consuming combination of t subsets of X via a pointwise product.

Now we sketch how to use this technique to extend our approach from the previous section to the clique predicate. The crucial observation is that for every clique in G , there is a root-to-leaf path in the elimination tree T on which every vertex of this clique lies. Thus, a clique predicate can be *falsified* in two ways. Either two vertices of the non-clique come from different subtrees of some node, or they are in the ancestor-descendant relation but there is no edge between them. Now consider a formula $\varphi \in \text{coreNEO}_2 + K$ and let $\text{CLS}(\varphi) \subseteq \text{var}^V(\varphi)$ (stands for “clique set”) denote the set of vertex set variables X such that $\text{clique}(X)$ occurs in the formula. We first extend the definition of the index of a triad tss to a set $S \subseteq \text{CLS}(\varphi)$. The set S then represents the vertex set variables X such that there is a vertex $v \in \hat{\mathbf{f}}(X) \cap \text{subtree}$ and a vertex $w \in \hat{\mathbf{f}}(X)$ with $v \neq w$ such that v and w are not adjacent. So in simple words, the vertex v “falsifies” the clique predicate $\text{clique}(X)$. Furthermore, to also capture the possibility that two vertices land in different subtrees and thus, falsify some clique, we further extend the index by a set $Z \subseteq \text{var}^V(\varphi)$ which reflects all vertex variables X satisfying $\hat{\mathbf{f}}(X) \cap \text{subtree} \neq \emptyset$.

With this idea we can proceed as follows algorithmically. At the root, we additionally iterate over all possible $S \subseteq \text{CLS}(\varphi)$ to fix the sets of clique-predicates we want to falsify: with this additional

information, we can check whether with this “guess” the formula evaluates to 1, otherwise we discard the candidate. Also at the root we iterate over the subsets $Z \subseteq \text{var}^V(\varphi)$ to fix which vertex sets are non-empty. At a leaf-node u (cf. Lemma 6.7) the set $\text{subtree}[u]$ is empty so we additionally check whether S or Z in the index is non-empty: in this case, we simply return 0, otherwise we proceed as before. At a forget-node (cf. Lemma 6.10), when we branch over all restrictions of the index to $\{u\}$ we proceed as follows. First, check which cliques, say S_u , are falsified by u in this branch, i.e., for which vertex set variables X there exists a vertex v such that have $u \in \mathbf{f}_u(X)$, $v \in \mathbf{f}(X)$, and $uv \notin E(G)$, and then (1) discard the branch if some clique outside S is falsified (i.e., if $S_u \not\subseteq S$ holds), otherwise (2) allow any subset of $Q \subseteq S_u$ to also be falsified in $\text{subtree}(u)$, and (3) enforce the remaining cliques in $S \setminus S_u$ to be falsified in the subtree (i.e., use $Q \cup (S \setminus S_u)$ in the recursive call). Second, for every vertex set variable X with $u \in \mathbf{f}(X)$ we can either keep X in the set Z for the recursive call (to reflect that the vertex set X is non-empty also in $\text{subtree}(u)$) or discard it. This increases the branching by a factor of $2^{\mathcal{O}(|\varphi|)}$.

At a join-node (cf. Lemma 6.15) u with children $v_1, \dots, v_t$ we proceed as follows. First, we iterate over all subsets $S_2 \subseteq S$: this set then reflects which cliques are falsified due to having two vertices from sets in $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$. Second, we iterate over all sets S_1 satisfying $S \setminus S_2 \subseteq S_1 \subseteq S$. The set S_1 reflects which of the cliques are falsified in the same subtree: it has to contain all cliques in $S \setminus S_2$ but is allowed to further falsify the elements of S_2 .

For a variable $X \in S_2$, to satisfy the constraint “*the vertices of X occur in at least two of $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$* ”, we can take the partial solutions where *the vertices of X occur in at least one of $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$* and subtract the ones where *the vertices of X occur in exactly one of $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$* . To compute the number of partial solutions where *the vertices of X occur in exactly one of $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$* we rely on the fast subset cover and we additionally equip the polynomials of the children of u with a new formal variable δ . The exponent of δ in the polynomial of the child v_i reflects whether or not a vertex of X occurs in $\text{subtree}[v_i]$. Thus, when we multiply the polynomials of the children of u , we can safely discard the partial products where the exponent of δ is larger than one as this corresponds to the case that vertices of X occur in at least two children. On the other hand, fast subset cover ensures that the vertices of X appear in *at least one of $\text{subtree}[v_1], \dots, \text{subtree}[v_t]$* . Crucially, the same variable δ (whose exponent is bounded by $\mathcal{O}(|\varphi|)$) can be used to correctly keep track of all variables in S_2 simultaneously. The usage of the variable δ increases the space complexity only by a factor of $\mathcal{O}(|\varphi|)$. For technical details, we refer to the work of Bergougnoux et al. [5] where they employ similar ideas relying on fast subset cover with one new formal variable to solve INDEPENDENT SET parameterized by shrubdepth .

As for the set $Z \setminus S_2$, we apply fast subset cover to efficiently combine all possibilities for $Z = Z_1 \cup \dots \cup Z_t$ where Z_i reflects the vertex variables that are non-empty in $\text{subtree}[v_i]$. Altogether, due to fast subset cover we get a total branching factor of $2^{\mathcal{O}(|\varphi|)}$ for a join-node (instead of using the same index for each child (cf. Lemma 6.15)). With this we obtain our logspace algorithm for $\text{NEO}_2[\text{FRec}] + \text{K}$:

Theorem 1.3. *There is an algorithm that given a formula $\varphi \in \text{NEO}_2[\text{FRec}] + \text{K}$ with ℓ size measurements, a graph G and an elimination tree T of G of depth at most td , runs in time $((d_\varphi + 1) \cdot p_\varphi)^{\mathcal{O}(|\varphi| \cdot \text{td})} \cdot n^{\mathcal{O}(\ell)}$ and space $\mathcal{O}(\text{td} \cdot |\varphi| \cdot (\log d_\varphi + \log p_\varphi + \log n + \log |\varphi|))$ and decides whether $G \models \varphi$ holds.*

8.2 Connectivity and acyclicity

Here we sketch how we can extend our results obtained so far to acyclicity and connectivity constraints. For this, let φ be a formula from $\text{NEO}_2 + \text{ACK}$. By Lemma 4.2 we may assume that φ is in

coreNEO₂+ACK. And let G be a graph and T an elimination forest of G of depth at most td . We carry out some transformations of φ following the ideas of Bergounoux et al. [6]. First, we bring the quantifier-free part of φ in disjunctive normal form to obtain the equivalent formula ψ . Now ψ is of form $\exists \bar{X} \psi_1 \vee \psi_2 \vee \dots \vee \psi_t$ where $\exists \bar{X}$ is a sequence of existential quantifiers and each ψ_i is a conjunction of (possibly negated) set equalities, size measurements, connectivity and acyclicity constraints, and clique constraints. And the task of model checking φ in a graph G is equivalent to checking if there exists an index $i \in [t]$ with $G \models \exists \bar{X} \psi_i$. Observe that we have $t \leq 2^{\mathcal{O}(|\varphi|)}$. Also for every $i \in [t]$, we have $|\psi_i| \in \mathcal{O}(|\varphi|)$, $p_{\psi_i} \leq p_\varphi$, and $d_{\psi_i} \leq d_\varphi$. This is because every set equality, every size measurement, and every and connectivity or acyclicity constraint occurring in φ occurs at most once in ψ_i , possibly (un)negated, and on the other hand, every such element occurring in ψ_i has to occur in φ as well. So we fix a value $i \in [t]$ and focus on checking $G \models \exists \bar{X} \psi_i$.

Connectivity A subgraph induced by a vertex set X is disconnected if and only if we can partition X into two non-empty sets with no edges between them. So for a vertex set variable X , the constraint $\neg \text{con}(X)$ is equivalent to

$$\begin{aligned} \exists A \exists B \exists C \exists Z_1 \exists Z_2 \neg(Z_1 = \emptyset) \wedge \neg(Z_2 = \emptyset) \wedge Z_1 \cup Z_2 = X \wedge Z_1 \cap Z_2 = A \wedge \emptyset = A \wedge \mathbf{E} = B \wedge \\ N^{\mathbb{N}^+}(Z_1, B) = C \wedge C \cap Z_2 = \emptyset. \end{aligned}$$

We remark that even though this equivalent formulation might seem unnecessarily complicated, it ensures that after the replacements carried out next, the formula will still be in coreNEO₂+ACK. A similar remark will apply to edge sets. Similarly, an edge set Y induces a disconnected graph if and only if the set, say Z , of endpoints of Y can be partitioned into two non-empty sets Z_1 and Z_2 such that no edges in Y run between Z_1 and Z_2 . So for an edge set variable Y , the constraint $\neg \text{con}(Y)$ is equivalent to

$$\begin{aligned} \exists A \exists B \exists C \exists Z \exists Z_1 \exists Z_2 \neg(Z_1 = \emptyset) \wedge \neg(Z_2 = \emptyset) \wedge Z_1 \cup Z_2 = Z \wedge Z_1 \cap Z_2 = A \wedge \emptyset = A \wedge \mathbf{V} = B \wedge \\ N^{\mathbb{N}^+}(B, Y) = Z \wedge N^{\mathbb{N}^+}(Z_1, Y) = C \wedge C \cap Z_2 = \emptyset. \end{aligned}$$

As a next step we therefore replace as sketched above all occurrences of $\neg \text{con}(X)$ in ψ_i where X is a vertex or an edge set variable. During each such replacement we introduce new vertex and edge set variables. After that, we push all existential quantifiers to the front to obtain a formula equivalent to ψ_i of form $\exists \bar{X} \xi'$ where ξ' is a quantifier-free formula in disjunctive normal form and in which all constraints of form $\text{con}(X)$ are not negated. For simplicity, we still denote this equivalent formula by ψ_i .

So to check if $G \models \psi_i$, we can discard the connectivity constraints and search for a “solution” where each of the required sets is connected. Note that the formula ψ_i is still in coreNEO₂+ACK. Now we sketch how to extend the framework presented in the paper so far with the Cut&Count technique by Cygan et al. [17] to ensure, with high probability, that for every constraint $\text{con}(X)$ in ψ_i , the vertex set $\text{con}(X)$ is connected. For now assume that the original formula φ did not contain acyclicity constraints, we consider them later.

Before describing the application of Cut&Count [17] to our setting, let us sketch the main idea behind it. Suppose we are looking for a vertex set X with a certain property α such that, additionally, the subgraph of G induced by X is connected. A partition (X^L, X^R) of a vertex set X is called a *consistent cut* if there is no edge in G with one endpoint in X^L and the other in X^R . In other words, every connected component of $G[X]$ is contained in one of X^L and X^R . Hence, there are 2^k consistent cuts of X where k denotes the number of connected components of $G[X]$. And this number is divisible by 4 if and only if X induces a disconnected subgraph of G . Thus, if we

may, for some reason, assume that there exists at most one vertex set X with the property α , then the task becomes equivalent to counting, modulo 4, the pairs (X^L, X^R) of disjoint vertex sets such that $X^L \cup X^R$ satisfies α and there are no edges between X^L and X^R . In other words, we replaced the time-consuming task of ensuring the connectivity to counting certain pairs of vertex sets with “simpler” constraints.

This idea can naturally be generalized to multiple connectivity constraints: as long as we may assume that there exists a unique “candidate solution” and the task involves connectivity constraints on vertex sets $Y_1, \dots, Y_r$ of the candidate solution, we can eliminate all connectivity constraints by counting the combinations of consistent cuts of each of $Y_1, \dots, Y_r$. Then deciding the existence of a solution where all of these r sets are connected is equivalent to checking whether the resulting number is not divisible by 2^{r+1} (i.e., whether for none of these sets the number of consistent cuts is divisible by four).

Similarly, for an edge set Y , the endpoints of Y are precisely the vertices in $V_Y := N^{\mathbb{N}^+}(V(G), Y)$, and this can be formalized in the logic. So to check connectivity of the edge set Y we can similarly count the consistent cuts of V_Y in the graph (V_Y, Y) . For this reason, in the remainder we assume that for every constraint $\text{con}(X)$ in the formula, X is a vertex set variable, and the edge set variables can be handled similarly.

The standard tool used to guarantee that, with high probability, the solution is unique (if it exists) is the classic Isolation lemma:

Theorem 8.2 (Isolation lemma, [30]). *For an integer n , let $\mathcal{F} \subseteq 2^{[n]}$ be a non-empty set family over the universe $[n]$. For each $i \in [n]$, choose a weight $w(i) \in [2n]$ uniformly and independently at random. Then with probability at least $1/2$ there exists a unique set of minimum weight in $\mathcal{F}$.*

Now we are ready to explain how to incorporate the required changes in our algorithm from Section 7 to handle the connectivity constraints. Let

$$L = 2 \cdot |\psi_i| \cdot n \cdot (1 + \text{td} + (d_{\psi_i} + 1) + \text{td} \cdot 2 \cdot d_{\psi_i}),$$

this value will be used as the largest weight. Recall that any partial solution in Section 7 has form $(\hat{\mathbf{f}}, \hat{\mathbf{g}}, \hat{\mathbb{E}}, \vec{\mathcal{W}})$ where the elements of this tuple satisfy the conditions from Definition 5.11. Thus we can view every partial solution as a subset of a universe of size $2^{|\psi_i| \cdot n \cdot (1 + \text{td} + (d_{\psi_i} + 1) + \text{td} \cdot 2 \cdot d_{\psi_i})}$ which consists of the following elements:

- at most $|\psi_i| \cdot n$ pairs of form (X, v) where X is a vertex set variable and v is a vertex of G ,
- at most $|\psi_i| \cdot n \cdot \text{td}$ pairs of form (Y, e) where Y is an edge set variable and e is an edge of G ,
- at most $|\psi_i| \cdot n \cdot (d_{\psi_i} + 1)$ triples of form (C, v, i) with $C \in \mathbb{X}$, v being a vertex of G , and $i \in [d_{\psi_i}]_0$,
- and at most $|\psi_i| \cdot n \cdot \text{td} \cdot 2 \cdot d_{\psi_i}$ pairs of form $(C, \vec{e})$ where $C \in \mathbb{X}$ and $\vec{e} \in \vec{A}(\vec{G}_\varphi)$.

So first, in our algorithm, we sample the following weight functions:

1. For every vertex set variable X , the weight function $w_X: V(G) \rightarrow [L]$.
2. For every edge set variable Y , the weight function $w_Y: E(G) \rightarrow [L]$.
3. For every $C \in \mathbb{X}$, the weight function $w_C^{\hat{\mathbb{E}}}: V(G) \times [d_{\psi_i}]_0 \rightarrow [L]$.
4. And for every element $C \in \mathbb{X}$, the weight function $w_C^{\vec{\mathcal{W}}}: \vec{A}(\vec{G}_\varphi) \rightarrow [L]$.

For every weight function and every element in its domain, we sample the weight of this element independently at random from the range of the values of the weight function. We naturally define a weight of a partial solution as the sum of the weights of all “building blocks” of a partial solution, we restrict the weighting of vertices to **subtree** only. In the next argument, we refer to the proof of Theorem 7.7. The choice of the weights together with Theorem 8.2 now imply the following. If the union of the sets $\text{Part}_r^-(\emptyset^4, \mathbf{F}, j, \bar{c}, j)$ over all $j \in \mathbb{N}$, $\bar{c} \in \mathbb{N}^\ell$, $\mathbf{F} \subseteq \text{EQ}(\mathbf{F})$ with $\llbracket \psi_i \rrbracket^{(\mathbf{F}, \bar{c})}$ restricted to partial solutions respecting all connectivity constraints is non-empty, then it contains an element of a unique weight with probability at least $1/2$. So now instead of counting partial solutions, we can count combinations of partial solutions and connected cuts of variables X such that ψ_i contains the constraint $\text{con}(X)$. Since connected cuts of a vertex set X can be easily modeled as

$$\exists A \exists B \exists X^L \exists X^R \quad X^L \cup X^R = X \wedge X^L \cap X^R = \emptyset \wedge \mathbf{E} = A \wedge N^{\mathbb{N}^+}(X^L, A) = B \wedge B \cap X^R = \emptyset, \quad (32)$$

we now replace every constraint of form $\text{con}(X)$ by such a formula, and then push all existential quantifiers to the front. For simplicity, we still denote the arising formula by ψ_i . Now we run our algorithm from the previous section with the following changes.

1. First, randomly generate the weight functions as described above and store them. Let us emphasize that we *do not* weight the variables representing the sides of the connected cut.
2. Then, extend all polynomials $\mathcal{P}$ and $\mathcal{Q}$ by a new formal variable, say τ , whose exponent keeps track of the weight of the partial solutions. Similarly to the definition of $\bar{c}$, we restrict the weight of vertex set variables to **subtree** only.
3. At the leaf computation, when the edge e is assigned to an edge set variable Y , we additionally multiply the polynomial by $\tau^{w_Y(e)}$. We proceed similarly for the arcs added to $\vec{\mathcal{W}}(C)$ for some $C \in \mathfrak{K}$.
4. At a forget computation of a vertex u , similarly to the multiplication by $\gamma^{\bar{c}}$ reflecting that u now belongs to the subtree, we multiply by τ^w for the total weight w of u in all vertex set variables it is assigned to together with the weights of its values in $\mathbb{E}_u$.
5. At the root, additionally to iterating through the values $j, \bar{c}, \mathbf{F}$ (cf. the proof of Theorem 7.7), we iterate over all possible weights. Note that the largest possible weight is at most

$$|\psi_i| \cdot n \cdot (1 + \text{td} + 1 + 2 \cdot \text{td}) \cdot L \in \mathcal{O}(n + |\psi_i| + d_{\psi_i})^{\mathcal{O}(1)}.$$

For each of those combinations, we reconstruct the corresponding coefficient of the polynomial and check whether it is divisible by 2^{k+1} where k denotes the number of connectivity constraints in ψ_i . If we find at least one combination where the coefficient is not divisible by 2^{k+1} , then we output that the formula is satisfiable.

This algorithm does not give false positives and gives false negatives with probability at most $1/2$. Finally, we analyse how these adaptations influence the running time and the space complexity.

1. We reduce the model checking of φ to the problem of model checking at most $2^{|\varphi|}$ formulas ψ_i such that the length of ψ_i is upper-bounded by $\mathcal{O}(|\varphi|)$, d_{ψ_i} is upper-bounded by $d_\varphi + 1$, and p_{ψ_i} is upper-bounded by p_φ .
2. To generate and store the weight functions we need space polynomial in n , d_φ , and $|\varphi|$.

3. The largest exponent of any formal variable of the root polynomial $\mathcal{P}$ is upper-bounded by a polynomial in n , d_φ , and $|\varphi|$ so the bit length of each prime number we use (cf. Section 7) remains logarithmic in n , d_φ , and $|\varphi|$.

Altogether we obtain the following intermediate result:

Theorem 8.3. *There is a Monte-Carlo algorithm that given a formula $\varphi \in \text{NEO}_2[\text{FRec}] + \text{ACK}$ without acyclicity constraints, a graph G and an elimination tree T of G of depth at most td , runs in time $((d_\varphi + 1) \cdot p_\varphi)^{\mathcal{O}(|\varphi| \cdot \text{td})} \cdot n^{\mathcal{O}(\ell)}$ and space polynomial in $n + |\varphi| + d_\varphi + p_\varphi$ and decides whether $G \models \varphi$ holds. The algorithm cannot give false positives and may give false negatives with probability at most $1/2$.*

Acyclicity Now we briefly sketch how to approach the acyclicity constraints following a similar but slightly more involved idea of Cut&Count [17]. Observe that for a vertex set X the subgraph $G[X]$ is not acyclic if and only if there exists a non-empty subset Y of X such that $G[Y]$ induces a subgraph of min-degree at least 2. For this reason, we replace every constraint of form $\neg \text{acy}(X)$ with a vertex set variable X in ψ_i by

$$\exists A \exists B \exists Y \ Y \cap X = Y \wedge \neg(Y = \emptyset) \wedge \mathbf{E} = A \wedge N^{\mathbb{N} \setminus \{0,1\}}(Y, A) = B \wedge Y \cap B = Y,$$

and the push all quantifiers to the front. Similarly, if X is an edge set variable, we replace the constraint $\neg \text{acy}(X)$ by

$$\exists A \exists B \exists Y \ \mathbf{v} = A \wedge Y \cap X = Y \wedge \neg(Y = \emptyset) \wedge N^{\mathbb{N}^+}(A, Y) = B \wedge N^{\mathbb{N} \setminus \{0,1\}}(B, Y) = B$$

and the push all quantifiers to the front. This is valid because an edge set X is not acyclic if it contains a non-empty subset Y which induces a subgraph of degree at least two (A is simply the whole vertex set and B is the set of endpoints of Y). For simplicity we still denote the arising formula by ψ_i and observe that it is in $\text{coreNEO}_2 + \text{ACK}$. We remark that this replacement ensures $d_{\psi_i} \leq d_\varphi + 1$. Now all acyclicity and connectivity constraints are non-negated and ψ_i is in disjunctive normal form so we may discard those constraints from ψ_i and instead, search for a solution which additionally satisfies them.

For this we follow the idea by Bodlaender et al. [8]. We artificially insert a universal (i.e., adjacent to all vertices in $V(G)$) vertex v^* into the graph G and obtain the graph G^* . Let E^* denote the set of edges incident with v^* . We can easily obtain an elimination forest of G^* of depth at most $\text{td} + 1$ by making v^* the new root and making the root of T to the unique child of v^* . For a vertex set X by $E[X]$ we denote the set of edges induced by X and by E_X^* we denote the edges from E^* with an endpoint in X . Then the following lemma is easy to verify:

Lemma 8.4 ([8] Subsection 3.4). *A vertex set $X \subseteq V(G)$ is acyclic in G if and only if there exists a set E_X satisfying $E[X] \subseteq E_X \subseteq E[X] \cup E_X^*$ such that the graph $(X \cup \{v^*\}, E_X)$ is a tree.*

The graph $(X \cup \{v^*\}, E_X)$ is a tree if and only if (1) it is connected and (2) the number of edges in E_X is precisely $|X|$ (this is because a tree on $|X| + 1$ vertices has $|X|$ edges). Crucially, observe that if this graph is connected, then the number of edges is at least $|X|$. So to ensure acyclicity it suffices to ensure connectivity and that the number of edges is no larger than $|X|$. It is not difficult to formulate an analogous criterion for acyclicity of an edge set so we focus on vertex set variables in the remainder.

We sketch what changes need to be made to incorporate this idea in our model checking algorithm. First, for every constraint $\text{acy}(X)$ occurring in ψ_i , we introduce an existentially quantified

variable E_X . Furthermore, we replace the constraint $\text{acy}(X)$ similarly to (32) to model the connected cuts of $X \cup \{v^*\}$ with the only difference that we now have $X^L \cup X^R = X \cup \{v^*\}$ —note that we can easily transform this condition into $\text{coreNEO}_2[\text{FRec}]$.

The vertex v^* and the edges in E^* are “transparent” to the algorithm except for the treatment of the acyclicity constraints: by this we mean that the algorithm does not assign the vertex v^* to any vertex sets other than the sides of the connected cuts, and similarly the edges in E^* are only allowed to be assigned to some of the edge set variables of form E_X . Additionally, we introduce two new formal variables ρ_1 and ρ_2 . The exponent of ρ_1 keeps track of the sum of cardinalities of the sets E_X and ρ_2 keeps track of the sum of the cardinalities of the sets X over all vertex set variables X for which we enforce acyclicity. At the leaf computation, when considering an edge, say vw , such that both v and w belong to $\mathbf{f}(X)$, the edge vw is necessarily placed in $\hat{\mathbf{g}}(E_X)$. On the other hand, we only allow such edges as well as the edges from E^* with an endpoint in $\hat{\mathbf{f}}(X)$ to be placed in E_X . And we multiply by ρ_1 whenever such an edge is added to some E_X . Then at the root we only consider the monomials in which the exponents of ρ_1 and ρ_2 are equal. Except for these changes, we proceed as described earlier for connectivity: in particular, we weight every variable E_X to ensure the uniqueness of a solution with high probability. As before we do not weight the sides of connected cuts. In the end, we check if there is a monomial such that the corresponding coefficient is not divisible by 2^{r+1} where r is the total number of connectivity constraints: those are both the original connectivity constraints of ψ_i together with connectivity constraints we introduced for acyclicity constraints. Note that the exponent of the variables ρ_1 and ρ_2 can be of order $\Theta(q \cdot n)$ where q denotes the number of acyclicity constraints in φ . This is the reason behind the increase by a factor of $n^{\mathcal{O}(q)}$ in the running time compared to Theorem 8.3. With this we obtain the following theorem:

Theorem 1.4. *There is a Monte-Carlo algorithm that given a formula $\varphi \in \text{NEO}_2[\text{FRec}]+\text{ACK}$ with ℓ size measurements and q acyclicity constraints, a graph G and an elimination tree T of G of depth at most td , runs in time $((d_\varphi+1) \cdot p_\varphi)^{\mathcal{O}(|\varphi| \cdot \text{td})} \cdot n^{\mathcal{O}(q+\ell)}$ and space polynomial in $n+|\varphi|+d_\varphi+p_\varphi$ and decides whether $G \models \varphi$ holds. The algorithm cannot give false positives and may give false negatives with probability at most $1/2$.*

9 Conclusions

We would like to conclude with some open questions. The exponent of n in the running time of our model checking algorithm depends on the numbers ℓ and q of size measurements and acyclicity constraints, respectively, in the formula. So the first natural question is whether this dependence can be avoided. For example, in the runtime of the model checking for DN in [6], the exponent of n depends only in the number of size measurements $|X| \prec m$ where m is not a constant.

The next open question is whether we can derandomize and/or improve the space complexity of our model checking algorithm for $\text{NEO}_2[\text{FRec}]+\text{ACK}$ (i.e., when connectivity and acyclicity constraints are allowed) to become logarithmic in n .

Also, it can be observed that if $d_\varphi = 1$, then our algorithm actually counts the models of φ . So we are able to count, for example, dominating sets or independent sets: in fact, for these problems, for each integer t , we can count the number of solutions of size t . However, in general, the same model of the formula yields multiple partial solutions at the root (due to the component $\vec{\mathcal{W}}$ of the partial solution). So is it possible to extend our algorithm for model counting?

Further, our logic cannot capture, for example, TRIANGLE PACKING. However, it is not difficult to use the techniques presented in this paper to solve this problem, and, more generally, handle the constraint of form “the edge set Y induces a disjoint union of cliques”. So it is also reasonable to

ask if there is a natural generalization of our logic (ideally not using too many ad-hoc operators) which places all problems under the same umbrella.

References

- [1] Eric Allender, Shiteng Chen, Tiancheng Lou, Periklis A. Papakonstantinou, and Bangsheng Tang. Width-Parametrized SAT: Time–Space Tradeoffs. *Theory Comput.*, 10:297–339, 2014.
- [2] Stefan Arnborg, Jens Lagergren, and Detlef Seese. Easy Problems for Tree-Decomposable Graphs. *Journal of Algorithms*, 12(2):308–340, 1991.
- [3] Rémy Belmonte and Ignasi Sau. On the Complexity of Finding Large Odd Induced Subgraphs and Odd Colorings. *Algorithmica*, 83(8):2351–2373, 2021.
- [4] Michael A. Bennett, Greg Martin, Kevin O’Bryant, and Andrew Rechnitzer. Explicit bounds for primes in arithmetic progressions. *Illinois Journal of Mathematics*, 62(1-4):427–532, 2018.
- [5] Benjamin Bergougnoux, Vera Chekan, Robert Ganian, Mamadou Moustapha Kanté, Matthias Mnich, Sang-il Oum, Michał Pilipczuk, and Erik Jan van Leeuwen. Space-Efficient Parameterized Algorithms on Graphs of Low Shrubdepth. *ACM Trans. Comput. Theory*, 17(3), June 2025.
- [6] Benjamin Bergougnoux, Jan Dreier, and Lars Jaffke. A logic-based algorithmic meta-theorem for mim-width. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 3282–3304. SIAM, 2023.
- [7] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto. Fourier meets Möbius: fast subset convolution. In *Proc. STOC 2007*, pages 67–74, 2007.
- [8] Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Information and Computation*, 243:86–111, 2015.
- [9] Hans L. Bodlaender, John R. Gilbert, Hjálmtýr Hafsteinsson, and Ton Kloks. Approximating Treewidth, Pathwidth, Frontsize, and Shortest Elimination Tree. *J. Algorithms*, 18(2):238–255, 1995.
- [10] Édouard Bonnet, Daniel Neuen, and Marek Sokolowski. Treedepth Inapproximability and Exponential ETH Lower Bound. *CoRR*, abs/2507.13818, 2025.
- [11] Richard B. Borie, R. Gary Parker, and Craig A. Tovey. Automatic Generation of Linear-Time Algorithms from Predicate Calculus Descriptions of Problems on Recursively Constructed Graph Families. *Algorithmica*, 7(5-6):555–581, 1992.
- [12] Andrew Chiu, George I. Davida, and Bruce E. Litow. Division in logspace-uniform NC^1 . *RAIRO - Theoretical Informatics and Applications*, 35(3):259–275, 2001.
- [13] Bruno Courcelle. The Monadic Second-Order Logic of Graphs. I. Recognizable Sets of Finite Graphs. *Information and Computation*, 85(1):12–75, 1990.

- [14] Bruno Courcelle. The monadic second-order logic of graphs III: tree-decompositions, minor and complexity issues. *RAIRO - Theoretical Informatics and Applications*, 26:257–286, 1992.
- [15] Bruno Courcelle. The Expression of Graph Properties and Graph Transformations in Monadic Second-Order Logic. In *Handbook of Graph Grammars and Computing by Graph Transformations, Volume 1: Foundations*, pages 313–400. World Scientific, 1997.
- [16] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [17] Marek Cygan, Jesper Nederlof, Marcin Pilipczuk, Michał Pilipczuk, Johan M. M. van Rooij, and Jakub Onufry Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. *ACM Trans. Algorithms*, 18(2):17:1–17:31, 2022.
- [18] Michael Elberfeld, Martin Grohe, and Till Tantau. Where First-Order and Monadic Second-Order Logic Coincide. *ACM Transactions on Computational Logic*, 17(4), September 2016.
- [19] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.
- [20] Jacob Focke, Dániel Marx, Fionn Mc Inerney, Daniel Neuen, Govind S. Sankar, Philipp Schep- per, and Philip Wellnitz. Tight Complexity Bounds for Counting Generalized Dominating Sets in Bounded-Treewidth Graphs - Part I: Algorithmic Results. *ACM Trans. Algorithms*, 21(3):27:1–27:45, 2025.
- [21] Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. *Annals of Pure and Applied Logic*, 130(1-3):3–31, 2004.
- [22] Martin Fürer and Huiwen Yu. Space Saving by Dynamic Algebraization Based on Tree-Depth. *Theory Comput. Syst.*, 61(2):283–304, 2017.
- [23] Jakob Greilhuber and Dániel Marx. The Price of Being Partial: Complexity of Partial Generalized Dominating Set on Bounded-Treewidth Graphs. *CoRR*, abs/2506.01645, 2025.
- [24] Jakob Greilhuber, Philipp Schepper, and Philip Wellnitz. Residue Domination in Bounded-Treewidth Graphs. In Olaf Beyersdorff, Michał Pilipczuk, Elaine Pimentel, and Kim Thang Nguyen, editors, *42nd International Symposium on Theoretical Aspects of Computer Science, STACS 2025, March 4-7, 2025, Jena, Germany*, volume 327 of *LIPICs*, pages 41:1–41:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.
- [25] Martin Grohe. Algorithmic Meta Theorems. In *Proc. of the 34th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, volume 5344 of *Lecture Notes in Computer Science*, page 30, 2008.
- [26] Martin Grohe and Stephan Kreutzer. Methods for Algorithmic Meta Theorems. In *Model Theoretic Methods in Finite Combinatorics - AMS-ASL Joint Special Session*, volume 558, pages 181–206. AMS, 2009.
- [27] Falko Hegerfeld and Stefan Kratsch. Solving Connectivity Problems Parameterized by Treedepth in Single-Exponential Time and Polynomial Space. In Christophe Paul and Markus Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPICs*, pages 29:1–29:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

- [28] Stephan Kreutzer. Algorithmic meta-theorems. In *Finite and Algorithmic Model Theory*, volume 379 of *London Mathematical Society Lecture Note Series*, pages 177–270. Cambridge University Press, 2011.
- [29] Daniel Lokshтанov, Matthias Mnich, and Saket Saurabh. Planar k -path in subexponential time and polynomial space. In Petr Kolman and Jan Kratochvíl, editors, *Graph-Theoretic Concepts in Computer Science - 37th International Workshop, WG 2011, Teplá Monastery, Czech Republic, June 21-24, 2011. Revised Papers*, volume 6986 of *Lecture Notes in Computer Science*, pages 262–270. Springer, 2011.
- [30] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [31] Wojciech Nadara, Michał Pilipczuk, and Marcin Smulewicz. Computing Treedepth in Polynomial Space and Linear FPT Time. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms, ESA 2022, September 5-9, 2022, Berlin/Potsdam, Germany*, volume 244 of *LIPIcs*, pages 79:1–79:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [32] Jesper Nederlof, Michał Pilipczuk, Céline M. F. Swennenhuis, and Karol Wegrzycki. Hamiltonian cycle parameterized by treedepth in single exponential time and polynomial space. *SIAM J. Discret. Math.*, 37(3):1566–1586, 2023.
- [33] Jaroslav Nešetřil and Patrice Ossona de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *European Journal of Combinatorics*, 27(6):1022–1041, 2006.
- [34] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*. Springer Berlin, Heidelberg, 2012.
- [35] Michał Pilipczuk. Problems Parameterized by Treewidth Tractable in Single Exponential Time: A Logical Approach. In Filip Murlak and Piotr Sankowski, editors, *Proceedings of the 36th International Symposium on Mathematical Foundations of Computer Science 2011 (MFCS 2011)*, volume 6907 of *Lecture Notes in Computer Science*, pages 520–531. Springer, 2011.
- [36] Michał Pilipczuk. Graph classes through the lens of logic. *arXiv preprint arXiv:2501.04166*, 2025.
- [37] Michał Pilipczuk and Marcin Wrochna. On Space Efficiency of Algorithms Working on Structural Decompositions of Graphs. *ACM Trans. Comput. Theory*, 9(4):18:1–18:36, 2018.
- [38] Sebastian Siebertz and Alexandre Vigny. Advances in Algorithmic Meta Theorems (Invited Paper). In Siddharth Barman and Slawomir Lasota, editors, *44th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2024, December 16-18, 2024, Gandhinagar, Gujarat, India*, volume 323 of *LIPIcs*, pages 2:1–2:29. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
- [39] Jan Arne Telle and Andrzej Proskurowski. Algorithms for vertex partitioning problems on partial k -trees. *SIAM Journal on Discrete Mathematics*, 10(4):529–550, 1997.
- [40] Johan M. M. van Rooij. Fast algorithms for join operations on tree decompositions. In Fedor V. Fomin, Stefan Kratsch, and Erik Jan van Leeuwen, editors, *Treewidth, Kernels, and*

Algorithms - Essays Dedicated to Hans L. Bodlaender on the Occasion of His 60th Birthday, volume 12160 of *Lecture Notes in Computer Science*, pages 262–297. Springer, 2020.

- [41] Johan M. M. van Rooij, Hans L. Bodlaender, and Peter Rossmanith. Dynamic programming on tree decompositions using generalised fast subset convolution. In Amos Fiat and Peter Sanders, editors, *Algorithms - ESA 2009, 17th Annual European Symposium, Copenhagen, Denmark, September 7-9, 2009. Proceedings*, volume 5757 of *Lecture Notes in Computer Science*, pages 566–577. Springer, 2009.

A Discrete Fourier Transforms for Polynomials

In this appendix, we carefully check that Theorems 3.3 and 3.4 holds for polynomials as well. We recall that D is a finite set, r is a positive integer, and p is a prime such that $\mathbb{F}_p$ admits the r .th root of unity ω . For $y, q: D \rightarrow \mathbb{Z}_r$ we define $y + q: D \rightarrow \mathbb{Z}_r$ as $(y + q)(d) \equiv_r y(d) + q(d)$ for every $d \in D$. For $h, g: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ we define also $h \cdot g: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ via $(f \cdot g)(y) \equiv_p h(y) \cdot g(y)$ for every $y \in D \rightarrow \mathbb{Z}_r$. In the following, we use DFT and DFT^{-1} as shortcuts for $\text{DFT}_{D,r,p}$ and $\text{DFT}_{D,r,p}^{-1}$, respectively.

Lemma A.1. For $k \in \mathbb{N}$, $h^1, \dots, h^k: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$, we have

$$\text{DFT}(h^1 + \dots + h^k) \equiv_p \text{DFT}(h^1) + \dots + \text{DFT}(h^k) \quad (33)$$

and

$$\text{DFT}^{-1}(h^1 + \dots + h^k) \equiv_p \text{DFT}^{-1}(h^1) + \dots + \text{DFT}^{-1}(h^k). \quad (34)$$

Proof. Let $y: D \rightarrow \mathbb{Z}_r$ be arbitrary: First, it holds that

$$\begin{aligned} \text{DFT}(h^1 + \dots + h^k)(y) &\equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} (h^1 + \dots + h^k)(q) \\ &\equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} (h^1(q) + \dots + h^k(q)) \equiv_p \left(\sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} (h^1(q)) \right) + \dots + \left(\sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} (h^k(q)) \right) \\ &\equiv_p (\text{DFT}(h^1))(y) + \dots + (\text{DFT}(h^k))(y) \equiv_p (\text{DFT}(h^1) + \dots + \text{DFT}(h^k))(y). \end{aligned}$$

And similarly

$$\begin{aligned} \text{DFT}^{-1}(h^1 + \dots + h^k)(y) &\equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} (h^1 + \dots + h^k)(q) \\ &\equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} (h^1(q) + \dots + h^k(q)) \equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} (h^1(q) + \dots + h^k(q)) \\ &\equiv_p \left(\frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} h^1(q) \right) + \dots + \left(\frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} h^k(q) \right) \\ &\equiv_p (\text{DFT}^{-1}(h^1))(y) + \dots + (\text{DFT}^{-1}(h^k))(y) \equiv_p (\text{DFT}^{-1}(h^1) + \dots + \text{DFT}^{-1}(h^k))(y). \end{aligned}$$

□

For a vector $\bar{a} = (a^1, \dots, a^\ell) \in \mathbb{N}^\ell$ and $h: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ we define the mapping $h \cdot \prod_{i=1}^\ell \gamma_i^{a^i}: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$ as

$$\left(h \cdot \prod_{i=1}^\ell \gamma_i^{a^i} \right)(y) \equiv_p h(y) \cdot \prod_{i=1}^\ell \gamma_i^{a^i}$$

for all $y: D \rightarrow \mathbb{Z}_r$.

Lemma A.2. Let $h: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$, $\bar{a} = (a^1, \dots, a^\ell) \in \mathbb{N}^\ell$. Then it holds that

$$\text{DFT}\left(h \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right) \equiv_p \text{DFT}(h) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \quad (35)$$

and

$$\text{DFT}^{-1}\left(h \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right) \equiv_p \text{DFT}^{-1}(h) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \quad (36)$$

Proof. Let $y: D \rightarrow \mathbb{Z}_r$ be arbitrary: First, it holds that:

$$\begin{aligned} \text{DFT}\left(h \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)(y) &\equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} \left(h \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)(q) \equiv_p \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} h(q) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \\ &\equiv_p \left(\sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{q \cdot y} h(q) \right) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \equiv_p \text{DFT}(h)(y) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}. \end{aligned}$$

And similarly,

$$\begin{aligned} \text{DFT}^{-1}\left(h \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)(y) &\equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} \left(h \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)(q) \equiv_p \frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} h(q) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \\ &\equiv_p \left(\frac{1}{r^{|D|}} \sum_{q: D \rightarrow \mathbb{Z}_r} \omega^{-q \cdot y} h(q) \right) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \equiv_p \text{DFT}^{-1}(h)(y) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}. \end{aligned}$$

□

Further, for all $\mathbf{a} = (a^1, \dots, a^\ell) \in \mathbb{N}^\ell$ we define the functions $h_{\bar{a}}: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p$ in such a way that

$$h(y) \equiv_p \sum_{\bar{a} \in \mathbb{N}^\ell} h_{\bar{a}}(y) \prod_{i=1}^{\ell} \gamma_i^{a^i}.$$

holds for every $y: D \rightarrow \mathbb{Z}_r$ —since every polynomial is uniquely determined by the coefficients at its monomials, these functions are well-defined. By definition we then have

$$h \equiv_p \sum_{\bar{a} \in \mathbb{N}^\ell} h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \quad (37)$$

Now we are ready to prove the two main results about DFTs for polynomials from $\mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$ by relying on their analogues for $\mathbb{F}_p$:

Theorem A.3. For every $h: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$ it holds that

$$\text{DFT}^{-1}(\text{DFT}(h)) \equiv_p h \text{ and } \text{DFT}(\text{DFT}^{-1}(h)) \equiv_p h. \quad (38)$$

Proof. It holds that:

$$\begin{aligned}
& \text{DFT}^{-1}(\text{DFT}(h)) \stackrel{(37)}{\equiv_p} \text{DFT}^{-1}\left(\text{DFT}\left(\sum_{\bar{a} \in \mathbb{N}^\ell} h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)\right) \stackrel{(33)}{\equiv_p} \text{DFT}^{-1}\left(\sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}\left(h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)\right) \\
& \stackrel{(35)}{\equiv_p} \text{DFT}^{-1}\left(\sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}(h_{\bar{a}}) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right) \stackrel{(34)}{\equiv_p} \sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}^{-1}\left(\text{DFT}(h_{\bar{a}}) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right) \\
& \stackrel{(36)}{=} \sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}^{-1}(\text{DFT}(h_{\bar{a}})) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \stackrel{\text{Theorem 3.3}}{\equiv_p} \sum_{\bar{a} \in \mathbb{N}^\ell} h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \stackrel{(37)}{\equiv_p} h.
\end{aligned}$$

And similarly,

$$\begin{aligned}
& \text{DFT}(\text{DFT}^{-1}(h)) \stackrel{(37)}{\equiv_p} \text{DFT}\left(\text{DFT}^{-1}\left(\sum_{\bar{a} \in \mathbb{N}^\ell} h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)\right) \stackrel{(34)}{\equiv_p} \text{DFT}\left(\sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}^{-1}\left(h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right)\right) \\
& \stackrel{(36)}{\equiv_p} \text{DFT}\left(\sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}^{-1}(h_{\bar{a}}) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right) \stackrel{(33)}{\equiv_p} \sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}\left(\text{DFT}^{-1}(h_{\bar{a}}) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i}\right) \\
& \stackrel{(35)}{\equiv_p} \sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}(\text{DFT}^{-1}(h_{\bar{a}})) \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \stackrel{\text{Theorem 3.3}}{\equiv_p} \sum_{\bar{a} \in \mathbb{N}^\ell} h_{\bar{a}} \cdot \prod_{i=1}^{\ell} \gamma_i^{a^i} \stackrel{(37)}{\equiv_p} h.
\end{aligned}$$

□

And it also holds that

Theorem A.4. Let $t \in \mathbb{N}$, let $h^1, \dots, h^t: (D \rightarrow \mathbb{Z}_r) \rightarrow \mathbb{F}_p[\gamma_1, \dots, \gamma_\ell]$, and let $y \in D \rightarrow \mathbb{Z}_r$. Then it holds that

$$(\text{DFT}^{-1}(\text{DFT}(h^1) \cdot \dots \cdot \text{DFT}(h^t)))(y) \equiv_p \sum_{\substack{y^1, \dots, y^t: D \rightarrow \mathbb{Z}_r \\ y^1 + \dots + y^t \equiv_r y}} h^1(y^1) \cdot \dots \cdot h^t(y^t).$$

Proof.

$$\begin{aligned}
& (\text{DFT}^{-1}(\text{DFT}(h^1) \cdot \dots \cdot \text{DFT}(h^t)))(y) \equiv_p (\text{DFT}^{-1}\left(\prod_{i=1}^t \text{DFT}(h^i)\right))(y) \\
& \stackrel{(37)}{\equiv_p} (\text{DFT}^{-1}\left(\prod_{i=1}^t \text{DFT}\left(\sum_{\bar{a}_i \in \mathbb{N}^\ell} (h_{\bar{a}_i}^i \cdot \prod_{j=1}^{\ell} \gamma_j^{a_j^i})\right)\right))(y) \stackrel{(33)}{\equiv_p} (\text{DFT}^{-1}\left(\prod_{i=1}^t \sum_{\bar{a}_i \in \mathbb{N}^\ell} \text{DFT}\left(h_{\bar{a}_i}^i \cdot \prod_{j=1}^{\ell} \gamma_j^{a_j^i}\right)\right))(y) \\
& \stackrel{(35)}{\equiv_p} (\text{DFT}^{-1}\left(\prod_{i=1}^t \sum_{\bar{a}_i \in \mathbb{N}^\ell} (\text{DFT}(h_{\bar{a}_i}^i) \cdot \prod_{j=1}^{\ell} \gamma_j^{a_j^i})\right))(y) \equiv_p (\text{DFT}^{-1}\left(\sum_{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell} \prod_{i=1}^t (\text{DFT}(h_{\bar{a}_i}^i) \cdot \prod_{j=1}^{\ell} \gamma_j^{a_j^i})\right))(y) \\
& \equiv_p (\text{DFT}^{-1}\left(\sum_{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell} \left(\prod_{j=1}^{\ell} \gamma_j^{\sum_{i=1}^t a_j^i}\right) \left(\prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i)\right)\right))(y)
\end{aligned}$$

$$\begin{aligned}
& \equiv_p (\text{DFT}^{-1} \left(\sum_{\bar{a} \in \mathbb{N}^\ell} \sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \left(\prod_{j=1}^\ell \gamma_j^{\sum_{i=1}^t a_i^j} \right) \left(\prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \right))(y) \\
& \equiv_p (\text{DFT}^{-1} \left(\sum_{\bar{a} \in \mathbb{N}^\ell} \sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \left(\prod_{j=1}^\ell \gamma_j^{a_j} \right) \left(\prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \right))(y) \\
& \equiv_p (\text{DFT}^{-1} \left(\sum_{\bar{a} \in \mathbb{N}^\ell} \left(\sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \cdot \prod_{j=1}^\ell \gamma_j^{a_j} \right))(y) \\
& \stackrel{(34)}{\equiv_p} \left(\sum_{\bar{a} \in \mathbb{N}^\ell} \text{DFT}^{-1} \left(\left(\sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \cdot \prod_{j=1}^\ell \gamma_j^{a_j} \right) \right) (y) \\
& \stackrel{(36)}{\equiv_p} \left(\sum_{\bar{a} \in \mathbb{N}^\ell} \left(\text{DFT}^{-1} \left(\sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \cdot \prod_{j=1}^\ell \gamma_j^{a_j} \right) \right) (y) \\
& \stackrel{(34)}{\equiv_p} \left(\sum_{\bar{a} \in \mathbb{N}^\ell} \left(\sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \text{DFT}^{-1} \left(\prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \right) \cdot \prod_{j=1}^\ell \gamma_j^{a_j} \right) (y) \\
& \equiv_p \sum_{\bar{a} \in \mathbb{N}^\ell} \left(\sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \left(\left(\text{DFT}^{-1} \left(\prod_{i=1}^t \text{DFT}(h_{\bar{a}_i}^i) \right) \right) (y) \right) \cdot \prod_{j=1}^\ell \gamma_j^{a_j} \right) \\
& \stackrel{\text{Theorem 3.4}}{\equiv_p} \sum_{\bar{a} \in \mathbb{N}^\ell} \left(\sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \left(\sum_{\substack{y_1, \dots, y_t: D \rightarrow \mathbb{Z}_r: \\ y_1 + \dots + y_t \equiv_r y}} \prod_{i=1}^t h_{\bar{a}_i}^i(y_i) \right) \right) \cdot \prod_{j=1}^\ell \gamma_j^{a_j} \\
& \equiv_p \sum_{\bar{a} \in \mathbb{N}^\ell} \sum_{\substack{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell: \\ \bar{a}_1 + \dots + \bar{a}_t = \bar{a}}} \sum_{\substack{y_1, \dots, y_t: D \rightarrow \mathbb{Z}_r: \\ y_1 + \dots + y_t \equiv_r y}} \prod_{i=1}^t \left(h_{\bar{a}_i}^i(y_i) \cdot \left(\prod_{j=1}^\ell \gamma_j^{a_j} \right) \right) \\
& \equiv_p \sum_{\substack{y_1, \dots, y_t: D \rightarrow \mathbb{Z}_r: \\ y_1 + \dots + y_t \equiv_r y}} \sum_{\bar{a}_1, \dots, \bar{a}_t \in \mathbb{N}^\ell} \prod_{i=1}^t \left(h_{\bar{a}_i}^i(y_i) \cdot \left(\prod_{j=1}^\ell \gamma_j^{a_j} \right) \right) \\
& \equiv_p \sum_{\substack{y_1, \dots, y_t: D \rightarrow \mathbb{Z}_r: \\ y_1 + \dots + y_t \equiv_r y}} \prod_{i=1}^t \sum_{\bar{a}_i \in \mathbb{N}^\ell} \left(h_{\bar{a}_i}^i(y_i) \prod_{j=1}^\ell \gamma_j^{a_j} \right) \\
& \stackrel{(37)}{\equiv_p} \sum_{\substack{y_1, \dots, y_t: D \rightarrow \mathbb{Z}_r: \\ y_1 + \dots + y_t \equiv_r y}} \prod_{i=1}^t h^i(y_i) \equiv_p \sum_{\substack{y_1, \dots, y_t: D \rightarrow \mathbb{Z}_r: \\ y_1 + \dots + y_t \equiv_r y}} h^1(y_1) \cdot \dots \cdot h^t(y_t)
\end{aligned}$$

□