Hierarchies of Minion Tests for PCSPs through Tensors*

Lorenzo Ciardo
TU Graz
lorenzo.ciardo@tugraz.at

Stanislav Živný
University of Oxford
standa.zivny@cs.ox.ac.uk

Abstract

We provide a unified framework to study hierarchies of relaxations for Constraint Satisfaction Problems and their Promise variant. The idea is to split the description of a hierarchy into an algebraic part, depending on a *minion* capturing the "base level", and a geometric part – which we call *tensorisation* – inspired by multilinear algebra. We exploit the structure of the tensor spaces arising from our construction to prove general properties of hierarchies. We identify certain classes of minions, which we call *linear* and *conic*, whose corresponding hierarchies have particularly fine features. We establish that the (combinatorial) bounded width, Sherali–Adams LP, affine IP, Sum-of-Squares SDP, and combined "LP + affine IP" hierarchies are all captured by this framework. In particular, in order to analyse the Sum-of-Squares SDP hierarchy, we also characterise the solvability of the standard SDP relaxation through a new minion.

1 Introduction

What are the limits of efficient algorithms and where is the precise borderline of tractability? The *constraint satisfaction problem* (CSP) offers a general framework for studying such fundamental questions for a large class of computational problems [49, 50, 81] but yet for a class that is amenable to identifying the mathematical structure governing tractability. Canonical examples of CSPs are satisfiability of 3-CNF formulas (3-Sat), "not-all-equal" satisfiability of 3-CNF formulas (3-Nae-Sat), linear equations, several variants of (hyper)graph colourings, and the graph clique problem. All CSPs can be seen as homomorphism problems between relational structures [61]: Given two relational structures $\mathbf{X}$ and $\mathbf{A}$, is there a homomorphism from $\mathbf{X}$ to $\mathbf{A}$? Intuitively, the structure $\mathbf{X}$ represents the variables of the CSP instance and their interactions, whereas the structure $\mathbf{A}$ represents the constraint language; i.e., the alphabet and the allowed constraint relations.

The most studied types of CSPs are so-called *non-uniform* CSPs [16, 61, 72, 78], in which the target structure $\mathbf{A}$ is fixed whereas the source structure $\mathbf{X}$ is given on input; this computational problem is denoted by $\mathrm{CSP}(\mathbf{A})$. From the examples above, 3-Sat, 3-Nae-Sat,

*An extended abstract of this work appeared in the Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23) [44]. The research leading to these results has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714532). The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein. This work was also supported by UKRI EP/X024431/1. The work was done while the first author was at the University of Oxford.

All data is provided in full in the results section of this paper.

(hyper)graph colourings with constantly many colours, linear equations of bounded arity over finite fields, and linear equations of bounded arity over the rationals are all examples of non-uniform CSPs, all on finite domains except the last one. For instance, in the graph c-colouring problem the target structure $\mathbf{A}$ is a c-clique and the structure $\mathbf{X}$ is the input graph. The existence of a homomorphism from a graph to a c-clique is equivalent to the existence of a colouring of the graph with c colours. The graph clique problem is an example of a CSP with a fixed class of source structures [65,91] but an arbitrary target structure and, thus, it is not a non-uniform CSP.

We will be concerned with polynomial-time tractability of CSPs. Studied research directions include investigating questions such as: Is there a solution [34,100]? How many solutions are there, exactly [33,48,59] or approximately [35,39]? What is the maximum number of simultaneously satisfied constraints, exactly [47,71,97] or approximately [7,56,95]? What is the minimum number of simultaneously unsatisfied constraints [54,74]? Given an almost satisfiable instance, can one find a somewhat satisfying solution [14,52,53]? In this paper, we will focus on the following question:

Given a satisfiable instance, can one find a solution that is satisfying in a weaker sense [9, 12, 23]?

This was formalised as *promise constraint satisfaction problems* (PCSPs) by Austrin, Guruswami and Håstad [9] and Brakensiek and Guruswami [23]. Let $\mathbf{A}$ and $\mathbf{B}$ be two fixed relational structures[1] such that there is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$, indicated by $\mathbf{A} \to \mathbf{B}$. Intuitively, the structure $\mathbf{A}$ represents the "strict" constraints and the structure $\mathbf{B}$ represents the corresponding "weak" constraints. An instance of the PCSP over the template $(\mathbf{A}, \mathbf{B})$, denoted by $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$, is a relational structure $\mathbf{X}$ such that there is a homomorphism from $\mathbf{X}$ to $\mathbf{A}$. The task is to find a homomorphism from $\mathbf{X}$ to $\mathbf{B}$, which exists since homomorphisms compose. What we described above is the *search* variant of the PCSP. In the *decision* variant, one is given a relational structure $\mathbf{X}$ and the task is to decide whether there is a homomorphism from $\mathbf{X}$ to $\mathbf{A}$ or whether there is not a homomorphism from $\mathbf{X}$ to $\mathbf{B}$. As $\mathbf{X} \to \mathbf{A}$ implies $\mathbf{X} \to \mathbf{B}$, the two cases cannot happen simultaneously. Clearly, the decision variant of the PCSP reduces to the search variant (see the discussion in [12] after Definition 2.6), but it is not known whether there is a reduction in the other direction for all PCSPs. In this paper, we shall focus on the decision variant.

PCSPs are a vast generalisation of CSPs including problems that cannot be expressed as CSPs. The work of Barto, Bulín, Krokhin, and Opršal [12] lifted and greatly extended the algebraic framework developed for CSPs [17,32,72] to the realm of PCSPs. Subsequently, there has been a series of recent works on the computational complexity of PCSPs building on [12], including applicability of local consistency and convex relaxations [5,22,27,36,43] and complexity of fragments of PCSPs [2,11,15,24,28,67,82,93]. Strong results on PCSPs have also been established via other techniques than those in [12], mostly analytical methods, e.g., hardness of various (hyper)graph colourings [8,58,70,75] and other PCSPs [19,20,25,30].

An example of a PCSP, identified in [9], is (in the search variant) finding a satisfying assignment to a k-CNF formula given that a g-satisfying assignment exists; i.e., an assignment that satisfies at least g literals in each clause. Austrin et al. [9] established that this problem is NP-hard if $g/k < 1/2$ and solvable via a constant level of the Sherali–Adams linear

[1] Unless otherwise stated, we shall use the word "structure" to mean finite-domain structures; if the domain is allowed to be infinite, we shall say it explicitly.

programming relaxation otherwise. This classification was later extended to problems over arbitrary finite domains by Brandts et al. [28].

A second example of a PCSP, identified in [23], is (in the search variant) finding a "not-all-equal" assignment to a monotone 3-CNF formula given that a "1-in-3" assignment is promised to exist; i.e., given a 3-CNF formula with positive literals only and the promise that an assignment exists that satisfies exactly one literal in each clause, the task is to find an assignment that satisfies one or two literals in each clause. This problem is solvable in polynomial time via a constant level of the Sherali–Adams linear programming relaxation [23] but not via a reduction to finite-domain CSPs [12].

A third example of a PCSP is the well-known *approximate graph colouring* problem: Given a c-colourable graph, find a d-colouring of it, for constants c and d with $c \leq d$. This corresponds to $\mathrm{PCSP}(\mathbf{K}_c, \mathbf{K}_d)$, where $\mathbf{K}_p$ is the clique on p vertices. Despite a long history dating back to 1976 [62], the complexity of this problem is only understood under stronger assumptions [29, 57, 67] and for special cases [12, 21, 66, 70, 73, 75, 82]. It is believed that the problem is NP-hard already in the decision variant [62], i.e., deciding whether a graph is c-colourable or not even d-colourable, for each $3 \leq c \leq d$. By using the framework developed in the current work, non-solvability of approximate graph colouring through standard algorithmic techniques was established in the follow-up works [41, 42].

PCSPs can be solved by designing *tests*. If a test, applied to a given instance of the problem, is positive then the answer is YES; if it is negative then the answer is NO. The challenge is then to find tests that are able to guarantee a low number – ideally, zero – of false positives and false negatives. Clearly, a test is itself a decision problem. However, its nature may be substantially different, and less complicated, than the nature of the original problem.

Given a PCSP template $(\mathbf{A}, \mathbf{B})$, we may use any (potentially infinite) structure $\mathbf{T}$ to make a test for $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$: We simply let the outcome of the test on an instance structure $\mathbf{X}$ be YES if $\mathbf{X} \to \mathbf{T}$, and NO if $\mathbf{X} \not\to \mathbf{T}$. In other words, $\mathrm{CSP}(\mathbf{T})$ is a test for $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$. Let $\mathbf{X}$ be an instance of $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$. If $\mathbf{X} \to \mathbf{T}$ whenever $\mathbf{X} \to \mathbf{A}$, the test is guaranteed not to generate false negatives, and we call it *complete*. If $\mathbf{X} \to \mathbf{B}$ whenever $\mathbf{X} \to \mathbf{T}$, the test is guaranteed not to generate false positives, and we call it *sound*. Since homomorphisms compose, the test is complete if and only if $\mathbf{A} \to \mathbf{T}$, and it is sound if and only if $\mathbf{T} \to \mathbf{B}$.[2] When both of these conditions hold, we say that the test *solves* $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$. In this case, one obtains a so-called "sandwich reduction" from $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ to $\mathrm{CSP}(\mathbf{T})$ (see [22, Section 3.1]).

To make a test $\mathbf{T}$ useful as a polynomial-time algorithm to solve a PCSP, one requires that $\mathrm{CSP}(\mathbf{T})$ should be tractable. It was conjectured in [22, Section 3.1] that every tractable (finite-domain) PCSP is solved by a tractable sandwich. In other words, if the conjecture is true, *sandwich reductions are the sole source of tractability for PCSPs*. For the conjecture to be true, one needs to admit structures $\mathbf{T}$ having infinite domains. For example, this is the case for the "1-in-3 vs. not-all-equal" problem, whose template we denote by $(\mathbf{1\text{-}in\text{-}3}, \mathbf{NAE})$: As shown in [12], there is no (finite) structure $\mathbf{T}$ such that $\mathbf{1\text{-}in\text{-}3} \to \mathbf{T} \to \mathbf{NAE}$ and $\mathrm{CSP}(\mathbf{T})$ is tractable.

The complexity of both CSPs and PCSPs was shown to be determined by higher-order symmetries of the solution sets of the problems, known as *polymorphisms*, denoted by $\mathrm{Pol}(\mathbf{A})$ for $\mathrm{CSP}(\mathbf{A})$ [32] and by $\mathrm{Pol}(\mathbf{A}, \mathbf{B})$ for $\mathrm{PCSP}(\mathbf{A}, \mathbf{B})$ [12, 23]. For CSPs, polymorphisms form *clones*; in particular, they are closed under composition. This means that some symmetries may be obtainable through compositions of other symmetries, so that one can hope to capture

[2] The second "only if" implication follows from a standard compactness argument, see Section 3.

properties of entire families of CSPs (e.g., bounded width, tractability, etc.) through the presence of a certain polymorphism and, more generally, to describe their complexity through universal-algebraic tools. A chief example of this approach is the positive resolution of the dichotomy conjecture for CSPs by Bulatov [34] and Zhuk [100], establishing that finite-domain non-uniform CSPs are either in P or are NP-complete. For PCSPs, however, polymorphisms are not closed under composition, and the algebraic structure they are endowed with – known as a *minion* – is much less rich than clones. In particular, the structure theory of finite algebras, a central topic of universal algebra, is not applicable in this setting.

For a PCSP template $(\mathbf{A}, \mathbf{B})$, one would ideally aim to build tests for $\text{PCSP}(\mathbf{A}, \mathbf{B})$ in a systematic way. One method to do so is by considering tests associated with minions and, in particular, their *free structures*. The free structure $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$ of a minion $\mathscr{M}$ generated by a structure $\mathbf{A}$ [12, Definition 4.1] is a (potentially infinite) structure obtained, essentially, by simulating the relations in $\mathbf{A}$ on a domain consisting of elements of $\mathscr{M}$. Then, we define $\text{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{YES}$ if $\mathbf{X} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A})$, and NO otherwise. (Note that $\mathbf{X}$ is the input to the problem; the minion $\mathscr{M}$ and the relational structure $\mathbf{A}$, coming from a PCSP template, are (fixed) parameters of the test.)

For certain choices of $\mathscr{M}$, $\text{Test}_{\mathscr{M}}$ is *always* a tractable test; i.e., $\text{CSP}(\mathbb{F}_{\mathscr{M}}(\mathbf{A}))$ is tractable for any $\mathbf{A}$. This is the case for the minions $\mathscr{H} = \text{Pol}(\text{HORN-3-SAT})$ (whose elements are nonempty subsets of a given set), $\mathscr{Q}_{\text{conv}}$ (whose elements are stochastic vectors), and $\mathscr{Z}_{\text{aff}}$ (whose elements are affine integers vectors). As it was shown in [12], these three minions correspond to three well-studied algorithmic relaxations: $\text{Test}_{\mathscr{H}}$ is Arc Consistency (AC) [89], $\text{Test}_{\mathscr{Q}_{\text{conv}}}$ is the Basic Linear Programming relaxation (BLP) [83], and $\text{Test}_{\mathscr{Z}_{\text{aff}}}$ is the Affine Integer Programming relaxation ($\mathscr{Z}_{\text{aff}}$) [22]. In [27], the algorithm BLP + AIP (which we shall call BA in this work) corresponding to a combination of linear and integer programming was shown to be captured by a certain minion $\mathscr{M}_{\text{BA}}$. In summary, several widely used algorithms for (P)CSPs are minion tests; in particular, Arc Consistency, which is the simplest example of consistency algorithms, and standard algorithms based on relaxations.

Convex relaxations have been instrumental in the understanding of the complexity of many variants of CSPs, including constant approximability of Min-CSPs [54, 60] and Max-CSPs [76, 95], robust satisfiability of CSPs [14, 83, 101], and exact solvability of optimisation CSPs [79, 98]. An important line of work focused on making convex relaxations stronger and stronger via the so-called "lift-and-project" method, which includes the Sherali–Adams LP hierarchy [96], the SDP hierarchy of Lovász and Schrijver [88], and the (stronger) SDP hierarchy of Lasserre [84], also known as the Sum-of-Squares hierarchy (see [85] for a comparison of these hierarchies). The study of the power of various convex hierarchies has led to several breakthroughs, e.g., [1, 37, 64, 80, 86, 99].

In the same spirit as lift-and-project hierarchies of convex relaxations, the (combinatorial) k-consistency algorithm (also known as the k-bounded width algorithm) has a central role in the study of tractability for constraint satisfaction problems [3, 61]. Here k is an integer bounding the number of variables considered in reasoning about partial solutions; the case $k = 1$ corresponds to Arc Consistency mentioned above. The notion of local consistency, in addition to being one of the key concepts in constraint satisfaction, has also emerged independently in finite model theory [77], graph theory [68], and proof complexity [4]. The power of local consistency for (non-promise) CSPs is now fully understood [10, 13, 31]. Recent works identified necessary conditions on local consistency to solve PCSPs [5, 38, 45]. Very recent work established the importance of infinite-domain sandwiches for promise CSPs [92, 94].

Contributions The main contribution of this work is the introduction of *a general framework for refining algorithmic relaxations of* (P)CSPs. Given a minion $\mathscr{M}$, we present a technique to systematically turn $\text{Test}_{\mathscr{M}}$ into the corresponding *hierarchy of minion tests*: a sequence of increasingly tight relaxations $\text{Test}^k_{\mathscr{M}}$ for $k \in \mathbb{N}$.

The technique we adopt to build hierarchies of minion tests is inspired by multilinear algebra. We describe a *tensorisation* construction that turns a given structure $\mathbf{X}$ into a structure $\mathbf{X}^{\textcircled{k}}$ on a different signature, where both the domain and the relations are multidimensional objects living in tensor spaces. Essentially, $\text{Test}^k_{\mathscr{M}}$ works by applying $\text{Test}_{\mathscr{M}}$ to *tensorised* versions of the structures $\mathbf{X}$ and $\mathbf{A}$ rather than to $\mathbf{X}$ and $\mathbf{A}$ themselves. This allows us to study the functioning of the algorithms in the hierarchy by describing the structure of a space of tensors – which can be accomplished by using multilinear algebra. This approach has not appeared in the literature on Sherali–Adams, bounded width, Sum-of-Squares, hierarchies of integer programming, and related algorithmic techniques such as the high-dimensional Weisfeiler–Leman algorithm [6, 36]. Butti and Dalmau [36] recently characterised for CSPs when the k-th level of the Sherali–Adams LP programming hierarchy accepts in terms of a construction different from the one introduced in this work. Unlike the tensorisation, the construction considered in [36] yields a relational structure whose domain includes the set of constraints of the original structure.

One key feature of our framework is that it is *modular*, in that it allows splitting the description of a hierarchy of minion tests into an *algebraic* part, corresponding to the minion $\mathscr{M}$, and a *geometric* part, entirely dependent on the tensorisation construction and hence common to all hierarchies. By considering certain well-behaved families of minions, which we call *linear* and *conic*, we can then deduce general properties of the corresponding hierarchies by only describing the structure of tensor spaces.

Letting the minion $\mathscr{M}$ be $\mathscr{H}$ (resp., $\mathscr{Q}_{\text{conv}}$, $\mathscr{Z}_{\text{aff}}$), we shall retrieve in this way the bounded width hierarchy (resp., the Sherali–Adams LP hierarchy, the affine integer programming hierarchy). Additionally, we describe a new minion $\mathscr{S}$ capturing the power of the basic semidefinite programming relaxation (SDP),[3] and we prove that $\text{Test}^k_{\mathscr{S}}$ coincides with the Sum-of-Squares hierarchy. As a consequence, our framework is able to provide a unified description of all these four well-known hierarchies of algorithmic relaxations. In addition to casting *known* hierarchies of relaxations as hierarchies of minion tests, this approach can be used to design *new* hierarchies. In particular, we describe an operation that we call *semi-direct product* of minions, which consists in combining multiple minions to form a new minion associated with a stronger relaxation. In practice, this method can be used to design an algorithm that combines the features of different known algorithmic techniques. We show that the minion $\mathscr{M}_{\text{BA}}$ associated with the BA relaxation from [27] is the semi-direct product of $\mathscr{Q}_{\text{conv}}$ and $\mathscr{Z}_{\text{aff}}$, and we formally introduce the BA^k hierarchy as the hierarchy $\text{Test}^k_{\mathscr{M}_{\text{BA}}}$.

The scope of this framework is potentially not limited to constraint satisfaction: The multilinear pattern that we found at the core of different algorithmic hierarchies appears to be transversal to the constraint satisfaction setting and, instead, inherently connected to the algorithmic techniques themselves, which can be applied to classes of computational problems living beyond the realms of (P)CSPs.

[3]We point out that Brakensiek, Guruswami, and Sandeep [26] independently provided a minion characterisation for the power of SDP that is essentially identical to the one we obtain in the current work. We also note that the complexity of testing SDP feasibility is an open problem, cf. [63].

Subsequent work The tensorisation methodology introduced in this paper has later been used by the authors in follow-up work on the applicability of relaxation hierarchies to specific problems. In particular, it has been used to prove that the approximate graph colouring problem is not solved by the hierarchy for the combined basic linear programming and affine integer programming relaxation [41, 42]. Recently, Dalmau and Opršal [55] showed that various concepts introduced in this work – in particular, the concept of conic minions – are naturally connected to the description of a certain type of reductions between (P)CSPs, see [55, Section 4.2] as well as Remark 50 in the current paper.

Organisation The remaining part of the paper is organised as follows. Section 2 contains relevant terminology for tensors and gives a formal description of promise CSPs as well as the relaxations and hierarchies used throughout the paper (see also Appendix A). Section 3 contains some initial, basic results on minion tests. Section 4 introduces a minion $\mathscr{S}$ capturing the power of SDP. Sections 5, 6, and 7 are the technical core of the paper; they provide a description of hierarchies of tests built on arbitrary minions, linear minions, and conic minions, respectively. Section 8 describes the semi-direct product of minions, needed to capture the BA^k hierarchy. The machinery assembled in the previous sections is finally used in Section 9 to prove that five well-known algorithmic hierarchies for (P)CSPs are captured by our framework.

2 Preliminaries

Notation We denote by $\mathbb{N}$ the set of positive integers. For $k \in \mathbb{N}$, we denote by $[k]$ the set $\{1, \dots, k\}$. We indicate by $\mathbf{e}_i$ the i-th standard unit vector of the appropriate size (which will be clear from the context); i.e., the i-th entry of $\mathbf{e}_i$ is 1, and all other entries are 0. $\mathbf{0}_p$ and $\mathbf{1}_p$ denote the all-zero and all-one vector of size p, respectively, while I_p and $O_{p,q}$ denote the $p \times p$ identity matrix and the $p \times q$ all-zero matrix, respectively. Given a matrix M, we let $\mathrm{tr}(M)$ and $\mathrm{csupp}(M)$ be the trace and the set of indices of nonzero columns of M, respectively. The symbol $\aleph_0$ denotes the cardinality of $\mathbb{N}$. Given a set S, we denote the identity map on S by id_S (or simply by id when S is clear from the context).

2.1 Terminology for tensors

Tuples Given a set S, two integers $k, \ell \in \mathbb{N}$, a tuple $\mathbf{s} = (s_1, \dots, s_k) \in S^k$, and a tuple $\mathbf{i} = (i_1, \dots, i_\ell) \in [k]^\ell$, $\mathbf{s}_\mathbf{i}$ shall denote the *projection* of $\mathbf{s}$ onto $\mathbf{i}$, i.e., the tuple in S^ℓ defined by $\mathbf{s}_\mathbf{i} = (s_{i_1}, \dots, s_{i_\ell})$. Given two tuples $\mathbf{s} = (s_1, \dots, s_k) \in S^k$ and $\tilde{\mathbf{s}} = (\tilde{s}_1, \dots, \tilde{s}_\ell) \in S^\ell$, their *concatenation* is the tuple $(\mathbf{s}, \tilde{\mathbf{s}}) = (s_1, \dots, s_k, \tilde{s}_1, \dots, \tilde{s}_\ell) \in S^{k+\ell}$. We also define $\{\mathbf{s}\} = \{s_1, \dots, s_k\}$. Given two sets $S, \tilde{S}$ and two tuples $\mathbf{s} = (s_1, \dots, s_k) \in S^k$, $\tilde{\mathbf{s}} = (\tilde{s}_1, \dots, \tilde{s}_k) \in \tilde{S}^k$, we write $\mathbf{s} \prec \tilde{\mathbf{s}}$ if, for any $\alpha, \beta \in [k]$, $s_\alpha = s_\beta$ implies $\tilde{s}_\alpha = \tilde{s}_\beta$. The expression $\mathbf{s} \not\prec \tilde{\mathbf{s}}$ shall mean the negation of $\mathbf{s} \prec \tilde{\mathbf{s}}$. Notice that the relation "$\prec$" is preserved under projections: If $\mathbf{s} \prec \tilde{\mathbf{s}}$ and $\mathbf{i} \in [k]^\ell$, then $\mathbf{s}_\mathbf{i} \prec \tilde{\mathbf{s}}_\mathbf{i}$.

Semirings A *semiring* $\mathcal{S}$ consists of a set S equipped with two binary operations "+" and "·" such that

- $(S, +)$ is a commutative monoid with an identity element "$0_\mathcal{S}$" (i.e., $(r+s)+t = r+(s+t)$, $0_\mathcal{S} + r = r + 0_\mathcal{S} = r$, and $r + s = s + r$);

- $(\mathcal{S}, \cdot)$ is a monoid with an identity element "$1_\mathcal{S}$" (i.e., $(r \cdot s) \cdot t = r \cdot (s \cdot t)$ and $1_\mathcal{S} \cdot r = r \cdot 1_\mathcal{S} = r$);
- "$\cdot$" distributes over "$+$" (i.e., $r \cdot (s + t) = (r \cdot s) + (r \cdot t)$ and $(r + s) \cdot t = (r \cdot t) + (s \cdot t)$);
- "$0_\mathcal{S}$" is a multiplicative absorbing element (i.e., $0_\mathcal{S} \cdot r = r \cdot 0_\mathcal{S} = 0_\mathcal{S}$).

Examples of semirings are $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ with the usual addition and multiplication operations, or the Boolean semiring $(\{0, 1\}, \vee, \wedge)$.

Let V be a finite set, and choose an element $s_v \in \mathcal{S}$ for each $v \in V$. We let the formal expression $\sum_{v \in V} s_v$ equal $0_\mathcal{S}$ if $V = \emptyset$. To increase the readability, we shall usually write 0 and 1 for $0_\mathcal{S}$ and $1_\mathcal{S}$; the relevant semiring $\mathcal{S}$ will always be clear from the context.

Tensors Given a set S, an integer $k \in \mathbb{N}$, and a tuple $\mathbf{n} = (n_1, \dots, n_k) \in \mathbb{N}^k$, by $\mathcal{T}^\mathbf{n}(S)$ we denote the set of functions from $[n_1] \times \cdots \times [n_k]$ to S, which we visualise as hypermatrices or tensors having k *modes*, where the i-th mode has size n_i for $i \in [k]$. If $\mathbf{n} = n \cdot \mathbf{1}_k = (n, \dots, n)$ is a constant tuple, $\mathcal{T}^\mathbf{n}(S)$ is a set of *cubical* tensors, each of whose modes has the same length n. For example, if $\mathbf{n} = n \cdot \mathbf{1}_2 = (n, n)$, $\mathcal{T}^\mathbf{n}(S)$ is the set of $n \times n$ matrices having entries in S. We sometimes denote an element of $\mathcal{T}^\mathbf{n}(S)$ by $T = (t_\mathbf{i})$, where $\mathbf{i} \in [n_1] \times \cdots \times [n_k]$ and $t_\mathbf{i}$ is the image of $\mathbf{i}$ under T. Moreover, given two tuples $\mathbf{n} \in \mathbb{N}^k$ and $\tilde{\mathbf{n}} \in \mathbb{N}^\ell$, we sometimes write $\mathcal{T}^{\mathbf{n},\tilde{\mathbf{n}}}(S)$ for $\mathcal{T}^{(\mathbf{n},\tilde{\mathbf{n}})}(S)$, where $(\mathbf{n}, \tilde{\mathbf{n}})$ is the concatenation of $\mathbf{n}$ and $\tilde{\mathbf{n}}$. Whenever $k \geq 2$ and $n_i = 1$ for some $i \in [k]$, we can (and will) identify $\mathcal{T}^\mathbf{n}(S)$ with $\mathcal{T}^{\hat{\mathbf{n}}}(S)$, where $\hat{\mathbf{n}} \in \mathbb{N}^{k-1}$ is obtained from $\mathbf{n}$ by deleting the i-th entry. Note that the definition of tensors straightforwardly extends to the case that $n_i = \aleph_0$ for some $i \in [k]$.

Contraction Take a semiring $\mathcal{S}$. For $k, \ell, m \in \mathbb{N}$, take $\mathbf{n} \in \mathbb{N}^k$, $\mathbf{p} \in \mathbb{N}^\ell$, and $\mathbf{q} \in \mathbb{N}^m$. The *contraction* of two tensors $T = (t_\mathbf{i}) \in \mathcal{T}^{\mathbf{n},\mathbf{p}}(\mathcal{S})$ and $\tilde{T} = (\tilde{t}_\mathbf{i}) \in \mathcal{T}^{\mathbf{p},\mathbf{q}}(\mathcal{S})$, denoted by $T \overset{\ell}{*} \tilde{T}$, is the tensor in $\mathcal{T}^{\mathbf{n},\mathbf{q}}(\mathcal{S})$ such that, for $\mathbf{i} \in [n_1] \times \cdots \times [n_k]$ and $\mathbf{j} \in [q_1] \times \cdots \times [q_m]$, the $(\mathbf{i}, \mathbf{j})$-th entry of $T \overset{\ell}{*} \tilde{T}$ is given by

$$\sum_{\mathbf{z} \in [p_1] \times \cdots \times [p_\ell]} t_{(\mathbf{i},\mathbf{z})} \tilde{t}_{(\mathbf{z},\mathbf{j})} \tag{1}$$

(where the addition and multiplication are meant in the semiring $\mathcal{S}$). This notation straightforwardly extends to the cases when $k = 0$ or $m = 0$, i.e., when we are contracting over all modes of T or $\tilde{T}$. In such cases, we write $T * \tilde{T}$ for $T \overset{\ell}{*} \tilde{T}$. We shall use the convention of always contracting from the left to the right, so that the expression $T_1 * T_2 * T_3$ is a synonym of the expression $(T_1 * T_2) * T_3$. The reason for this is that the contraction operation is not associative in general. For instance, for T and $\tilde{T}$ as above and $\hat{T} \in \mathcal{T}^{\mathbf{n},\mathbf{q}}(\mathcal{S})$, the expression $(T \overset{\ell}{*} \tilde{T}) \overset{k+m}{*} \hat{T} = (T \overset{\ell}{*} \tilde{T}) * \hat{T}$ is well defined, while changing the order of the contractions results in an expression that is not well defined in general. On the other hand, it is not hard to check that the order in which the contractions are performed is irrelevant if the contractions are taken over disjoint sets of modes.

Example 1. Given two vectors $\mathbf{u}, \mathbf{v} \in \mathcal{T}^p(\mathbb{R})$ and two matrices $M \in \mathcal{T}^{n,p}(\mathbb{R})$, $N \in \mathcal{T}^{p,q}(\mathbb{R})$, we have that $\mathbf{u} \overset{1}{*} \mathbf{v} = \mathbf{u} * \mathbf{v} = \mathbf{u}^T \mathbf{v}$ (the dot product of $\mathbf{u}$ and $\mathbf{v}$), $M \overset{1}{*} \mathbf{u} = M * \mathbf{u} = M\mathbf{u}$ (the matrix-vector product of M and $\mathbf{u}$), and $M \overset{1}{*} N = MN$ (the matrix product of M and N), as can be easily checked by applying (1).

For $\mathbf{i} \in [n_1] \times \cdots \times [n_k]$, we denote by $E_{\mathbf{i}}$ the $\mathbf{i}$-*th standard unit tensor*; i.e., the tensor in $\mathcal{T}^{\mathbf{n}}(\mathcal{S})$ all of whose entries are $0_{\mathcal{S}}$, except the $\mathbf{i}$-th entry that is $1_{\mathcal{S}}$. Given $T \in \mathcal{T}^{\mathbf{n}}(\mathcal{S})$, notice that $E_{\mathbf{i}} \overset{k}{*} T = E_{\mathbf{i}} * T$ is the $\mathbf{i}$-th entry of T (in fact, we shall always indicate the entries of a tensor in this way). The *support* of T is the set of indices of all nonzero entries of T; i.e., the set $\text{supp}(T) = \{\mathbf{i} \in [n_1] \times \cdots \times [n_k] : E_{\mathbf{i}} * T \neq 0_{\mathcal{S}}\}$.

2.2 Promise CSPs

Structures A *signature* σ is a finite set of relation symbols R, each having an *arity* $\text{ar}(R) \in \mathbb{N}$. A σ-*structure* $\mathbf{A}$ consists of a set A (called the *domain*) and, for each $R \in \sigma$, a relation $R^{\mathbf{A}} \subseteq A^{\text{ar}(R)}$. A σ-structure $\mathbf{A}$ is finite if the size $|A|$ of its domain A is finite. In this case, we often assume that the domain of $\mathbf{A}$ is $A = [n]$. (In general, we will reserve the letter n to denote the domain size of $\mathbf{A}$.)

Let $\mathbf{A}$ and $\mathbf{B}$ be σ-structures. A *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a map $h : A \to B$ such that, for each $R \in \sigma$ with $r = \text{ar}(R)$ and for each $\mathbf{a} = (a_1, \ldots, a_r) \in A^r$, if $\mathbf{a} \in R^{\mathbf{A}}$ then $h(\mathbf{a}) = (h(a_1), \ldots, h(a_r)) \in R^{\mathbf{B}}$. We denote the existence of a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ by $\mathbf{A} \to \mathbf{B}$. Fix a pair $(\mathbf{A}, \mathbf{B})$ of σ-structures such that $\mathbf{A} \to \mathbf{B}$. The *promise constraint satisfaction problem* parameterised by the *template* $(\mathbf{A}, \mathbf{B})$, denoted by $\text{PCSP}(\mathbf{A}, \mathbf{B})$, is the following decision problem:[4] The input is a σ-structure $\mathbf{X}$ and the goal is to answer YES if $\mathbf{X} \to \mathbf{A}$ and NO if $\mathbf{X} \not\to \mathbf{B}$. The promise is that it is not the case that $\mathbf{X} \not\to \mathbf{A}$ and $\mathbf{X} \to \mathbf{B}$. We write $\text{CSP}(\mathbf{A})$ for $\text{PCSP}(\mathbf{A}, \mathbf{A})$, the classic (non-promise) constraint satisfaction problem.

Polymorphisms and minions The algebraic theory of PCSPs developed in [12, 23] relies on the notions of polymorphism and minion. Let $\mathbf{A}$ be a σ-structure. For $L \in \mathbb{N}$, the L-*th power* of $\mathbf{A}$ is the σ-structure $\mathbf{A}^L$ with domain A^L whose relations are defined as follows: Given $R \in \sigma$ and an $L \times \text{ar}(R)$ matrix M such that all rows of M are tuples in $R^{\mathbf{A}}$, the columns of M form a tuple in $R^{\mathbf{A}^L}$.

An L-ary *polymorphism* of a PCSP template $(\mathbf{A}, \mathbf{B})$ is a homomorphism from $\mathbf{A}^L$ to $\mathbf{B}$. Minions were defined in [12] as sets of functions with certain properties. We shall use here the abstract definition of minions, as first done in [27], cf. also [43]. A *minion* $\mathscr{M}$ consists in the disjoint union of nonempty sets $\mathscr{M}^{(L)}$ for $L \in \mathbb{N}$ equipped with (so-called *minor*) operations $(\cdot)_{/\pi} : \mathscr{M}^{(L)} \to \mathscr{M}^{(L')}$ for all functions $\pi : [L] \to [L']$, which satisfy $M_{/\text{id}} = M$ and, for $\pi : [L] \to [L']$ and $\tilde{\pi} : [L'] \to [L'']$, $(M_{/\pi})_{/\tilde{\pi}} = M_{/\tilde{\pi}\circ\pi}$ for all $M \in \mathscr{M}^{(L)}$.

Example 2. The set $\text{Pol}(\mathbf{A}, \mathbf{B})$ of all polymorphisms of a PCSP template $(\mathbf{A}, \mathbf{B})$ is a minion with the minor operations defined by $f_{/\pi}(a_1, \ldots, a_{L'}) = f(a_{\pi(1)}, \ldots, a_{\pi(L)})$ for $f : \mathbf{A}^L \to \mathbf{B}$ and $\pi : [L] \to [L']$. In this minion, the minor operations correspond to identifying coordinates, permuting coordinates, and introducing dummy coordinates (of polymorphisms).

Example 3. Other examples of minions that shall appear frequently in this work are $\mathscr{Q}_{\text{conv}}$, $\mathscr{Z}_{\text{aff}}$, and $\mathscr{H}$, capturing the power of the algorithms BLP, AIP, and AC, respectively. The L-ary elements of $\mathscr{Q}_{\text{conv}}$ are rational vectors of size L that are *stochastic* (i.e., whose entries are nonnegative and sum up to 1), with the minor operations defined as follows: For $\mathbf{q} \in \mathscr{Q}_{\text{conv}}{}^{(L)}$ and $\pi : [L] \to [L']$, $\mathbf{q}_{/\pi} = P\mathbf{q}$, where P is the $L' \times L$ matrix whose (i, j)-th entry is 1 if $\pi(j) = i$, and 0 otherwise. $\mathscr{Z}_{\text{aff}}$ is defined similarly to $\mathscr{Q}_{\text{conv}}$, the only difference being that its L-ary elements are affine integer vectors (i.e., their entries are integer – possibly negative – numbers

[4] From now on, we shall work only with the *decision* version of the PCSP.

and sum up to 1). $\mathscr{H}$ is the minion of polymorphisms of the CSP template Horn-3-Sat, i.e., the Boolean structure whose three relations are "$x \wedge y \Rightarrow z$", $\{0\}$, and $\{1\}$.[5] Equivalently (cf. [12]), $\mathscr{H}$ can be described as follows: For any $L \in \mathbb{N}$, the L-ary elements of $\mathscr{H}$ are Boolean functions of the form $f_Z(x_1, \ldots, x_L) = \bigwedge_{z \in Z} x_z$ for any $Z \subseteq [L]$, $Z \neq \emptyset$; the minor operations are defined as in Example 2. We shall also mention the minion $\mathscr{M}_{\text{BA}}$ capturing the algorithm BA described in the Introduction. Its L-ary elements are $L \times 2$ matrices whose first column $\mathbf{u}$ belongs to $\mathscr{Q}_{\text{conv}}{}^{(L)}$ and whose second column $\mathbf{v}$ belongs to $\mathscr{Z}_{\text{aff}}{}^{(L)}$, and such that if the i-th entry of $\mathbf{u}$ is zero then the i-th entry of $\mathbf{v}$ is also zero, for each $i \in [L]$. The minor operation is defined on each column individually; i.e., $[\ \mathbf{u} \quad \mathbf{v}\ ]_{/\pi} = [\ \mathbf{u}_{/\pi} \quad \mathbf{v}_{/\pi}\ ]$.

For two minions $\mathscr{M}$ and $\mathscr{N}$, a *minion homomorphism* $\xi : \mathscr{M} \to \mathscr{N}$ is a map that preserves arities and minors: Given $M \in \mathscr{M}^{(L)}$ and $\pi : [L] \to [L']$, $\xi(M) \in \mathscr{N}^{(L)}$ and $\xi(M_{/\pi}) = \xi(M)_{/\pi}$. We denote the existence of a minion homomorphism from $\mathscr{M}$ to $\mathscr{N}$ by $\mathscr{M} \to \mathscr{N}$. If a minion homomorphism is invertible as a function – in which case, its inverse must also be a minion homomorphism – we say that it is a *minion isomorphism*.

We will also need the concept of free structure from [12]. Let $\mathscr{M}$ be a minion and let $\mathbf{A}$ be a (finite) σ-structure. The *free structure* of $\mathscr{M}$ generated by $\mathbf{A}$ is a σ-structure $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$ with domain $\mathscr{M}^{(|A|)}$ (potentially infinite). Given a relation symbol $R \in \sigma$ of arity r, a tuple $(M_1, \ldots, M_r)$ of elements of $\mathscr{M}^{(|A|)}$ belongs to $R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A})}$ if and only if there is some $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ such that $M_i = Q_{/\pi_i}$ for each $i \in [r]$, where $\pi_i : R^{\mathbf{A}} \to A$ maps $\mathbf{a} \in R^{\mathbf{A}}$ to its i-th coordinate a_i.[6] The definition of free structure may at this point strike the reader as rather technical. We shall see that, if we consider certain quite general classes of minions, this object unveils an interesting geometric description of linear and multilinear nature.

2.3 Relaxations and hierarchies

The following relaxations of (P)CSPs shall appear in this paper: *Arc consistency* (AC) is a propagation algorithm that checks for the existence of assignments satisfying the local constraints of the given (P)CSP instance [89]; the *basic linear programming* (BLP) relaxation looks for compatible probability distributions on assignments [83]; the *affine integer programming* (AIP) relaxation turns the constraints into linear equations, that can be solved over the integers using (a variant of) Gaussian elimination [22]; the *basic semidefinite programming* (SDP) relaxation is essentially a strengthening of BLP, where probabilities are replaced by vectors satisfying orthogonality requirements [95]; the *combined basic linear programming and affine integer programming* (BA) relaxation is a hybrid algorithm blending BLP and AIP [27]. We refer the reader to [12] for a formal description of AC, BLP, and AIP, and to [27] for BA, while SDP shall be formally defined later in this section.

In this work, we shall mainly focus on algorithmic *hierarchies*. The bounded width (BW^k) hierarchy (also known as the *local consistency checking algorithm*) refines AC by propagating local solutions over bigger and bigger portions of the instance, while the Sherali–Adams LP (SA^k), affine integer programming (AIP^k), Sum-of-Squares (SoS^k), and combined basic linear programming and affine integer programming (BA^k) hierarchies strengthen the BLP, AIP, SDP, and BA relaxations, respectively, by looking for compatible distributions of assignments over sub-instances of some fixed size. Below, we give a formal description of the five hierarchies mentioned above, as well as the SDP relaxation (see also the discussion in Appendix A for

[5]Horn-3-SAT is often defined in the literature with one more relation "$x \wedge y \Rightarrow \neg z$" but this is redundant.
[6]Here and henceforth, we are implicitly identifying $R^{\mathbf{A}}$ with $|R^{\mathbf{A}}|$ and A with $|A|$ for readability.

a comparison with different formulations of these algorithms appearing in the literature on (P)CSPs).

$\mathbf{BW^k}$ Given two σ-structures $\mathbf{X}$ and $\mathbf{A}$ and a subset $S \subseteq X$, a *partial homomorphism* from $\mathbf{X}$ to $\mathbf{A}$ with domain S is a homomorphism from $\mathbf{X}[S]$ to $\mathbf{A}$, where $\mathbf{X}[S]$ is the substructure of $\mathbf{X}$ *induced* by S – i.e., it is the σ-structure whose domain is S and, for any $R \in \sigma$, $R^{\mathbf{X}[S]} = R^{\mathbf{X}} \cap S^{\operatorname{ar}(R)}$. We say that the k-th level of the *bounded width algorithm* accepts when applied to $\mathbf{X}$ and $\mathbf{A}$, and we write $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$, if there exists a nonempty collection $\mathcal{F}$ of partial homomorphisms from $\mathbf{X}$ to $\mathbf{A}$ with at most k-element domains such that (i) $\mathcal{F}$ is closed under restrictions, i.e., for every $f \in \mathcal{F}$ and every $V \subseteq \operatorname{dom}(f)$, $f|_V \in \mathcal{F}$, and (ii) $\mathcal{F}$ has the extension property up to k, i.e., for every $f \in \mathcal{F}$ and every $V \subseteq X$ with $|V| \leq k$ and $\operatorname{dom}(f) \subseteq V$, there exists $g \in \mathcal{F}$ such that g extends f and $\operatorname{dom}(g) = V$.

We say that BW^k *solves* a PCSP template $(\mathbf{A}, \mathbf{B})$ if $\mathbf{X} \to \mathbf{B}$ whenever $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. Note that the algorithm is always complete: If $\mathbf{X} \to \mathbf{A}$ then $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$.

$\mathbf{SA^k}$, $\mathbf{AIP^k}$, and $\mathbf{BA^k}$ For $k \in \mathbb{N}$, we say that a σ-structure $\mathbf{A}$ is *k-enhanced* if the signature σ contains a k-ary symbol R_k such that $R_k^{\mathbf{A}} = A^k$. Given two k-enhanced σ-structures $\mathbf{X}, \mathbf{A}$, we introduce a variable $\lambda_{R,\mathbf{x},\mathbf{a}}$ for every $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, $\mathbf{a} \in R^{\mathbf{A}}$. Consider the following system of equations:

$$
\begin{array}{llcll}
(\clubsuit 1) & \sum_{\mathbf{a} \in R^{\mathbf{A}}} \lambda_{R,\mathbf{x},\mathbf{a}} & = & 1 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}} \\
(\clubsuit 2) & \sum_{\mathbf{a} \in R^{\mathbf{A}},\, \mathbf{a}_{\mathbf{i}} = \mathbf{b}} \lambda_{R,\mathbf{x},\mathbf{a}} & = & \lambda_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{b}} & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{i} \in [\operatorname{ar}(R)]^k, \mathbf{b} \in A^k \\
(\clubsuit 3) & \lambda_{R,\mathbf{x},\mathbf{a}} & = & 0 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}, \mathbf{x} \not\prec \mathbf{a}.^{7}
\end{array}
\qquad (\clubsuit)
$$

We say that the k-th level of the *Sherali–Adams linear programming hierarchy* accepts when applied to $\mathbf{X}$ and $\mathbf{A}$, and we write $\mathrm{SA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$, if the system ($\clubsuit$) admits a solution such that all variables take rational nonnegative values.[8] Similarly, we say that the k-th level of the *affine integer programming hierarchy* accepts when applied to $\mathbf{X}$ and $\mathbf{A}$, and we write $\mathrm{AIP}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$, if the system above admits a solution such that all variables take integer values. Moreover, we say that the k-th level of the *combined basic linear programming and affine integer programming hierarchy* accepts when applied to $\mathbf{X}$ and $\mathbf{A}$, and we write $\mathrm{BA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$, if the system above admits both a solution such that all variables take rational nonnegative values and a solution such that all variables take integer values, and the following *refinement condition* holds: Denoting the rational nonnegative and the integer solutions by the superscripts (B) and (A), respectively, we require that $\lambda^{(\mathrm{A})}_{R,\mathbf{x},\mathbf{a}} = 0$ whenever $\lambda^{(\mathrm{B})}_{R,\mathbf{x},\mathbf{a}} = 0$, for each $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}$.

We say that SA^k *solves* a PCSP template $(\mathbf{A}, \mathbf{B})$ if $\mathbf{X} \to \mathbf{B}$ whenever $\mathrm{SA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. The definition for AIP^k and BA^k is analogous. Note that the three algorithms are always complete: If $\mathbf{X} \to \mathbf{A}$ then $\mathrm{SA}^k(\mathbf{X}, \mathbf{A}) = \mathrm{AIP}^k(\mathbf{X}, \mathbf{A}) = \mathrm{BA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$.

[7] The condition $\mathbf{x} \not\prec \mathbf{a}$ formalises the requirement that the same variables (elements of $\mathbf{x}$) should not be assigned different values (elements of $\mathbf{a}$). Some papers avoid this requirement by imposing that (P)CSP instances should have no repetition of variables in the constraint scopes; i.e., elements of $\mathbf{x}$ should all be distinct.

[8] See Appendix A.1 (in particular, Lemma 60) for an equivalent description of SA^k that does not involve the enhancement operation.

SDP Given two σ-structures $\mathbf{X}, \mathbf{A}$, let $\gamma = |X| \cdot |A| + \sum_{R\in\sigma} |R^{\mathbf{X}}| \cdot |R^{\mathbf{A}}|$. We introduce a variable $\boldsymbol{\lambda}_{x,a}$ taking values in $\mathbb{R}^\gamma$ for every $x \in X$, $a \in A$, and a variable $\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}$ taking values in $\mathbb{R}^\gamma$ for every $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}$.[9] Consider the following system of equations:

$$
(\blacklozenge)\quad
\begin{array}{llcll}
(\blacklozenge 1) & \displaystyle\sum_{a\in A} \|\boldsymbol{\lambda}_{x,a}\|^2 & = & 1 & x \in X \\
(\blacklozenge 2) & \boldsymbol{\lambda}_{x,a} \cdot \boldsymbol{\lambda}_{x,a'} & = & 0 & x \in X, a \neq a' \in A \\
(\blacklozenge 3) & \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'} & = & 0 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \neq \mathbf{a}' \in R^{\mathbf{A}} \\
(\blacklozenge 4) & \displaystyle\sum_{\mathbf{a}\in R^{\mathbf{A}},\, a_i = a} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} & = & \boldsymbol{\lambda}_{x_i,a} & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, a \in A, i \in [\mathrm{ar}(R)].
\end{array}
$$

We say that the *standard semidefinite programming relaxation* accepts when applied to $\mathbf{X}$ and $\mathbf{A}$, and we write $\mathrm{SDP}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$, if the system ($\blacklozenge$) admits a solution. We say that SDP *solves* a PCSP template $(\mathbf{A}, \mathbf{B})$ if $\mathbf{X} \to \mathbf{B}$ whenever $\mathrm{SDP}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. Note that the algorithm is always complete: If $\mathbf{X} \to \mathbf{A}$ then $\mathrm{SDP}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$.

$\mathbf{SoS^k}$ Given two k-enhanced σ-structures $\mathbf{X}, \mathbf{A}$, let $\gamma = \sum_{R\in\sigma} |R^{\mathbf{X}}| \cdot |R^{\mathbf{A}}|$. We introduce a variable $\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}$ taking values in $\mathbb{R}^\gamma$ for every $R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}$. Consider the following system of equations:

$$
(\spadesuit)\quad
\begin{array}{llcll}
(\spadesuit 1) & \displaystyle\sum_{\mathbf{a}\in R^{\mathbf{A}}} \|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 & = & 1 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}} \\
(\spadesuit 2) & \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'} & = & 0 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \neq \mathbf{a}' \in R^{\mathbf{A}} \\
(\spadesuit 3) & \displaystyle\sum_{\mathbf{a}\in R^{\mathbf{A}},\, \mathbf{a}_{\mathbf{i}} = \mathbf{b}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} & = & \boldsymbol{\lambda}_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{b}} & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{i} \in [\mathrm{ar}(R)]^k, \mathbf{b} \in A^k \\
(\spadesuit 4) & \|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 & = & 0 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}, \mathbf{x} \not\prec \mathbf{a}.
\end{array}
$$

We say that the k-th level of the *Sum-of-Squares semidefinite programming hierarchy* accepts when applied to $\mathbf{X}$ and $\mathbf{A}$, and we write $\mathrm{SoS}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$, if the system ($\spadesuit$) admits a solution. We say that SoS^k *solves* a PCSP template $(\mathbf{A}, \mathbf{B})$ if $\mathbf{X} \to \mathbf{B}$ whenever $\mathrm{SoS}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. Note that the algorithm is always complete: If $\mathbf{X} \to \mathbf{A}$ then $\mathrm{SoS}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$.

Remark 4. Note that the definitions of the SA^k, AIP^k, BA^k, and SoS^k hierarchies given above require that the structures $\mathbf{X}$ and $\mathbf{A}$ to which the hierarchies are applied should be k-enhanced. This allows expressing marginalisation requirements between partial assignments of possibly different sizes via the same condition expressing marginalisation between constraints and variables appearing in the constraint (namely, condition $\clubsuit$2 in ($\clubsuit$) and condition $\spadesuit$3 in ($\spadesuit$)), applied to the extra complete relation R_k. See, for example, Lemma 60 in Appendix A.1 in the case of the Sherali–Adams hiererchy. In contrast, the definition of BW^k does not explicitly require k-enhancement. In fact, it is clear from that description that, for any two structures $\mathbf{X}$ and $\mathbf{A}$, it holds that $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \mathrm{BW}^k(\tilde{\mathbf{X}}, \tilde{\mathbf{A}})$, where $\tilde{\mathbf{X}}$ (resp. $\tilde{\mathbf{A}}$) is the k-enhanced version of $\mathbf{X}$ (resp. $\mathbf{A}$). Note that this is merely a consequence of the specific formulation we use to define the hierarchies. The reason why we adopt the current formulations is that they are formally closer to the tensorisation construction, as we shall see later in the paper.

[9] We point out that requiring the SDP vectors to be elements of $\mathbb{R}^\gamma$ is equivalent to letting them belong to an arbitrary finite-dimensional real vector space. This is because a real square matrix has a Cholesky decomposition if and only if it has a square Cholesky decomposition (if and only if it is positive semidefinite). The same holds for the SoS^k system ($\spadesuit$).

3 Minion tests

Let $(\mathbf{A}, \mathbf{B})$ be a PCSP template. As discussed in Section 1, any (potentially infinite) structure $\mathbf{T}$ on the same signature as $\mathbf{A}$ and $\mathbf{B}$ can be viewed as a test for the decision problem $\text{PCSP}(\mathbf{A}, \mathbf{B})$: Given an instance $\mathbf{X}$, the test returns YES if $\mathbf{X} \to \mathbf{T}$, and NO otherwise. As the next definition illustrates, minions provide a systematic method to build tests for PCSPs.

Definition 5. Let $\mathscr{M}$ be a minion. The *minion test* $\text{Test}_{\mathscr{M}}$ is the decision problem defined as follows: Given two σ-structures $\mathbf{X}$ and $\mathbf{A}$, return YES if $\mathbf{X} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A})$, and NO otherwise.

If $\mathbf{X}$ is an instance of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ for some template $(\mathbf{A}, \mathbf{B})$, we write $\text{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{YES}$ if $\text{Test}_{\mathscr{M}}$ applied to $\mathbf{X}$ and $\mathbf{A}$ returns YES (i.e., if $\mathbf{X} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A})$), and we write $\text{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{NO}$ otherwise. Note that, in the expression "$\text{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{A})$", $\mathbf{X}$ is the input structure of the PCSP, while $\mathbf{A}$ is the fixed structure from the PCSP template.

Leaving SDP aside for the moment, it turns out that the algebraic structure lying at the core of all relaxations mentioned in Section 2.3, of seemingly different nature, is the same, as all of them are minion tests for specific minions.

Theorem 6 ([12, Theorems 7.4, 7.9, and 7.19],[27, Lemma 5.4])**.** $\text{AC} = \text{Test}_{\mathscr{H}}$, $\text{BLP} = \text{Test}_{\mathscr{Q}_{\text{conv}}}$, $\text{AIP} = \text{Test}_{\mathscr{Z}_{\text{aff}}}$, $\text{BA} = \text{Test}_{\mathscr{M}_{\text{BA}}}$.

One reason why minion tests are an interesting type of tests is that they are always complete.

Proposition 7. *$\text{Test}_{\mathscr{M}}$ is complete for any minion $\mathscr{M}$; i.e., for any $\mathbf{X}$ and $\mathbf{A}$ with $\mathbf{X} \to \mathbf{A}$, we have $\mathbf{X} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A})$.*

This immediately follows from the next lemma, implicitly proved in [12] for the case of function minions (see Remark 27). For completeness, we include below the simple proof (which closely follows the proof in [12]).

Lemma 8 ([12])**.** *Let $\mathscr{M}$ be a minion and let $\mathbf{A}$ be a σ-structure. Then, $\mathbf{A} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A})$.*

Proof. Take a unary element $M \in \mathscr{M}^{(1)}$, and consider the map

$$\begin{aligned} f : A &\to \mathscr{M}^{(n)} \\ a &\mapsto M_{/\rho_a} \end{aligned}$$

where $\rho_a : [1] \to [n] = A$ is defined by $\rho_a(1) = a$. Take $R \in \sigma$ of arity r, and consider a tuple $\mathbf{a} = (a_1, \ldots, a_r) \in R^{\mathbf{A}}$. Let $m = |R^{\mathbf{A}}|$, and consider the function $\pi : [1] \to [m]$ defined by $\pi(1) = \mathbf{a}$. Let $Q = M_{/\pi} \in \mathscr{M}^{(m)}$. For each $i \in [r]$, recall the function $\pi_i : [m] \to [n]$ defined by $\pi_i(\mathbf{b}) = b_i$, where $\mathbf{b} = (b_1, \ldots, b_r) \in R^{\mathbf{A}}$. Observe that $\rho_{a_i} = \pi_i \circ \pi$ for each $i \in [r]$. We obtain

$$\begin{aligned} f(\mathbf{a}) &= (f(a_1), \ldots, f(a_r)) = (M_{/\rho_{a_1}}, \ldots, M_{/\rho_{a_r}}) = (M_{/\pi_1 \circ \pi}, \ldots, M_{/\pi_r \circ \pi}) \\ &= ((M_{/\pi})_{/\pi_1}, \ldots, (M_{/\pi})_{/\pi_r}) = (Q_{/\pi_1}, \ldots, Q_{/\pi_r}) \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A})}, \end{aligned}$$

thus showing that f is a homomorphism from $\mathbf{A}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$. □

A second feature of minion tests is that their soundness can be characterised algebraically, as stated in the next proposition.

Proposition 9. *Let* $\mathscr{M}$ *be a minion and let* $(\mathbf{A}, \mathbf{B})$ *be a* PCSP *template. Then,* $\text{Test}_{\mathscr{M}}$ *solves* $\text{PCSP}(\mathbf{A}, \mathbf{B})$ *if and only if* $\mathscr{M} \to \text{Pol}(\mathbf{A}, \mathbf{B})$*.*

Our proof of Proposition 9 uses a standard compactness argument from [90] (see also [69]), and it follows the lines of [12, Theorem 7.9 and Remark 7.13], where the same result is derived from Kőnig's Lemma in the restricted case of *locally countable minions* – i.e., minions $\mathscr{M}$ having the property that the set $\mathscr{M}^{(L)}$ is countable for any L. Since the minion $\mathscr{S}$ described in Definition 11 is not locally countable, we shall need this stronger version of the result when proving that $\mathscr{S}$ provides an algebraic characterisation of the power of SDP (cf. Theorem 14).

We say that a potentially infinite σ-structure $\mathbf{B}$ is *compact* if, for any potentially infinite σ-structure $\mathbf{A}$, $\mathbf{A} \to \mathbf{B}$ if and only if $\mathbf{A}' \to \mathbf{B}$ for every finite substructure $\mathbf{A}'$ of $\mathbf{A}$. The following result is a direct consequence of the uncountable version of the *compactness theorem of logic* ([69, Theorem 6.1.1]).

Theorem 10 ([69, 90])**.** *Every finite* σ*-structure is compact.*

Proof of Proposition 9. We first observe that the condition $\mathscr{M} \to \text{Pol}(\mathbf{A}, \mathbf{B})$ is equivalent to the condition $\mathbb{F}_{\mathscr{M}}(\mathbf{A}) \to \mathbf{B}$ by [12, Lemma 4.4] (see also [43] for the proof for abstract minions).

Suppose that $\mathbb{F}_{\mathscr{M}}(\mathbf{A}) \to \mathbf{B}$. Given an instance $\mathbf{X}$, if $\text{Test}_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{YES}$ then $\mathbf{X} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A})$, and composing the two homomorphisms yields $\mathbf{X} \to \mathbf{B}$. Hence, $\text{Test}_{\mathscr{M}}$ is sound on the template $(\mathbf{A}, \mathbf{B})$. Since, as noted above, $\text{Test}_{\mathscr{M}}$ is always complete, we deduce that $\text{Test}_{\mathscr{M}}$ solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$.

Conversely, suppose that $\text{Test}_{\mathscr{M}}$ solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$. Let $\mathbf{F}$ be a finite substructure of $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$, and notice that the inclusion map yields a homomorphism from $\mathbf{F}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$. Hence, $\text{Test}_{\mathscr{M}}(\mathbf{F}, \mathbf{A}) = \text{YES}$, so $\mathbf{F} \to \mathbf{B}$. Since $\mathbf{B}$ is compact by Theorem 10, we deduce that $\mathbb{F}_{\mathscr{M}}(\mathbf{A}) \to \mathbf{B}$, as required. □

4 A minion for SDP

The goal of this section is to design a minion $\mathscr{S}$ capturing the power of SDP, thus showing that, similarly to AC, BLP, AIP, and BLP + AIP, also SDP is a minion test. We remark that an equivalent characterisation for the power of SDP was also obtained independently by Brakensiek, Guruswami, and Sandeep in [26].

Definition 11. For $L \in \mathbb{N}$, let $\mathscr{S}^{(L)}$ be the set of real $L \times \aleph_0$ matrices M such that

$$\text{(C1) } \operatorname{csupp}(M) \text{ is finite} \qquad \text{(C2) } MM^T \text{ is a diagonal matrix} \qquad \text{(C3) } \operatorname{tr}(MM^T) = 1. \tag{2}$$

Given a function $\pi : [L] \to [L']$ and a matrix $M \in \mathscr{S}^{(L)}$, we let $M_{/\pi} = PM$, where P is the $L' \times L$ matrix whose (i, j)-th entry is 1 if $\pi(j) = i$, and 0 otherwise. We set $\mathscr{S} = \bigsqcup_{L \in \mathbb{N}} \mathscr{S}^{(L)}$.

First of all, we show that $\mathscr{S}$ is closed with respect to the minor maps described above and, thus, it is indeed a minion.

Proposition 12. $\mathscr{S}$ *is a minion.*

Proof. For $\pi : [L] \to [L']$ and $M \in \mathscr{S}^{(L)}$, we have that $M_{/\pi} = PM \in \mathcal{T}^{L', \aleph_0}(\mathbb{R})$ and $\operatorname{csupp}(PM)$ is finite (where P is the $L' \times L$ matrix associated with π, as per Definition 11).

One easily checks that (i) $P^T\mathbf{1}_{L'} = \mathbf{1}_L$, and (ii) PP^T is a diagonal matrix. Using that both MM^T and PP^T are diagonal, we find that $M_{/\pi}(M_{/\pi})^T = PMM^TP^T$ is diagonal, too. Moreover, since the trace of a diagonal matrix equals the sum of its entries, we obtain

$$\operatorname{tr}(M_{/\pi}(M_{/\pi})^T) = \mathbf{1}_{L'}^T M_{/\pi}(M_{/\pi})^T\mathbf{1}_{L'} = \mathbf{1}_{L'}^T PMM^TP^T\mathbf{1}_{L'} = \mathbf{1}_L^T MM^T\mathbf{1}_L = \operatorname{tr}(MM^T) = 1.$$

It follows that $M_{/\pi} \in \mathscr{S}^{(L')}$. Furthermore, one easily checks that $M_{/\mathrm{id}} = M$ and, given $\tilde{\pi} : [L'] \to [L'']$, $(M_{/\pi})_{/\tilde{\pi}} = M_{/\tilde{\pi}\circ\pi}$, which concludes the proof that $\mathscr{S}$ is a minion. □

In the remaining part of this section, we prove that $\mathscr{S}$ captures the power of the SDP relaxation, as stated below.

Proposition 13. SDP $=$ Test$_{\mathscr{S}}$. *In other words, given two σ-structures $\mathbf{X}$ and $\mathbf{A}$,* SDP$(\mathbf{X},\mathbf{A}) =$ YES *if and only if $\mathbf{X} \to \mathbb{F}_{\mathscr{S}}(\mathbf{A})$.*

Combining Propositions 13 and 9, we immediately obtain the following algebraic characterisation of the power of SDP.

Theorem 14. *Let $(\mathbf{A},\mathbf{B})$ be a* PCSP *template. Then,* SDP *solves* PCSP$(\mathbf{A},\mathbf{B})$ *if and only if $\mathscr{S} \to \operatorname{Pol}(\mathbf{A},\mathbf{B})$.*

For a σ-structure $\mathbf{A}$, a symbol $R \in \sigma$ of arity r, and a number $i \in [r]$, we consider the $|A| \times |R^{\mathbf{A}}|$ matrix P_i whose $(a,\mathbf{a})$-th entry is 1 if $a_i = a$, and 0 otherwise. We shall use the following simple description of the entries of P_i.

Lemma 15 ([43, Section 4.1])**.** *Let $\mathbf{A}$ be a σ-structure, let $R \in \sigma$ of arity r, and let $i \in [r], a \in A$. Then,*

$$\mathbf{e}_a^T P_i = \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_i=a}} \mathbf{e}_{\mathbf{a}}^T.$$

We shall also need the following observation, whose simple proof is deferred to Appendix A.2.

Proposition 16. *Let $\mathbf{X},\mathbf{A}$ be two σ-structures. The system* (♦) *implies the following facts:*

$$(i)\quad \|\sum_{a\in A}\boldsymbol{\lambda}_{x,a}\|^2 = 1 \qquad x\in X;$$

$$(ii)\quad \sum_{\mathbf{a}\in R^{\mathbf{A}}}\|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \|\sum_{\mathbf{a}\in R^{\mathbf{A}}}\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = 1 \qquad R\in\sigma, \mathbf{x}\in R^{\mathbf{X}};$$

$$(iii)\quad \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_i=a,\ a_j=a'}}\|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \boldsymbol{\lambda}_{x_i,a}\cdot\boldsymbol{\lambda}_{x_j,a'} \qquad R\in\sigma, \mathbf{x}\in R^{\mathbf{X}}, a,a'\in A, i,j\in[\operatorname{ar}(R)].$$

If, in addition, $\mathbf{X}$ and $\mathbf{A}$ are 2-enhanced,

$$(iv)\quad \sum_{a\in A}\boldsymbol{\lambda}_{x,a} = \sum_{a\in A}\boldsymbol{\lambda}_{x',a} \qquad x,x'\in X.$$

In order to prove that SDP $=$ Test$_{\mathscr{S}}$, we essentially need to encode the vectors $\boldsymbol{\lambda}$ witnessing an SDP solution as rows of matrices belonging to $\mathscr{S}$. Note that the vectors $\boldsymbol{\lambda}$ live in a vector space having a finite dimension – namely, the number $\gamma = |X|\cdot|A| + \sum_{R\in\sigma}|R^{\mathbf{X}}|\cdot|R^{\mathbf{A}}|$ (cf. the description of SDP in Section 2.3). On the other hand, the matrices in $\mathscr{S}$ have rows of infinite size, living in $\mathbb{R}^{\aleph_0}$. This issue is easily solved by considering a γ-dimensional subspace of $\mathbb{R}^{\aleph_0}$ and working in an orthonormal basis of such subspace, through a standard orthonormalisation argument.

Proof of Proposition 13. Suppose that $\mathrm{SDP}(\mathbf{X},\mathbf{A}) = \textsc{Yes}$, and let the family of vectors $\boldsymbol{\lambda}_{x,a}, \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \in \mathbb{R}^\gamma$ witness it, for $x \in X, a \in A, R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}$, where $\gamma = |X|\cdot|A| + \sum_{R\in\sigma}|R^{\mathbf{X}}|\cdot|R^{\mathbf{A}}|$. Consider the map $\xi : X \to \mathscr{S}^{(|A|)}$ assigning to each $x \in X$ the $|A| \times \aleph_0$ matrix whose rows are the vectors $\boldsymbol{\lambda}_{x,a}$ for each $a \in A$, padded with infinitely many zeroes; i.e.,

$$\mathbf{e}_a^T\xi(x)\mathbf{e}_j = \begin{cases} \mathbf{e}_j^T\boldsymbol{\lambda}_{x,a} & \text{if } j \leq \gamma \\ 0 & \text{otherwise} \end{cases} \qquad x \in X, a \in A, j \in \mathbb{N}.$$

We claim that ξ is well defined. First, $\xi(x)\mathbf{e}_j = \mathbf{0}$ for each $j > \gamma$, so condition (C1) from Definition 11 is satisfied. Given $a, a' \in A$, it holds that $\mathbf{e}_a^T\xi(x)\xi(x)^T\mathbf{e}_{a'} = \boldsymbol{\lambda}_{x,a}\cdot\boldsymbol{\lambda}_{x,a'}$. If $a \neq a'$, this quantity is zero by ◆2, so (C2) is satisfied. Finally,

$$\mathrm{tr}(\xi(x)\xi(x)^T) = \sum_{a\in A}\mathbf{e}_a^T\xi(x)\xi(x)^T\mathbf{e}_a = \sum_{a\in A}\|\boldsymbol{\lambda}_{x,a}\|^2 = 1$$

by ◆1, so (C3) is also satisfied and the claim is true. We now show that ξ yields a homomorphism from $\mathbf{X}$ to $\mathbb{F}_{\mathscr{S}}(\mathbf{A})$. Take $R \in \sigma$ of arity r and $\mathbf{x} = (x_1,\ldots,x_r) \in R^{\mathbf{X}}$. Consider the $|R^{\mathbf{A}}| \times \aleph_0$ matrix Q whose rows are the vectors $\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}$ for each $\mathbf{a} \in R^{\mathbf{A}}$ padded with infinitely many zeroes. Using the same arguments as above, we have that Q satisfies (C1) and $\mathbf{e}_{\mathbf{a}}^TQQ^T\mathbf{e}_{\mathbf{a}'} = \boldsymbol{\lambda}_{R,\mathbf{x}\,\mathbf{a}}\cdot\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'}$, so (C2) follows from ◆3 and (C3) from point (ii) of Proposition 16. Therefore, $Q \in \mathscr{S}^{(|R^{\mathbf{A}}|)}$. We now claim that $\xi(x_i) = Q_{/\pi_i}$ for each $i \in [r]$. Indeed, for $a \in A$ and $j \in [\gamma]$, we have

$$\mathbf{e}_a^T\xi(x_i)\mathbf{e}_j = \mathbf{e}_j^T\boldsymbol{\lambda}_{x_i,a} = \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_i=a}}\mathbf{e}_j^T\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} = \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_i=a}}\mathbf{e}_{\mathbf{a}}^TQ\mathbf{e}_j = \mathbf{e}_a^TP_iQ\mathbf{e}_j = \mathbf{e}_a^TQ_{/\pi_i}\mathbf{e}_j,$$

where the second and fourth equalities follow from ◆4 and Lemma 15, respectively. Also, clearly, $\mathbf{e}_a^T\xi(x_i)\mathbf{e}_j = \mathbf{e}_a^TQ_{/\pi_i}\mathbf{e}_j = 0$ if $j \in \mathbb{N}\setminus[\gamma]$. As a consequence, the claim holds. It follows that $\xi(\mathbf{x}) \in R^{\mathbb{F}_{\mathscr{S}}(\mathbf{A})}$, so that ξ is a homomorphism.

Conversely, let $\xi : \mathbf{X} \to \mathbb{F}_{\mathscr{S}}(\mathbf{A})$ be a homomorphism. For $R \in \sigma$ of arity r and $\mathbf{x} = (x_1,\ldots,x_r) \in R^{\mathbf{X}}$, we can fix a matrix $Q_{R,\mathbf{x}} \in \mathscr{S}^{(|R^{\mathbf{A}}|)}$ satisfying $\xi(x_i) = Q_{R,\mathbf{x}/\pi_i}$ for each $i \in [r]$. Consider the sets $S_1 = \{\xi(x)^T\mathbf{e}_a : x \in X, a \in A\}$ and $S_2 = \{Q_{R,\mathbf{x}}^T\mathbf{e}_{\mathbf{a}} : R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}\}$, and the vector space $\mathcal{U} = \mathrm{span}(S_1 \cup S_2) \subseteq \mathbb{R}^{\aleph_0}$. Observe that $\dim\mathcal{U} \leq |S_1 \cup S_2| \leq |S_1| + |S_2| \leq |X|\cdot|A| + \sum_{R\in\sigma}|R^{\mathbf{X}}|\cdot|R^{\mathbf{A}}| = \gamma$. Consider a vector space $\mathcal{V}$ of dimension γ such that $\mathcal{U} \subseteq \mathcal{V} \subseteq \mathbb{R}^{\aleph_0}$. Using the Gram–Schmidt process,[10] we find a projection matrix $Z \in \mathcal{T}^{\aleph_0,\gamma}(\mathbb{R})$ such that $Z^TZ = I_\gamma$ and $ZZ^T\mathbf{v} = \mathbf{v}$ for any $\mathbf{v} \in \mathcal{V}$. Consider the family of vectors

$$\begin{aligned} \boldsymbol{\lambda}_{x,a} &= Z^T\xi(x)^T\mathbf{e}_a & & x \in X, a \in A, \\ \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} &= Z^TQ_{R,\mathbf{x}}^T\mathbf{e}_{\mathbf{a}} & & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}. \end{aligned} \tag{3}$$

[10]We note that the Gram–Schmidt process also applies to vector spaces of countably infinite dimension.

We claim that (3) witnesses that $\mathrm{SDP}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. To check ♦1, observe that

$$\sum_{a \in A} \|\boldsymbol{\lambda}_{x,a}\|^2 = \sum_{a \in A} \mathbf{e}_a^T \xi(x) Z Z^T \xi(x)^T \mathbf{e}_a = \sum_{a \in A} \mathbf{e}_a^T \xi(x) \xi(x)^T \mathbf{e}_a = \mathrm{tr}(\xi(x)\xi(x)^T) = 1,$$

where the second equality follows from the fact that $\xi(x)^T \mathbf{e}_a \in S_1 \subseteq \mathcal{U} \subseteq \mathcal{V}$ and the fourth from (C3). In a similar way, using (C2), we obtain

$$\boldsymbol{\lambda}_{x,a} \cdot \boldsymbol{\lambda}_{x,a'} = \mathbf{e}_a^T \xi(x) Z Z^T \xi(x)^T \mathbf{e}_{a'} = \mathbf{e}_a^T \xi(x) \xi(x)^T \mathbf{e}_{a'} = 0,$$
$$\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'} = \mathbf{e}_{\mathbf{a}}^T Q_{R,\mathbf{x}} Z Z^T Q_{R,\mathbf{x}}^T \mathbf{e}_{\mathbf{a}'} = \mathbf{e}_{\mathbf{a}}^T Q_{R,\mathbf{x}} Q_{R,\mathbf{x}}^T \mathbf{e}_{\mathbf{a}'} = 0$$

if $a \neq a' \in A$ and $\mathbf{a} \neq \mathbf{a}' \in R^{\mathbf{A}}$. This shows that ♦2 and ♦3 hold. Finally, to prove ♦4, we observe that

$$\sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ a_i = a}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} = \sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ a_i = a}} Z^T Q_{R,\mathbf{x}}^T \mathbf{e}_{\mathbf{a}} = \Big(\sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ a_i = a}} \mathbf{e}_{\mathbf{a}}^T Q_{R,\mathbf{x}} Z \Big)^T = \left(\mathbf{e}_a^T P_i Q_{R,\mathbf{x}} Z\right)^T = \left(\mathbf{e}_a^T \xi(x_i) Z\right)^T$$
$$= Z^T \xi(x_i)^T \mathbf{e}_a = \boldsymbol{\lambda}_{x_i,a},$$

where the third equality follows from Lemma 15. Therefore, the claim is true and the proof is complete. □

5 Hierarchies of minion tests

As we have seen in Section 3, minions give a systematic method for designing tests for (P)CSPs. We now describe a construction, which we call *tensorisation*, that provides a technique to systematically refine minion tests and turning them into algorithmic *hierarchies.*

Let S be a set and let $k \in \mathbb{N}$. Recall that, for a tuple $\mathbf{n} = (n_1, \dots, n_k) \in \mathbb{N}^k$, $\mathcal{T}^{\mathbf{n}}(S)$ denotes the set of all functions from $[n_1] \times \cdots \times [n_k]$ to S, which we visualise as hypermatrices or tensors. Furthermore, given a signature σ, we denote by $\sigma^{\textcircled{k}}$ the signature consisting of the same symbols as σ such that each symbol R of arity r in σ has arity r^k in $\sigma^{\textcircled{k}}$.

Definition 17. The *k-th tensor power* of a σ-structure $\mathbf{A}$ is the $\sigma^{\textcircled{k}}$-structure $\mathbf{A}^{\textcircled{k}}$ having domain A^k and relations defined as follows: For each symbol $R \in \sigma$ of arity r in σ, we set $R^{\mathbf{A}^{\textcircled{k}}} = \left\{\mathbf{a}^{\textcircled{k}} : \mathbf{a} \in R^{\mathbf{A}}\right\}$, where, for $\mathbf{a} \in R^{\mathbf{A}}$, $\mathbf{a}^{\textcircled{k}}$ is the tensor in $\mathcal{T}^{r \cdot \mathbf{1}_k}(A^k)$ defined as follows: For any $\mathbf{i} \in [r]^k$, the $\mathbf{i}$-th element of $\mathbf{a}^{\textcircled{k}}$ is $\mathbf{a}_{\mathbf{i}}$.

In other words, it holds that $E_{\mathbf{i}} * \mathbf{a}^{\textcircled{k}} = \mathbf{a}_{\mathbf{i}}$ for any $\mathbf{i} \in [r]^k$. Note that $\mathbf{a}^{\textcircled{k}}$ can be visualised as the formal Segre outer product of k copies of $\mathbf{a}$ (cf. [87]).

It is easy to check that $\mathbf{A}^{\textcircled{1}} = \mathbf{A}$. Moreover, the function $R^{\mathbf{A}} \to R^{\mathbf{A}^{\textcircled{k}}}$ given by $\mathbf{a} \mapsto \mathbf{a}^{\textcircled{k}}$ is a bijection, so the cardinality of $R^{\mathbf{A}^{\textcircled{k}}}$ equals the cardinality of $R^{\mathbf{A}}$.

Example 18. Let us describe the third tensor power of the 3-clique – i.e., the structure $\mathbf{K}_3^{\textcircled{3}}$. The domain of $\mathbf{K}_3^{\textcircled{3}}$ is $[3]^3$, i.e., the set of tuples of elements in $[3]$ having length 3. Let R be the symbol corresponding to the binary edge relation in $\mathbf{K}_3$, so that

$R^{\mathbf{K}_3} = \{(1,2),(2,1),(2,3),(3,2),(3,1),(1,3)\}$. Then, $R^{\mathbf{K}_3^{③}}$ has arity $2^3 = 8$ and it is a subset of $\mathcal{T}^{(2,2,2)}([3]^3)$. Specifically, $R^{\mathbf{K}_3^{③}} = \{(1,2)^{③},(2,1)^{③},(2,3)^{③},(3,2)^{③},(3,1)^{③},(1,3)^{③}\}$ where, e.g.,

$$(2,3)^{③} = \left[\begin{matrix}(2,2,2) & (2,2,3) \\ (2,3,2) & (2,3,3)\end{matrix}\,\middle|\,\begin{matrix}(3,2,2) & (3,2,3) \\ (3,3,2) & (3,3,3)\end{matrix}\right].$$

Here, the vertical line separates the two 2×2 layers of the $2 \times 2 \times 2$ tensor: The left block contains the layer whose entries have first coordinate 1, while the right block contains the layer whose entries have first coordinate 2.

Recall that a σ-structure $\mathbf{A}$ is *k-enhanced* if σ contains a k-ary symbol R_k such that $R_k^{\mathbf{A}} = A^k$. Observe that any two σ-structures $\mathbf{A}$ and $\mathbf{B}$ are homomorphic if and only if the structures $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ obtained by adding R_k to their signatures are homomorphic (thus, PCSP$(\mathbf{A},\mathbf{B})$ is equivalent to PCSP$(\tilde{\mathbf{A}},\tilde{\mathbf{B}})$). We now give the main definition of this work.

Definition 19. For a minion $\mathscr{M}$ and an integer $k \in \mathbb{N}$, the *k-th level of the minion test* $\text{Test}_{\mathscr{M}}$, denoted by $\text{Test}^k_{\mathscr{M}}$, is the decision problem defined as follows: Given two k-enhanced σ-structures $\mathbf{X}$ and $\mathbf{A}$, return YES if $\mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$, and NO otherwise.

We point out that, in the definition above, it is crucial to require that the structures $\mathbf{X}$ and $\mathbf{A}$ be k-enhanced in order for the test to capture the hierarchies of (P)CSP relaxations studied in the literature, as we shall see later. Such hierarchies involve marginalisation constraints over sets of variables of possibly different sizes (see, for example, the condition (♥2) in the equivalent description (♥) of the Sherali–Adams hierarchy in Appendix A.1). Our goal is then to simulate these constraints in the homomorphism $\mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ of $\sigma^{ⓚ}$-structures witnessing the truth of the test. To that end, we ask that the homomorphism should preserve an extra relation containing *all* possible tuples of variables of length k—which is precisely the relation R_k.

Comparing Definition 19 with Definition 5, we see that $\text{Test}^k_{\mathscr{M}}(\mathbf{X},\mathbf{A}) = \text{Test}_{\mathscr{M}}(\mathbf{X}^{ⓚ},\mathbf{A}^{ⓚ})$. In other words, the k-th level of a minion test is just the minion test applied to the tensor power of the (k-enhanced) structures. We have seen (cf. Proposition 7) that a minion test is always complete. It turns out that this property keeps holding for any level of a minion test.

Proposition 20. *$\text{Test}^k_{\mathscr{M}}$ is complete for any minion $\mathscr{M}$ and any integer $k \in \mathbb{N}$.*

The proof of Proposition 20 relies on the fact that homomorphisms between structures are in some sense invariant under the tensorisation construction, as formally stated in the next proposition. We let $\text{Hom}(\mathbf{A},\mathbf{B})$ denote the set of homomorphisms from $\mathbf{A}$ to $\mathbf{B}$.

Proposition 21. *Let $k \in \mathbb{N}$ and let $\mathbf{A},\mathbf{B}$ be two σ-structures. Then*

(i) $\mathbf{A} \to \mathbf{B}$ *if and only if* $\mathbf{A}^{ⓚ} \to \mathbf{B}^{ⓚ}$;

(ii) *if $\mathbf{A}$ is k-enhanced, there is a bijection* $\rho : \text{Hom}(\mathbf{A},\mathbf{B}) \to \text{Hom}(\mathbf{A}^{ⓚ},\mathbf{B}^{ⓚ})$.

Proof. Let $f : \mathbf{A} \to \mathbf{B}$ be a homomorphism, and consider the function $f^* : A^k \to B^k$ defined by $f^*((a_1,\dots,a_k)) = (f(a_1),\dots,f(a_k))$.[11] Take $R \in \sigma$ of arity r, and consider $\mathbf{a}^{ⓚ} \in R^{\mathbf{A}^{ⓚ}}$,

[11] Note that, throughout the paper, we use the same symbol f to denote both a function $a \mapsto f(a)$ for $a \in A$ and its component-wise application $\mathbf{a} \mapsto f(\mathbf{a}) = (f(a_1),\dots,f(a_p))$ for $\mathbf{a} \in A^p$. Only in this proof, however, it is convenient to introduce the new symbol f^* to specifically denote the component-wise application of f to tuples of length k.

where $\mathbf{a} \in R^{\mathbf{A}}$. Let $\mathbf{b} = f(\mathbf{a})$. Since f is a homomorphism, $\mathbf{b} \in R^{\mathbf{B}}$, so $\mathbf{b}^{ⓚ} \in R^{\mathbf{B}^{ⓚ}}$. For any $\mathbf{i} \in [r]^k$, we have

$$E_{\mathbf{i}} * f^* \left(\mathbf{a}^{ⓚ}\right) = f^* \left(E_{\mathbf{i}} * \mathbf{a}^{ⓚ}\right) = f^*(\mathbf{a}_{\mathbf{i}}) = \mathbf{b}_{\mathbf{i}} = E_{\mathbf{i}} * \mathbf{b}^{ⓚ},$$

which yields $f^*(\mathbf{a}^{ⓚ}) = \mathbf{b}^{ⓚ} \in R^{\mathbf{B}^{ⓚ}}$. Hence, $f^* : \mathbf{A}^{ⓚ} \to \mathbf{B}^{ⓚ}$ is a homomorphism.

Conversely, let $g : \mathbf{A}^{ⓚ} \to \mathbf{B}^{ⓚ}$ be a homomorphism. We define the function $g_* : A \to B$ by setting $g_*(a) = \mathbf{e}_1^T g((a, \dots, a))$ for each $a \in A$. Take $R \in \sigma$ of arity r, and consider a tuple $\mathbf{a} = (a_1, \dots, a_r) \in R^{\mathbf{A}}$. Since $\mathbf{a}^{ⓚ} \in R^{\mathbf{A}^{ⓚ}}$ and g is a homomorphism, we have that $g(\mathbf{a}^{ⓚ}) \in R^{\mathbf{B}^{ⓚ}}$. Therefore, $g(\mathbf{a}^{ⓚ}) = \mathbf{b}^{ⓚ}$ for some $\mathbf{b} = (b_1, \dots, b_r) \in R^{\mathbf{B}}$. For each $j \in [r]$, consider the tuple $\mathbf{i} = (j, \dots, j) \in [r]^k$ and observe that

$$g((a_j, \dots, a_j)) = g(\mathbf{a}_{\mathbf{i}}) = g\left(E_{\mathbf{i}} * \mathbf{a}^{ⓚ}\right) = E_{\mathbf{i}} * g(\mathbf{a}^{ⓚ}) = E_{\mathbf{i}} * \mathbf{b}^{ⓚ} = \mathbf{b}_{\mathbf{i}} = (b_j, \dots, b_j).$$

Hence, we find

$$g_*(\mathbf{a}) = \left(\mathbf{e}_1^T g((a_1, \dots, a_1)), \dots, \mathbf{e}_1^T g((a_r, \dots, a_r))\right) = (b_1, \dots, b_r) = \mathbf{b} \in R^{\mathbf{B}}.$$

Therefore, $g_* : \mathbf{A} \to \mathbf{B}$ is a homomorphism. This concludes the proof of (i).

To prove (ii), observe first that, if $\mathbf{A} \not\to \mathbf{B}$, then $\mathrm{Hom}(\mathbf{A}, \mathbf{B}) = \mathrm{Hom}(\mathbf{A}^{ⓚ}, \mathbf{B}^{ⓚ}) = \emptyset$, so there is a trivial bijection in this case. If $\mathbf{A} \to \mathbf{B}$, consider the map $\rho : \mathrm{Hom}(\mathbf{A}, \mathbf{B}) \to \mathrm{Hom}(\mathbf{A}^{ⓚ}, \mathbf{B}^{ⓚ})$ defined by $f \mapsto f^*$ and the map $\rho' : \mathrm{Hom}(\mathbf{A}^{ⓚ}, \mathbf{B}^{ⓚ}) \to \mathrm{Hom}(\mathbf{A}, \mathbf{B})$ defined by $g \mapsto g_*$. For $f : \mathbf{A} \to \mathbf{B}$ and $a \in A$, we have

$$(f^*)_*(a) = \mathbf{e}_1^T f^*((a, \dots, a)) = \mathbf{e}_1^T (f(a), \dots, f(a)) = f(a)$$

so that $\rho' \circ \rho = \mathrm{id}_{\mathrm{Hom}(\mathbf{A}, \mathbf{B})}$. Consider now $g : \mathbf{A}^{ⓚ} \to \mathbf{B}^{ⓚ}$, and take $\mathbf{a} = (a_1, \dots, a_k) \in A^k$. Using the assumption that $\mathbf{A}$ is k-enhanced, we have $\mathbf{a} \in R_k^{\mathbf{A}}$, which implies $\mathbf{a}^{ⓚ} \in R_k^{\mathbf{A}^{ⓚ}}$. Hence, $g(\mathbf{a}^{ⓚ}) \in R_k^{\mathbf{B}^{ⓚ}}$, so $g(\mathbf{a}^{ⓚ}) = \mathbf{b}^{ⓚ}$ for some $\mathbf{b} = (b_1, \dots, b_k) \in R_k^{\mathbf{B}} \subseteq B^k$. For $j \in [k]$ and $\mathbf{i} = (j, \dots, j) \in [k]^k$, we have

$$g((a_j, \dots, a_j)) = g(\mathbf{a}_{\mathbf{i}}) = g\left(E_{\mathbf{i}} * \mathbf{a}^{ⓚ}\right) = E_{\mathbf{i}} * g\left(\mathbf{a}^{ⓚ}\right) = E_{\mathbf{i}} * \mathbf{b}^{ⓚ} = \mathbf{b}_{\mathbf{i}} = (b_j, \dots, b_j).$$

Letting $\mathbf{i}' = (1, \dots, k) \in [k]^k$, we obtain

$$\begin{aligned}(g_*)^*(\mathbf{a}) &= (g_*(a_1), \dots, g_*(a_k)) = \left(\mathbf{e}_1^T g((a_1, \dots, a_1)), \dots, \mathbf{e}_1^T g((a_k, \dots, a_k))\right) = (b_1, \dots, b_k) \\ &= \mathbf{b} = \mathbf{b}_{\mathbf{i}'} = E_{\mathbf{i}'} * \mathbf{b}^{ⓚ} = E_{\mathbf{i}'} * g\left(\mathbf{a}^{ⓚ}\right) = g\left(E_{\mathbf{i}'} * \mathbf{a}^{ⓚ}\right) = g(\mathbf{a}_{\mathbf{i}'}) = g(\mathbf{a}),\end{aligned}$$

so that $\rho \circ \rho' = \mathrm{id}_{\mathrm{Hom}(\mathbf{A}^{ⓚ}, \mathbf{B}^{ⓚ})}$, which concludes the proof of (ii). □

Remark 22. Part (ii) of Proposition 21 does not hold if we relax the requirement that $\mathbf{A}$ be k-enhanced. More precisely, in this case, the function $\rho : \mathrm{Hom}(\mathbf{A}, \mathbf{B}) \to \mathrm{Hom}(\mathbf{A}^{ⓚ}, \mathbf{B}^{ⓚ})$ defined in the proof of Proposition 21 still needs to be injective, but may not be surjective. Therefore, we have $|\mathrm{Hom}(\mathbf{A}, \mathbf{B})| \leq \left|\mathrm{Hom}(\mathbf{A}^{ⓚ}, \mathbf{B}^{ⓚ})\right|$, and the inequality may be strict.

For example, consider the Boolean structure $\mathbf{A}$ having a single unary relation $R_1^{\mathbf{A}} = A = \{0,1\}$. So, $\mathbf{A}$ is 1-enhanced but not 2-enhanced. Observe that $|\mathrm{Hom}(\mathbf{A},\mathbf{A})| = 4$. The tensorised structure $\mathbf{A}^{②}$ has domain $\{0,1\}^2 = \{(0,0),(0,1),(1,0),(1,1)\}$, and its (unary) relation is $R_1^{\mathbf{A}^{②}} = \{0^{②}, 1^{②}\} = \{(0,0),(1,1)\}$. Therefore, each map $f : \{0,1\}^2 \to \{0,1\}^2$ such that $f((0,0)) \in \{(0,0),(1,1)\}$ and $f((1,1)) \in \{(0,0),(1,1)\}$ yields a proper homomorphism $\mathbf{A}^{②} \to \mathbf{A}^{②}$. It follows that $\left|\mathrm{Hom}(\mathbf{A}^{②},\mathbf{A}^{②})\right| = 64$, so $\mathrm{Hom}(\mathbf{A},\mathbf{A})$ and $\mathrm{Hom}(\mathbf{A}^{②},\mathbf{A}^{②})$ are not in bijection.

It readily follows from Proposition 21 that hierarchies of minion tests are always complete.

Proof of Proposition 20. Let $\mathbf{X}$ and $\mathbf{A}$ be two k-enhanced σ-structures and suppose that $\mathbf{X} \to \mathbf{A}$. Proposition 21 yields $\mathbf{X}^{ⓚ} \to \mathbf{A}^{ⓚ}$, while Lemma 8 yields $\mathbf{A}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$. The composition of the two homomorphisms witnesses that $\mathrm{Test}^k_{\mathscr{M}}(\mathbf{X},\mathbf{A}) = \textsc{Yes}$, as required. □

Before continuing, we now discuss some of the basic algebraic properties of the tensorisation construction of Definition 17.

Remark 23. We now recall the notion of *pp-power* of relational structures, following the presentation in [12]. A *pp-formula* over a signature σ is a formal expression ψ consisting of an existentially quantified conjunction of predicates of the form (i) "$x = y$", or (ii) "$(x_{i_1},\dots,x_{i_r}) \in R$" for some $R \in \sigma$ of arity r, where $x, y, x_{i_1},\dots,x_{i_r}$ are variables. Let f be the number of free (i.e., unquantified) variables in ψ. Given a σ-structure $\mathbf{A}$, the interpretation of ψ in $\mathbf{A}$ is the set $\psi(\mathbf{A}) \subseteq A^f$ containing all tuples $(a_1,\dots,a_f) \in A^f$ that satisfy ψ, where each symbol R appearing in ψ is interpreted in $\mathbf{A}$. Let now $\mathbf{A}'$ be a σ'-structure for some possibly different signature σ', such that $A' = A$. We say that $\mathbf{A}'$ is *pp-definable* from $\mathbf{A}$ if for each symbol $S \in \sigma'$ it holds that $S^{\mathbf{A}'} = \psi_S(\mathbf{A})$ for some pp-formula ψ_S over σ. Suppose now that $A' = A^m$ for some $m \in \mathbb{N}$, and let $\mathrm{vec}_m(\mathbf{A}')$ be the structure with domain A and relations defined as follows: For each r-ary symbol $S \in \sigma'$, $\mathrm{vec}_m(\mathbf{A}')$ has an rm-ary relation containing the tuple $(b_1^{(1)},\dots,b_m^{(1)},\dots,b_1^{(r)},\dots,b_m^{(r)})$ for each tuple $(\mathbf{b}^{(1)},\dots,\mathbf{b}^{(r)}) \in S^{\mathbf{A}'}$. We say that $\mathbf{A}'$ is an m-fold *pp-power* of $\mathbf{A}$ if $\mathrm{vec}_m(\mathbf{A}')$ is pp-definable from $\mathbf{A}$.

It is not hard to verify that the k-th tensor power $\mathbf{A}^{ⓚ}$ of a σ-structure $\mathbf{A}$ is a specific (k-fold) pp-power of $\mathbf{A}$. Indeed, consider the structure $\mathrm{vec}_k(\mathbf{A}^{ⓚ})$, and notice that the relation corresponding to a symbol $R \in \sigma$ (having arity r in σ and arity r^k in $\sigma^{ⓚ}$) has arity $k \cdot r^k$ in $\mathrm{vec}_k(\mathbf{A}^{ⓚ})$. Such relation is easily seen to be pp-definable from $\mathbf{A}$. As an example, suppose for concreteness that $k = 2$ and σ contains a single ternary symbol R. Then, we have

$$R^{\mathbf{A}^{②}} = \left\{ \begin{bmatrix} (a_1,a_1) & (a_1,a_2) & (a_1,a_3) \\ (a_2,a_1) & (a_2,a_2) & (a_2,a_3) \\ (a_3,a_1) & (a_3,a_2) & (a_3,a_3) \end{bmatrix} : (a_1,a_2,a_3) \in R^{\mathbf{A}} \right\} \quad \text{and}$$

$$R^{\mathrm{vec}_2(\mathbf{A}^{②})} = \{(a_1,a_1,a_1,a_2,a_1,a_3,a_2,a_1,a_2,a_2,a_2,a_3,a_3,a_1,a_3,a_2,a_3,a_3) : (a_1,a_2,a_3) \in R^{\mathbf{A}}\}.$$

Clearly, there exists a pp-formula ψ over σ such that the 18-ary relation of $\mathrm{vec}_2(\mathbf{A}^{②})$ (where $18 = 2 \cdot 3^2$) satisfies $R^{\mathrm{vec}_2(\mathbf{A}^{②})} = \psi(\mathbf{A})$. This easily generalises to arbitrary structures and arbitrary powers k. In the general case, the pp-formula ψ_R satisfying $R^{\mathrm{vec}_k(\mathbf{A}^{ⓚ})} = \psi_R(\mathbf{A})$ (for some symbol R of arity r in σ) is

$$\psi_R(x_1,\dots,x_{k\cdot r^k}) = \exists_{z_1,\dots,z_r} \;\; (z_1,\dots,z_r) \in R \;\; \wedge \bigwedge_{\mathbf{i}\in[r]^k} \bigwedge_{j\in[k]} x_{k\cdot(p(\mathbf{i})-1)+j} = z_{i_j},$$

where p is a fixed bijection from $[r]^k$ to $[r^k]$.

It also follows that, given a PCSP template $(\mathbf{A}, \mathbf{B})$, the PCSP template $(\mathbf{A}^{\textcircled{k}}, \mathbf{B}^{\textcircled{k}})$ is a pp-power of $(\mathbf{A}, \mathbf{B})$ in the sense of [12, Definition 4.7]—i.e., $\text{vec}_k(\mathbf{A}^{\textcircled{k}})$ and $\text{vec}_k(\mathbf{B}^{\textcircled{k}})$ are definable from $\mathbf{A}$ and $\mathbf{B}$, respectively, using the *same* pp-formulas. (Note that $(\mathbf{A}^{\textcircled{k}}, \mathbf{B}^{\textcircled{k}})$ is a proper PCSP template by virtue of part (i) of Proposition 21.) It then immediately follows from [12, Corollary 4.10] that there exists a minion homomorphism $\text{Pol}(\mathbf{A}, \mathbf{B}) \to \text{Pol}(\mathbf{A}^{\textcircled{k}}, \mathbf{B}^{\textcircled{k}})$. In fact, it was proved in [40, Theorem 37] that $\text{Pol}(\mathbf{A}, \mathbf{B})$ and $\text{Pol}(\mathbf{A}^{\textcircled{k}}, \mathbf{B}^{\textcircled{k}})$ are actually homomorphically equivalent (and even isomorphic when $\mathbf{A}$ is k-enhanced).

Recall now the description of the free structure of a minion $\mathscr{M}$ given in Section 2.2. In the case that the free structure is applied to the k-th tensor power $\mathbf{A}^{\textcircled{k}}$ of a σ-structure $\mathbf{A}$, it follows from Definition 17 that the domain of $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})$ is the set $\mathscr{M}^{(n^k)}$ (where, as usual, we are denoting $n = |A|$). Furthermore, given a relation symbol $R \in \sigma$ whose arity in σ is r, a tuple $(M_{\mathbf{i}})_{\mathbf{i} \in [r]^k}$ of elements of $\mathscr{M}^{(n^k)}$ belongs to $R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})}$ if and only if there is some $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ such that $M_{\mathbf{i}} = Q_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$, where $\pi_{\mathbf{i}} : R^{\mathbf{A}} \to A^k$ maps $\mathbf{a} \in R^{\mathbf{A}}$ to its $\mathbf{i}$-th projection $\mathbf{a}_{\mathbf{i}}$.

It is well known that each of the hierarchies of relaxations described in Section 2.3 has the property that higher levels are at least as powerful as lower levels. As the next result shows, this is in fact a property of all hierarchies of minion tests.

Proposition 24. *Let $\mathscr{M}$ be a minion, let $k, p \in \mathbb{N}$ be such that $k > p$, and let $\mathbf{X}, \mathbf{A}$ be two k- and p-enhanced σ-structures. If $\text{Test}^k_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{Yes}$ then $\text{Test}^p_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{Yes}$.*

Proof. Let $\xi : \mathbf{X}^{\textcircled{k}} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})$ be a homomorphism witnessing that $\text{Test}^k_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{Yes}$. Since $k > p$, we can choose two tuples $\mathbf{v} \in [k]^p$ and $\mathbf{w} \in [p]^k$ such that $\mathbf{w}_{\mathbf{v}} = (1, \dots, p)$. (For instance, we may take $\mathbf{v} = (1, \dots, p)$ and $\mathbf{w} = (1, \dots, p, \dots, p)$.) Consider the map $\tau : A^k \to A^p$ defined by $\mathbf{a} \mapsto \mathbf{a}_{\mathbf{v}}$. We claim that the map $\vartheta : X^p \to \mathscr{M}^{(n^p)}$ defined by $\mathbf{x} \mapsto \xi(\mathbf{x}_{\mathbf{w}})_{/\tau}$ yields a homomorphism from $\mathbf{X}^{\textcircled{p}}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{p}})$, thus witnessing that $\text{Test}^p_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) = \text{Yes}$. To that end, for $R \in \sigma$, take $\mathbf{x} \in R^{\mathbf{X}}$ and observe that, since ξ is a homomorphism and $\mathbf{x}^{\textcircled{k}} \in R^{\mathbf{X}^{\textcircled{k}}}$, $\xi(\mathbf{x}^{\textcircled{k}}) \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})}$. Therefore, there exists $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ satisfying $\xi(\mathbf{x}_{\mathbf{i}}) = Q_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$. If we manage to show that $\vartheta(\mathbf{x}_{\mathbf{j}}) = Q_{/\pi_{\mathbf{j}}}$ for each $\mathbf{j} \in [r]^p$, we would deduce that $\vartheta(\mathbf{x}^{\textcircled{p}}) \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{p}})}$, thus proving the claim. Observe that

$$\vartheta(\mathbf{x}_{\mathbf{j}}) = \xi(\mathbf{x}_{\mathbf{j}_{\mathbf{w}}})_{/\tau} = (Q_{/\pi_{\mathbf{j}_{\mathbf{w}}}})_{/\tau} = Q_{/\tau \circ \pi_{\mathbf{j}_{\mathbf{w}}}},$$

so we are left to show that $\tau \circ \pi_{\mathbf{j}_{\mathbf{w}}} = \pi_{\mathbf{j}}$. Indeed, given any $\mathbf{a} \in R^{\mathbf{A}}$,

$$(\tau \circ \pi_{\mathbf{j}_{\mathbf{w}}})(\mathbf{a}) = \tau(\pi_{\mathbf{j}_{\mathbf{w}}}(\mathbf{a})) = \tau(\mathbf{a}_{\mathbf{j}_{\mathbf{w}}}) = \mathbf{a}_{\mathbf{j}_{\mathbf{w}_{\mathbf{v}}}} = \mathbf{a}_{\mathbf{j}} = \pi_{\mathbf{j}}(\mathbf{a}),$$

as required. □

It follows from Proposition 24 that, if some level of a minion test is sound for a template $(\mathbf{A}, \mathbf{B})$ (equivalently, if it solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$), then any higher level is sound for $(\mathbf{A}, \mathbf{B})$ (equivalently, it solves $\text{PCSP}(\mathbf{A}, \mathbf{B})$).

The next theorem shows that the framework defined above is general enough to capture each of the five hierarchies for (P)CSPs described in Section 2.3.

Theorem 25 (informal). *If $k \in \mathbb{N}$ is at least the maximum arity of the template,*

- $\text{BW}^k = \text{Test}^k_{\mathscr{H}}$
- $\text{SA}^k = \text{Test}^k_{\mathscr{Q}_{\text{conv}}}$
- $\text{AIP}^k = \text{Test}^k_{\mathscr{Z}_{\text{aff}}}$
- $\text{SoS}^k = \text{Test}^k_{\mathscr{S}}$
- $\text{BA}^k = \text{Test}^k_{\mathscr{M}_{\text{BA}}}$.

We will prove Theorem 25 in Section 9. First, it will be convenient to focus on hierarchies of tests corresponding to minions having specific characteristics—which we call linear and conic.

6 Linear minions

By the Definition 19 of a hierarchy of minion tests, $\text{Test}^k_{\mathscr{M}}$ applied to an instance $\mathbf{X}$ of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ checks for the existence of a homomorphism from $\mathbf{X}^{\textcircled{k}}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})$. Therefore, to describe the hierarchy and get knowledge on its functioning it is necessary to study the structure $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})$. Certain features of the hierarchies of minion tests – in particular, that they are complete (Proposition 20) and their power does not decrease while the level increases (Proposition 24) – hold true for any minion, as they only depend on basic properties of the tensorisation construction. In order to prove Theorem 25, however, it is necessary to dig deeper by investigating how the tensorisation construction interacts with the free structure. To that end, we isolate a property shared by all minions mentioned in this work: Their objects can be interpreted as matrices, and their minor operations can be expressed as matrix multiplications. We call such minions *linear*.

Definition 26. A minion $\mathscr{M}$ is *linear* if there exists a semiring $\mathcal{S}$ with additive identity $0_{\mathcal{S}}$ and multiplicative identity $1_{\mathcal{S}}$ and a number $d \in \mathbb{N} \cup \{\aleph_0\}$ (called *depth*) such that

1. the elements of $\mathscr{M}^{(L)}$ are $L \times d$ matrices whose entries belong to $\mathcal{S}$, for each $L \in \mathbb{N}$;
2. given $L, L' \in \mathbb{N}$, $\pi : [L] \to [L']$, and $M \in \mathscr{M}^{(L)}$, $M_{/\pi} = PM$, where P is the $L' \times L$ matrix such that, for $i \in [L']$ and $j \in [L]$, the (i, j)-th entry of P is $1_{\mathcal{S}}$ if $\pi(j) = i$, and $0_{\mathcal{S}}$ otherwise.

Observe that pre-multiplying a matrix M by P amounts to performing a combination of the following three elementary operations to the rows of M: swapping two rows, replacing two rows with their sum, and inserting a zero row. Hence, we may equivalently define a linear minion as a collection of matrices over $\mathcal{S}$ that is closed under such elementary operations.

Remark 27. We now show that any linear minion can be naturally interpreted as a minion of functions. Given two (potentially infinite) sets A, B and an integer $L \in \mathbb{N}$, let $\mathscr{F}^{(L)}_{A,B}$ be the set of all functions $f : A^L \to B$. For $\pi : [L] \to [L']$, we define $f_{/\pi} \in \mathscr{F}^{(L')}_{A,B}$ to be the function given by

$$(a_1, \ldots, a_{L'}) \mapsto f(a_{\pi(1)}, \ldots, a_{\pi(L)}).$$

It is easy to verify that the disjoint union $\mathscr{F}_{A,B} = \bigsqcup_{L\in\mathbb{N}} \mathscr{F}_{A,B}^{(L)}$ equipped with such minor operations is a minion. We let a *function minion* over A, B be any subminion of $\mathscr{F}_{A,B}$ (i.e., a non-empty subset of $\mathscr{F}_{A,B}$ that is closed under the minor operations) [12]. For example, the polymorphism minion $\text{Pol}(\mathbf{A}, \mathbf{B})$ of a PCSP template $(\mathbf{A}, \mathbf{B})$ described in Example 2 is a function minion over the sets A, B. Now, given a linear minion $\mathscr{M}$ of depth d over a semiring $\mathcal{S}$, we can naturally see $\mathscr{M}$ as a function minion over the sets $\mathcal{S}, \mathcal{S}^d$ as follows. Consider the map $\xi : \mathscr{M} \to \mathscr{F}_{\mathcal{S},\mathcal{S}^d}$ that associates with a matrix $M \in \mathscr{M}^{(L)}$ the linear operator $\xi(M)$ from $\mathcal{S}^L$ to $\mathcal{S}^d$ corresponding to the matrix M^T. It is not hard to verify that ξ preserves arities and minors, and it is thus a minion homomorphism. Indeed, given a map $\pi : [L] \to [L']$ and a tuple $\mathbf{s} \in \mathcal{S}^{L'}$, and letting P be the $L' \times L$ Boolean matrix associated with π as per Definition 26, we have

$$\xi(M_{/\pi})(\mathbf{s}) = (M_{/\pi})^T\mathbf{s} = M^T P^T \mathbf{s} = \xi(M)(P^T\mathbf{s}) = \xi(M)_{/\pi}(\mathbf{s}).$$

Furthermore, ξ is injective, so it induces a minion isomorphism from $\mathscr{M}$ to a subminion of $\mathscr{F}_{\mathcal{S},\mathcal{S}^d}$. As a consequence, ξ witnesses that $\mathscr{M}$ is isomorphic to a function minion.

As illustrated by the next proposition, the family of linear minions is rich enough to include the minions associated with all minion tests studied in the literature on PCSPs, including SDP.

Proposition 28. *The following minions are linear:*

- $\mathscr{H}$, *with* $\mathcal{S} = (\{0,1\}, \vee, \wedge)$ *and* $d = 1$
- $\mathscr{Q}_{\text{conv}}$, *with* $\mathcal{S} = \mathbb{Q}$ *and* $d = 1$
- $\mathscr{Z}_{\text{aff}}$, *with* $\mathcal{S} = \mathbb{Z}$ *and* $d = 1$
- $\mathscr{S}$, *with* $\mathcal{S} = \mathbb{R}$ *and* $d = \aleph_0$
- $\mathscr{M}_{\text{BA}}$, *with* $\mathcal{S} = \mathbb{Q}$ *and* $d = 2$.

Proof. The result for $\mathscr{Q}_{\text{conv}}$, $\mathscr{Z}_{\text{aff}}$, and $\mathscr{M}_{\text{BA}}$ directly follows from their definitions in Example 3, while the result for $\mathscr{S}$ is clear from Definition 11.

We now turn to $\mathscr{H}$. Recall that, in Example 3, we described $\mathscr{H}$ as a set of functions rather than a set of matrices. Hence, formally, we prove here that $\mathscr{H}$ is *isomorphic* to a linear minion. Given $L \in \mathbb{N}$ and $\emptyset \neq Z \subseteq [L]$, we identify the Boolean function $f_Z = \bigwedge_{z\in Z} x_z \in \mathscr{H}^{(L)}$ with the indicator vector $\mathbf{v}_Z \in \{0,1\}^L$ whose i-th entry, for $i \in [L]$, is 1 if $i \in Z$, and 0 otherwise. To conclude, we need to show that, under this identification, the minor operations of $\mathscr{H}$ correspond to the minor operations given in Definition 26. In other words, we claim that the function $f_{Z/\pi}$ corresponds to the vector $P\mathbf{v}_Z$ for any $L' \in \mathbb{N}$ and any $\pi : [L] \to [L']$, where P is the $L' \times L$ matrix described in Definition 26. First, observe that

$$f_{Z/\pi}(x_1, \ldots, x_{L'}) = f_Z(x_{\pi(1)}, \ldots, x_{\pi(L)}) = \bigwedge_{z\in Z} x_{\pi(z)} = \bigwedge_{t\in\pi(Z)} x_t = f_{\pi(Z)}(x_1, \ldots, x_{L'}),$$

so $f_{Z/\pi} = f_{\pi(Z)}$. To conclude, we need to show that $P\mathbf{v}_Z = \mathbf{v}_{\pi(Z)}$, where $\mathbf{v}_{\pi(Z)}$ is the indicator vector of the nonempty set $\pi(Z) \subseteq [L']$. Notice that the matrix multiplication is performed in the semiring $\mathcal{S} = (\{0,1\}, \vee, \wedge)$. For any $i \in [L']$, we have

$$\mathbf{e}_i^T P\mathbf{v}_Z = \bigvee_{j\in[L]} \left((\mathbf{e}_i^T P\mathbf{e}_j) \wedge (\mathbf{e}_j^T \mathbf{v}_Z)\right) = \bigvee_{\substack{j\in Z \\ \pi(j)=i}} 1 = \bigvee_{i\in\pi(Z)} 1 = \mathbf{e}_i^T \mathbf{v}_{\pi(Z)},$$

as required. □

As a consequence, the machinery we build in this section (and in Section 7, where we consider an even more specialised minion class) shall be crucial to show that the framework of hierarchies of minion tests captures all hierarchies of relaxations in Theorem 25.

Recall that, as per Definition 5, the minion test associated with a minion $\mathscr{M}$ works by checking whether a given instance is homomorphic to the free structure of $\mathscr{M}$; in other words, $\text{Test}_{\mathscr{M}}$ for a template $(\mathbf{A}, \mathbf{B})$ is $\text{CSP}(\mathbb{F}_{\mathscr{M}}(\mathbf{A}))$. It is then worth checking what the latter object looks like in the case that $\mathscr{M}$ is linear. The next remark shows that, in this case, $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$ has a simple matrix-theoretic description.

Remark 29. Given a linear minion $\mathscr{M}$ with semiring $\mathcal{S}$ and depth d, and a σ-structure $\mathbf{A}$, the free structure $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$ of $\mathscr{M}$ generated by $\mathbf{A}$ has the following description:

- The elements of its domain $\mathscr{M}^{(|A|)}$ are $|A| \times d$ matrices having entries in $\mathcal{S}$.
- For $R \in \sigma$ of arity r, the elements of $R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A})}$ are tuples of the form $(P_1 Q, \dots, P_r Q)$, where $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ is a $|R^{\mathbf{A}}| \times d$ matrix having entries in $\mathcal{S}$ and, for $i \in [r]$, P_i is the $|A| \times |R^{\mathbf{A}}|$ matrix whose $(a, \mathbf{a})$-th entry is $1_{\mathcal{S}}$ if $a_i = a$, and $0_{\mathcal{S}}$ otherwise.

6.1 Multilinear tests

We say that a test is *multilinear* if it can be expressed as $\text{Test}^k_{\mathscr{M}}$ for some linear minion $\mathscr{M}$ and some integer k. In the same way as, for a template $(\mathbf{A}, \mathbf{B})$, $\text{Test}_{\mathscr{M}}$ is $\text{CSP}(\mathbb{F}_{\mathscr{M}}(\mathbf{A}))$, it follows from Definition 19 that $\text{Test}^k_{\mathscr{M}}$ corresponds to $\text{CSP}(\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ}))$, as it checks for the existence of a homomorphism between the tensor power of the instance and the free structure of $\mathscr{M}$ generated by the tensor power of $\mathbf{A}$. (However, recall that $\text{Test}^k_{\mathscr{M}}$ requires that $\mathbf{X}$ and $\mathbf{A}$ be k-enhanced, as per Definition 19.) As we have seen in Remark 29, when $\mathscr{M}$ is linear, the structure $\mathbb{F}_{\mathscr{M}}(\mathbf{A})$ consists in a space of matrices with relations defined through specific matrix products. Similarly, we now show that $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ is a space of tensors, endowed with relations that can be described through the tensor contraction operation.

Given a semiring $\mathcal{S}$, a symbol $R \in \sigma$ of arity r, and a tuple $\mathbf{i} \in [r]^k$, consider the tensor $P_{\mathbf{i}} \in \mathcal{T}^{n \cdot \mathbf{1}_k, |R^{\mathbf{A}}|}(\mathcal{S})$ defined by

$$E_{\mathbf{a}} * P_{\mathbf{i}} * E_{\mathbf{a}'} = \begin{cases} 1 & \text{if } \mathbf{a}'_{\mathbf{i}} = \mathbf{a} \\ 0 & \text{otherwise} \end{cases} \qquad \forall \mathbf{a} \in A^k, \mathbf{a}' \in R^{\mathbf{A}}. \tag{4}$$

Observe that the tensor $P_{\mathbf{i}}$ is the multilinear equivalent of the matrix P_i from Remark 29. We point out that, just like P_i, the tensor $P_{\mathbf{i}}$ depends on the symbol R. We leave this dependence implicit to avoid introducing additional notation; the symbol R shall always be clear from the context. We also observe that, in the expression (4), the tuples $\mathbf{a}$ and $\mathbf{a}'$ have different roles as coordinates of $P_{\mathbf{i}}$: The former is a tuple of k-many coordinates in $[n] = A$, while the latter is a single coordinate in $[|R^{\mathbf{A}}|]$.

Let $\mathscr{M}$ be a linear minion with semiring $\mathcal{S}$ and depth d. The domain of $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ is $\mathscr{M}^{(n^k)}$, which we visualise as a subset of $\mathcal{T}^{n \cdot \mathbf{1}_k, d}(\mathcal{S})$. Given a symbol $R \in \sigma$ of arity r, consider a block tensor $M = (M_{\mathbf{i}})_{\mathbf{i} \in [r]^k} \in \mathcal{T}^{r \cdot \mathbf{1}_k}(\mathcal{T}^{n \cdot \mathbf{1}_k, d}(\mathcal{S})) = \mathcal{T}^{rn \cdot \mathbf{1}_k, d}(\mathcal{S})$. From the definition of free structure, we have that $M \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$ if and only if there exists $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ such that $M_{\mathbf{i}} = Q_{/\pi_{\mathbf{i}}} = P_{\mathbf{i}} \overset{1}{*} Q$ for each $\mathbf{i} \in [r]^k$.

To give a first glance of this object, we illustrate below the structure of $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ in the case that $\mathscr{M} = \mathscr{Q}_{\text{conv}}$, $k = 3$, and $\mathbf{A} = \mathbf{K}_3$.

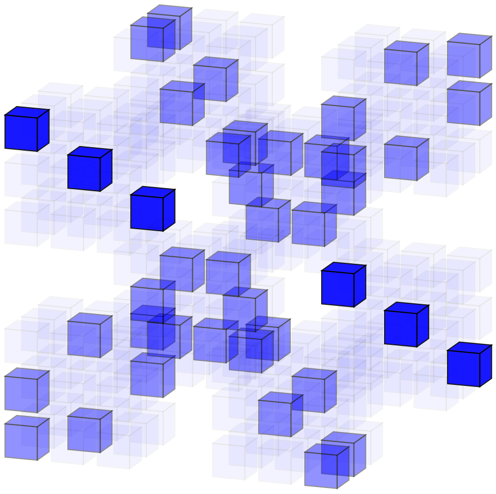

Figure 1: A tensor $M \in R^{\mathbf{F}}$ from Example 30, corresponding to the uniform distribution on the set of edges of $\mathbf{K}_3$. The opacity of a cell is proportional to the value of the corresponding entry: $= \frac{1}{3}$, $= \frac{1}{6}$, $= 0$.

Example 30. Let us denote $\mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{K}_3^{③})$ by $\mathbf{F}$. The domain of $\mathbf{F}$ is the set of nonnegative tensors in $\mathcal{T}^{3\cdot\mathbf{1}_3}(\mathbb{Q})$ whose entries sum up to 1. The relation $R^{\mathbf{F}}$ is the set of those tensors $M \in \mathcal{T}^{2\cdot\mathbf{1}_3}(\mathcal{T}^{3\cdot\mathbf{1}_3}(\mathbb{Q})) = \mathcal{T}^{6\cdot\mathbf{1}_3}(\mathbb{Q})$ such that there exists a stochastic vector $\mathbf{q} = (q_1, \dots, q_6) \in \mathscr{Q}_{\text{conv}}{}^{(6)}$ (which should be interpreted as a probability distribution over the elements of $R^{\mathbf{K}_3}$, i.e., over the directed edges in $\mathbf{K}_3$) for which the $\mathbf{i}$-th block $M_{\mathbf{i}}$ of M satisfies $M_{\mathbf{i}} = \mathbf{q}_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [2]^3$. The $(1,1,1)$-th and the $(2,1,2)$-th entries of M are given below:

$$M_{(1,1,1)} = \left[\begin{array}{ccc|ccc|ccc} q_1+q_6 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & q_2+q_3 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q_4+q_5 \end{array}\right],$$

$$M_{(2,1,2)} = \left[\begin{array}{ccc|ccc|ccc} 0 & 0 & 0 & 0 & q_1 & 0 & 0 & 0 & q_6 \\ q_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q_3 \\ q_5 & 0 & 0 & 0 & q_4 & 0 & 0 & 0 & 0 \end{array}\right].$$

Figure 1 illustrates the tensor $M \in R^{\mathbf{F}}$ corresponding to the uniform distribution $\mathbf{q} = \frac{1}{6} \cdot \mathbf{1}_6$.

We now show that the entries of $P_{\mathbf{i}}$ satisfy the following simple equality, analogous to Lemma 15.

Lemma 31. *Let $k \in \mathbb{N}$, let $\mathbf{A}$ be a σ-structure, let $R \in \sigma$ of arity r, and consider the tuples $\mathbf{a} \in A^k$ and $\mathbf{i} \in [r]^k$. Then*

$$E_{\mathbf{a}} * P_{\mathbf{i}} = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a}}} E_{\mathbf{b}}.$$

Proof. For any $\mathbf{a}' \in R^{\mathbf{A}}$, we have

$$\left(\sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a}}} E_{\mathbf{b}} \right) * E_{\mathbf{a}'} = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a}}} (E_{\mathbf{b}} * E_{\mathbf{a}'}) = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a} \\ \mathbf{b} = \mathbf{a}'}} 1 = \begin{cases} 1 & \text{if } \mathbf{a}'_{\mathbf{i}} = \mathbf{a} \\ 0 & \text{otherwise} \end{cases} = (E_{\mathbf{a}} * P_{\mathbf{i}}) * E_{\mathbf{a}'},$$

from which the result follows. $\square$

The next lemma shows that certain entries of a tensor in the relation $R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$ (i.e., the interpretation of R in the free structure of $\mathscr{M}$ generated by $\mathbf{A}^{ⓚ}$) need to be zero.

Lemma 32. *Let $\mathscr{M}$ be a linear minion of depth d, let $k \in \mathbb{N}$, let $\mathbf{A}$ be a σ-structure, let $R \in \sigma$ of arity r, and suppose $M = (M_{\mathbf{i}})_{\mathbf{i}\in[r]^k} \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$. Then $E_{\mathbf{a}} * M_{\mathbf{i}} = \mathbf{0}_d$ for any $\mathbf{i} \in [r]^k$, $\mathbf{a} \in A^k$ such that $\mathbf{i} \not\prec \mathbf{a}$.*[12]

Proof. Observe that there exists $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ such that $M_{\mathbf{i}} = Q_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$. Using Lemma 31, we obtain

$$E_{\mathbf{a}} * M_{\mathbf{i}} = E_{\mathbf{a}} * Q_{/\pi_{\mathbf{i}}} = E_{\mathbf{a}} * (P_{\mathbf{i}} \overset{1}{*} Q) = (E_{\mathbf{a}} * P_{\mathbf{i}}) * Q = \sum_{\substack{\mathbf{b}\in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}}=\mathbf{a}}} E_{\mathbf{b}} * Q = \sum_{\mathbf{b}\in\emptyset} E_{\mathbf{b}} * Q = \mathbf{0}_d,$$

where the fifth equality follows from the fact that $\mathbf{b}_{\mathbf{i}} = \mathbf{a}$ implies $\mathbf{i} \prec \mathbf{a}$; indeed, in that case, $i_\alpha = i_\beta$ implies $a_\alpha = b_{i_\alpha} = b_{i_\beta} = a_\beta$. $\square$

It follows from Lemma 32 and the previous discussion that, if $\mathscr{M}$ is linear, $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ can be visualised as a space of sparse and highly symmetric tensors, whose nonzero entries form regular patterns. This feature becomes more evident for higher values of the level k. In turn, the structure of $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ is reflected in the properties of the homomorphisms ξ from $\mathbf{X}^{ⓚ}$ to it – which, by virtue of Definition 19, are precisely the solutions sought by $\text{Test}^k_{\mathscr{M}}$. Next, we highlight certain features of such homomorphisms that will be used to prove Theorem 25 in Section 9.

Lemma 33. *Let $\mathscr{M}$ be a linear minion, let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two k-enhanced σ-structures, and let $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ be a homomorphism. Then $E_{\mathbf{a}} * \xi(\mathbf{x}) = \mathbf{0}_d$ for any $\mathbf{x} \in X^k$, $\mathbf{a} \in A^k$ such that $\mathbf{x} \not\prec \mathbf{a}$.*

Proof. From $\mathbf{x} \in X^k = R_k^{\mathbf{X}}$, we derive $\mathbf{x}^{ⓚ} \in R_k^{\mathbf{X}^{ⓚ}}$; since ξ is a homomorphism, this yields $\xi(\mathbf{x}^{ⓚ}) \in R_k^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$. Writing $\xi(\mathbf{x}^{ⓚ})$ in block form as $\xi(\mathbf{x}^{ⓚ}) = (\xi(\mathbf{x}_{\mathbf{i}}))_{\mathbf{i}\in[k]^k}$ and applying Lemma 32, we obtain $E_{\mathbf{a}} * \xi(\mathbf{x}_{\mathbf{i}}) = \mathbf{0}_d$ for any $\mathbf{i} \in [k]^k$ such that $\mathbf{i} \not\prec \mathbf{a}$. Write $\mathbf{x} = (x_1, \ldots, x_k)$ and $\mathbf{a} = (a_1, \ldots, a_k)$. Since $\mathbf{x} \not\prec \mathbf{a}$, there exist $\alpha, \beta \in [k]$ such that $x_\alpha = x_\beta$ and $a_\alpha \neq a_\beta$. Let $\mathbf{i}' \in [k]^k$ be the tuple obtained from $(1, \ldots, k)$ by replacing the β-th entry with α. Observe that $\mathbf{x}_{\mathbf{i}'} = \mathbf{x}$ and $\mathbf{i}' \not\prec \mathbf{a}$. Hence,

$$\mathbf{0}_d = E_{\mathbf{a}} * \xi(\mathbf{x}_{\mathbf{i}'}) = E_{\mathbf{a}} * \xi(\mathbf{x}),$$

as required. $\square$

Using Lemma 33, we obtain some more information on the image of a homomorphism from $\mathbf{X}^{ⓚ}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$. Given a signature σ, we let $\text{armax}(\sigma)$ denote the maximum arity of a relation symbol in σ.

[12] Recall that $\mathbf{i} \prec \mathbf{a}$ means that, for each $\alpha, \beta \in [k]$, $i_\alpha = i_\beta$ implies $a_\alpha = a_\beta$.

Lemma 34. *Let $\mathscr{M}$ be a linear minion, let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two k-enhanced σ-structures such that $k \geq \operatorname{armax}(\sigma)$, and let $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ be a homomorphism. For $R \in \sigma$ of arity r, let $\mathbf{x} \in R^{\mathbf{X}}$ and $\mathbf{a} \in R^{\mathbf{A}}$ be such that $\mathbf{x} \not\prec \mathbf{a}$. Let $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ be such that $Q_{/\pi_{\mathbf{i}}} = \xi(\mathbf{x_i})$ for each $\mathbf{i} \in [r]^k$. Then $E_{\mathbf{a}} * Q = \mathbf{0}_d$.*

Proof. Write $\mathbf{x} = (x_1, \dots, x_r)$ and $\mathbf{a} = (a_1, \dots, a_r)$. Since $\mathbf{x} \not\prec \mathbf{a}$, there exist $\alpha, \beta \in [r]$ such that $x_\alpha = x_\beta$ and $a_\alpha \neq a_\beta$. Using that $k \geq r$, we can take the tuple $\mathbf{i} = (1, 2, \dots, r, r, \dots, r) \in [r]^k$. Consider $\mathbf{x_i} \in X^k$, $\mathbf{a_i} \in A^k$. Notice that $x_{i_\alpha} = x_\alpha = x_\beta = x_{i_\beta}$ and $a_{i_\alpha} = a_\alpha \neq a_\beta = a_{i_\beta}$, so $\mathbf{x_i} \not\prec \mathbf{a_i}$. Applying Lemma 33, we find

$$\mathbf{0}_d = E_{\mathbf{a_i}} * \xi(\mathbf{x_i}) = E_{\mathbf{a_i}} * Q_{/\pi_{\mathbf{i}}} = E_{\mathbf{a_i}} * P_{\mathbf{i}} * Q.$$

Using Lemma 31, we conclude that

$$\mathbf{0}_d = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b_i} = \mathbf{a_i}}} E_{\mathbf{b}} * Q = E_{\mathbf{a}} * Q,$$

where the last equality follows from the fact that $\mathbf{b_i} = \mathbf{a_i}$ if and only if $\mathbf{b} = \mathbf{a}$. □

If $\mathbf{X}$ and $\mathbf{A}$ are k-enhanced σ-structures, any homomorphism ξ from $\mathbf{X}^{ⓚ}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ must satisfy certain symmetries that ultimately depend on the fact that ξ preserves R_k. As shown below in Proposition 37, these symmetries can be concisely expressed through a list of tensor equations. Given a tuple $\mathbf{i} \in [k]^k$, we let $\Pi_{\mathbf{i}} \in \mathcal{T}^{n \cdot \mathbf{1}_{2k}}(\mathcal{S})$ be the tensor defined by

$$E_{\mathbf{a}} * \Pi_{\mathbf{i}} * E_{\mathbf{a}'} = \begin{cases} 1 & \text{if } \mathbf{a}'_{\mathbf{i}} = \mathbf{a} \\ 0 & \text{otherwise} \end{cases} \qquad \forall \mathbf{a}, \mathbf{a}' \in A^k. \tag{5}$$

Observe that, unlike for the tensor $P_{\mathbf{i}}$ defined in (4), the tuples $\mathbf{a}$ and $\mathbf{a}'$ have now the same role as coordinates of $\Pi_{\mathbf{i}}$. The tensor defined above satisfies the following simple identity, which should be compared to the one in Lemma 31 concerning $P_{\mathbf{i}}$.

Lemma 35. *For any $\mathbf{a} \in A^k$ and $\mathbf{i} \in [k]^k$,*

$$E_{\mathbf{a}} * \Pi_{\mathbf{i}} = \sum_{\substack{\mathbf{b} \in A^k \\ \mathbf{b_i} = \mathbf{a}}} E_{\mathbf{b}}.$$

Proof. For any $\mathbf{a}' \in A^k$, we have

$$\left(\sum_{\substack{\mathbf{b} \in A^k \\ \mathbf{b_i} = \mathbf{a}}} E_{\mathbf{b}} \right) * E_{\mathbf{a}'} = \sum_{\substack{\mathbf{b} \in A^k \\ \mathbf{b_i} = \mathbf{a}}} (E_{\mathbf{b}} * E_{\mathbf{a}'}) = \sum_{\substack{\mathbf{b} \in A^k \\ \mathbf{b_i} = \mathbf{a} \\ \mathbf{b} = \mathbf{a}'}} 1 = \begin{cases} 1 & \text{if } \mathbf{a}'_{\mathbf{i}} = \mathbf{a} \\ 0 & \text{otherwise} \end{cases} = (E_{\mathbf{a}} * \Pi_{\mathbf{i}}) * E_{\mathbf{a}'},$$

from which the result follows. □

Remark 36. It is clear from the expressions (4) and (5) that, for any $\mathbf{i} \in [k]^k$, $\Pi_{\mathbf{i}}$ coincides with the tensor $P_{\mathbf{i}}$ associated with the relation symbol R_k.

Proposition 37. *Let $\mathscr{M}$ be a linear minion, let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two k-enhanced σ-structures, and let $\xi : X^k \to \mathscr{M}^{(n^k)}$ be a map. Then, ξ preserves R_k (interpreted as a symbol in $\sigma^{ⓚ}$) if and only if*

$$\xi(\mathbf{x_i}) = \Pi_{\mathbf{i}} \overset{k}{*} \xi(\mathbf{x}) \qquad \textit{for any } \mathbf{x} \in X^k, \mathbf{i} \in [k]^k. \tag{6}$$

Proof. Suppose that ξ preserves R_k, and take $\mathbf{x} \in X^k = R_k^{\mathbf{X}}$. It follows that $\mathbf{x}^{ⓚ} \in R_k^{\mathbf{X}^{ⓚ}}$, so $\xi(\mathbf{x}^{ⓚ}) \in R_k^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$. This means that there exists $Q \in \mathscr{M}^{(|R_k^{\mathbf{A}}|)} = \mathscr{M}^{(n^k)}$ such that $\xi(\mathbf{x_i}) = Q_{/\pi_{\mathbf{i}}} = \Pi_{\mathbf{i}} \overset{k}{*} Q$ for each $\mathbf{i} \in [k]^k$ (where we have used Remark 36). Consider now the tuple $\mathbf{j} = (1, \dots, k) \in [k]^k$, and observe that $\mathbf{x_j} = \mathbf{x}$. Noticing that the contraction by $\Pi_{\mathbf{j}}$ acts as the identity, we conclude that

$$\xi(\mathbf{x}) = \xi(\mathbf{x_j}) = \Pi_{\mathbf{j}} \overset{k}{*} Q = Q,$$

which concludes the proof of (6).

Conversely, suppose (6) holds and take $\mathbf{x} \in R_k^{\mathbf{X}} = X^k$. We need to show that $\xi(\mathbf{x}^{ⓚ}) \in R_k^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$. Take $Q = \xi(\mathbf{x}) \in \mathscr{M}^{(n^k)} = \mathscr{M}^{(|R_k^{\mathbf{A}}|)}$. Using again Remark 36, we observe that, for any $\mathbf{i} \in [k]^k$,

$$\xi(\mathbf{x_i}) = \Pi_{\mathbf{i}} \overset{k}{*} \xi(\mathbf{x}) = \xi(\mathbf{x})_{/\pi_{\mathbf{i}}} = Q_{/\pi_{\mathbf{i}}},$$

whence the result follows. □

Proposition 37 distils the requirements of the BW^k, SA^k, AIP^k, SoS^k, and BA^k hierarchies enforcing compatibility between partial assignments[13] from $\mathbf{X}$ to $\mathbf{A}$ into a single list of tensor equations

$$\xi(\mathbf{x_i}) = \Pi_{\mathbf{i}} \overset{k}{*} \xi(\mathbf{x})$$

for $\mathbf{x} \in X^k$ and $\mathbf{i} \in [k]^k$. For $k = 1$, the equation is vacuous, since in this case $\Pi_{\mathbf{i}}$ is the identity matrix and $\mathbf{x_i} = \mathbf{x}$. As k increases, it produces a progressively richer system of symmetries that must be satisfied by ξ, which corresponds to a progressively stronger relaxation.

7 Conic minions

A primary message of this work is that the tensorisation construction establishes a correspondence between the algebraic properties of a minion and the algorithmic properties of the hierarchy of tests built on the minion. For example, we have seen that if the minion is linear some general properties of the solutions of the hierarchy can be deduced by studying the structure of the tensors in $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$. Now, the bounded width hierarchy has the property that it only seeks assignments that are partial homomorphisms; similarly, the Sherali–Adams, Sum-of-Squares, and BA^k hierarchies only assign a positive weight to solutions satisfying local constraints. The next definition identifies the minion property guaranteeing this algorithmic feature.

[13] Cf. the "closure under restriction" property of BW^k and the requirements ♣2 and ♠3 in Section 2.3 applied to $R = R_k$.

Definition 38. A linear minion $\mathscr{M}$ of depth d is *conic* if, for any $L \in \mathbb{N}$ and for any $M \in \mathscr{M}^{(L)}$, (i) $M \neq O_{L,d}$, and (ii) for any $V \subseteq [L]$, the following implication is true:[14]

$$\sum_{i \in V} M^T \mathbf{e}_i = \mathbf{0}_d \quad \Rightarrow \quad M^T \mathbf{e}_i = \mathbf{0}_d \;\; \forall i \in V.$$

Paraphrasing Definition 38, a linear minion $\mathscr{M}$ is conic if any matrix in $\mathscr{M}$ is nonzero and has the property that, whenever some of its rows sum up to the zero vector, each of those rows is the zero vector. It turns out that all minions appearing in Proposition 28 are conic, with the notable exception of $\mathscr{Z}_{\text{aff}}$.

Proposition 39. *The minions $\mathscr{H}$, $\mathscr{Q}_{\text{conv}}$, $\mathscr{S}$, and $\mathscr{M}_{\text{BA}}$ are conic, while the minion $\mathscr{Z}_{\text{aff}}$ is not.*

Proof. The fact that $\mathscr{Q}_{\text{conv}}$ is conic trivially follows by noting that its elements are nonnegative vectors whose entries sum up to 1.

Similarly, using the description of $\mathscr{H}$ as a linear minion on the semiring $(\{0,1\}, \vee, \wedge)$ (cf. the proof of Proposition 28), the fact that $\mathscr{H}$ is conic follows from the fact that $\bigvee_{i \in V} x_i = 0$ means that $x_i = 0$ for each $i \in V$. (Observe also that the vectors in $\mathscr{H}$ are nonzero, as they are the indicator vectors of nonempty sets.)

To show that $\mathscr{S}$ is conic, take $L \in \mathbb{N}$ and $M \in \mathscr{S}^{(L)}$, and notice first that $M \neq O_{L,\aleph_0}$ by (C3) in Definition 11. Take now $V \subseteq [L]$. If $\sum_{i \in V} M^T \mathbf{e}_i = \mathbf{0}_{\aleph_0}$, using that MM^T is a diagonal matrix by (C2), we find

$$0 = (\sum_{i \in V} M^T \mathbf{e}_i)^T (\sum_{j \in V} M^T \mathbf{e}_j) = \sum_{i,j \in V} \mathbf{e}_i^T M M^T \mathbf{e}_j = \sum_{i \in V} \mathbf{e}_i^T M M^T \mathbf{e}_i = \sum_{i \in V} \|M^T \mathbf{e}_i\|^2,$$

which means that $M^T \mathbf{e}_i = \mathbf{0}_{\aleph_0}$ for any $i \in V$, as required.

As for $\mathscr{M}_{\text{BA}}$, we shall see in Section 8 (cf. Example 48) that this minion can be obtained as the semi-direct product of $\mathscr{Q}_{\text{conv}}$ and $\mathscr{Z}_{\text{aff}}$. Then, the fact that $\mathscr{M}_{\text{BA}}$ is conic is a direct consequence of the fact that semi-direct products of minions are always conic (cf. Proposition 45).

Finally, the element $(1, -1, 1) \in \mathscr{Z}_{\text{aff}}$ witnesses that $\mathscr{Z}_{\text{aff}}$ is not conic. □

Remark 40. It is not hard to verify that also the minion capturing the power of the CLAP algorithm from [43] is linear (with $\mathcal{S} = \mathbb{Q}$ and $d = \aleph_0$) and, in fact, conic. Since an algorithmic hierarchy built on top of CLAP has never been studied in the literature, we do not consider the analogue of Theorem 25 for CLAP.

Remark 41. We point out that the concept of conic minions can be extended to arbitrary abstract minions (as opposed to just linear) via the notion of essential coordinates. Given an element M of a minion $\mathscr{M}$ of arity, say, L, we say that a coordinate $i \in [L]$ is *essential* for M if there exist an integer $L' \in \mathbb{N}$ and two minor maps $\pi, \pi' : [L] \to [L']$ such that $\pi|_{[L] \setminus \{i\}} = \pi'|_{[L] \setminus \{i\}}$ but $M_{/\pi} \neq M_{/\pi'}$. (This straightforwardly extends the analogous notion described in [12, Definition 5.14] for function minions.) Now, if $\mathscr{M}$ is linear of width d, it is easy to check that a coordinate i is essential for $M \in \mathscr{M}$ precisely when $M^T \mathbf{e}_i \neq \mathbf{0}_d$. Hence, the requirements (i) and (ii) in Definition 38 can be extended to abstract minions as follows: (i) asks that each $M \in \mathscr{M}$ should have at least one essential coordinate, and (ii) asks that the image of each essential coordinate of M under a minor map π should be essential in $M_{/\pi}$.

[14] As usual, the sum, product, 0, and 1 operations appearing in this definition are to be meant in the semiring $\mathcal{S}$ associated with the linear minion $\mathscr{M}$.

It turns out that this simple property guarantees that the hierarchies of tests built on conic minions only look at assignments yielding partial homomorphisms, as we shall see in Proposition 43. It also follows that conic hierarchies are not fooled by small instances: Proposition 44 establishes that the k-th level of such hierarchies is able to correctly classify instances on k (or fewer) elements, as it is well known for the bounded width, Sherali–Adams, and Sum-of-Squares hierarchies; we may informally express this property by saying that conic hierarchies are "sound in the limit". Moreover, we shall see in the next section that any linear minion can be transformed into a conic minion – whose hierarchy enjoys the features mentioned above – via the *semi-direct product* construction.

First of all, we prove that Lemma 34 can be slightly strengthened if we are dealing with conic minions, in that the level k for which it holds can be decreased down to 2. As a consequence, the algebraic description of the Sherali–Adams and Sum-of-Squares hierarchies in terms of the tensorisation construction can be extended to lower levels, cf. Remark 58. As in Section 6, the letter d shall denote the depth of the relevant minion in all results of this section.

Lemma 42. *Let $\mathscr{M}$ be a conic minion, let $2 \leq k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two k-enhanced σ-structures, and let $\xi : \mathbf{X}^{\textcircled{k}} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})$ be a homomorphism. For $R \in \sigma$ of arity r, let $\mathbf{x} \in R^{\mathbf{X}}$ and $\mathbf{a} \in R^{\mathbf{A}}$ be such that $\mathbf{x} \not\prec \mathbf{a}$. Let $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ be such that $Q_{/\pi_{\mathbf{i}}} = \xi(\mathbf{x}_{\mathbf{i}})$ for each $\mathbf{i} \in [r]^k$. Then $E_{\mathbf{a}} * Q = \mathbf{0}_d$.*

Proof. Take $\alpha, \beta \in [r]$ such that $x_\alpha = x_\beta$ and $a_\alpha \neq a_\beta$, and consider the tuple $\mathbf{j} = (\alpha, \dots, \alpha, \beta) \in [r]^k$. Using that $k \geq 2$, we have $\mathbf{x}_{\mathbf{j}} \not\prec \mathbf{a}_{\mathbf{j}}$, since $x_{j_{k-1}} = x_\alpha = x_\beta = x_{j_k}$ and $a_{j_{k-1}} = a_\alpha \neq a_\beta = a_{j_k}$. From Lemma 33 and Lemma 31 we obtain

$$\mathbf{0}_d = E_{\mathbf{a}_{\mathbf{j}}} * \xi(\mathbf{x}_{\mathbf{j}}) = E_{\mathbf{a}_{\mathbf{j}}} * Q_{/\pi_{\mathbf{j}}} = E_{\mathbf{a}_{\mathbf{j}}} * P_{\mathbf{j}} * Q = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{j}} = \mathbf{a}_{\mathbf{j}}}} E_{\mathbf{b}} * Q.$$

Using that $\mathscr{M}$ is a conic minion, we deduce that $E_{\mathbf{b}} * Q = \mathbf{0}_d$ for any $\mathbf{b} \in R^{\mathbf{A}}$ such that $\mathbf{b}_{\mathbf{j}} = \mathbf{a}_{\mathbf{j}}$. In particular, $E_{\mathbf{a}} * Q = \mathbf{0}_d$, as required. □

The following result shows that hierarchies of tests built on conic minions only give a nonzero weight to those assignments that yield partial homomorphisms.

Proposition 43. *Let $\mathscr{M}$ be a conic minion, let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two k-enhanced σ-structures such that $k \geq \min(2, \operatorname{armax}(\sigma))$, and let $\xi : \mathbf{X}^{\textcircled{k}} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})$ be a homomorphism. Let $R \in \sigma$ have arity r, and take $\mathbf{x} \in X^k$, $\mathbf{a} \in A^k$, and $\mathbf{i} \in [k]^r$. If $\mathbf{x}_{\mathbf{i}} \in R^{\mathbf{X}}$ and $\mathbf{a}_{\mathbf{i}} \notin R^{\mathbf{A}}$, then $E_{\mathbf{a}} * \xi(\mathbf{x}) = \mathbf{0}_d$.*

Proof. From $\mathbf{x}_{\mathbf{i}} \in R^{\mathbf{X}}$ we have $\mathbf{x}_{\mathbf{i}}^{\textcircled{k}} \in R^{\mathbf{X}^{\textcircled{k}}}$ and, thus, $\xi(\mathbf{x}_{\mathbf{i}}^{\textcircled{k}}) \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{\textcircled{k}})}$. It follows that there exists $Q \in \mathscr{M}^{(|R^{\mathbf{A}}|)}$ such that $\xi(\mathbf{x}_{\mathbf{i}_{\mathbf{j}}}) = Q_{/\pi_{\mathbf{j}}}$ for each $\mathbf{j} \in [r]^k$. Proposition 37 then yields

$$\Pi_{\mathbf{i}_{\mathbf{j}}} \overset{k}{*} \xi(\mathbf{x}) = \xi(\mathbf{x}_{\mathbf{i}_{\mathbf{j}}}) = Q_{/\pi_{\mathbf{j}}} = P_{\mathbf{j}} \overset{1}{*} Q. \tag{7}$$

Consider, for each $\alpha \in [k]$, the set $S_\alpha = \{\beta \in [r] : i_\beta = \alpha\}$, and fix an element $\hat{\beta} \in [r]$. The tuple $\mathbf{j} \in [r]^k$ defined by setting $j_\alpha = \min S_\alpha$ if $S_\alpha \neq \emptyset$, $j_\alpha = \hat{\beta}$ otherwise satisfies $\mathbf{i}_{\mathbf{j}_{\mathbf{i}}} = \mathbf{i}$.

Indeed, for any $\beta \in [r]$, we have $S_{i_\beta} \neq \emptyset$ since $\beta \in S_{i_\beta}$, so $j_{i_\beta} = \min S_{i_\beta} \in S_{i_\beta}$, which means that $i_{j_{i_\beta}} = i_\beta$, as required. We obtain

$$E_{\mathbf{a}_{\mathbf{i}_\mathbf{j}}} * P_\mathbf{j} * Q = \sum_{\substack{\mathbf{b} \in R^\mathbf{A} \\ \mathbf{b}_\mathbf{j} = \mathbf{a}_{\mathbf{i}_\mathbf{j}}}} E_\mathbf{b} * Q = \sum_{\substack{\mathbf{b} \in R^\mathbf{A} \\ \mathbf{x}_\mathbf{i} \prec \mathbf{b} \\ \mathbf{b}_\mathbf{j} = \mathbf{a}_{\mathbf{i}_\mathbf{j}}}} E_\mathbf{b} * Q, \tag{8}$$

where the first equality comes from Lemma 31 and the second from Lemma 34 or Lemma 42 (depending on whether $k \geq \mathrm{armax}(\sigma)$ or $k \geq 2$). We claim that the sum on the right-hand side of (8) equals $\mathbf{0}_d$. Indeed, let $\mathbf{b} \in A^r$ satisfy $\mathbf{x}_\mathbf{i} \prec \mathbf{b}$ and $\mathbf{b}_\mathbf{j} = \mathbf{a}_{\mathbf{i}_\mathbf{j}}$. Since $\mathbf{i}_{\mathbf{j}_\mathbf{i}} = \mathbf{i}$, for any $\alpha \in [r]$ we have $x_{i_\alpha} = x_{i_{j_{i_\alpha}}}$ and, hence, $b_\alpha = b_{j_{i_\alpha}}$. It follows that $\mathbf{b} = \mathbf{b}_{\mathbf{j}_\mathbf{i}} = \mathbf{a}_{\mathbf{i}_{\mathbf{j}_\mathbf{i}}} = \mathbf{a}_\mathbf{i} \notin R^\mathbf{A}$, which proves the claim. Combining this with (7), (8), and Lemma 35, we find

$$\begin{aligned}\mathbf{0}_d &= E_{\mathbf{a}_{\mathbf{i}_\mathbf{j}}} * P_\mathbf{j} * Q = E_{\mathbf{a}_{\mathbf{i}_\mathbf{j}}} * (P_\mathbf{j} \overset{1}{*} Q) = E_{\mathbf{a}_{\mathbf{i}_\mathbf{j}}} * (\Pi_{\mathbf{i}_\mathbf{j}} \overset{k}{*} \xi(\mathbf{x})) = E_{\mathbf{a}_{\mathbf{i}_\mathbf{j}}} * \Pi_{\mathbf{i}_\mathbf{j}} * \xi(\mathbf{x}) \\ &= \sum_{\substack{\mathbf{b} \in A^k \\ \mathbf{b}_{\mathbf{i}_\mathbf{j}} = \mathbf{a}_{\mathbf{i}_\mathbf{j}}}} E_\mathbf{b} * \xi(\mathbf{x})\end{aligned}$$

so, in particular, $E_\mathbf{a} * \xi(\mathbf{x}) = \mathbf{0}_d$ since $\mathscr{M}$ is a conic minion. □

The next result shows that hierarchies of tests built on conic minions are sound in the limit, in that they correctly classify all instances whose domain size is less than or equal to the hierarchy level.

Proposition 44. *Let $\mathscr{M}$ be a conic minion, let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two k-enhanced σ-structures such that $k \geq \min(2, \mathrm{armax}(\sigma))$ and $k \geq |X|$, and suppose that $\mathrm{Test}^k_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) =$ YES. Then $\mathbf{X} \to \mathbf{A}$.*

Proof. Let $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ be a homomorphism witnessing that $\mathrm{Test}^k_{\mathscr{M}}(\mathbf{X}, \mathbf{A}) =$ YES, and assume without loss of generality that $X = [\ell]$ with $\ell \in [k]$. Take the tuple $\mathbf{v} = (1, \ldots, \ell, \ell, \ldots, \ell) \in [\ell]^k$, and notice that $\xi(\mathbf{v}) \neq O_{n^k,d}$ since $\mathscr{M}$ is a conic minion. Therefore, there exists some $\mathbf{a} \in A^k$ such that $E_\mathbf{a} * \xi(\mathbf{v}) \neq \mathbf{0}_d$. Consider the function $f : X \to A$ defined by $x \to a_x$ for each $x \in X$. We claim that f yields a homomorphism from $\mathbf{X}$ to $\mathbf{A}$. Let $R \in \sigma$ be a relation symbol of arity r and take a tuple $\mathbf{x} \in R^\mathbf{X}$. Notice that $\mathbf{x} \in [\ell]^r \subseteq [k]^r$ and $f(\mathbf{x}) = \mathbf{a}_\mathbf{x}$. Furthermore, it holds that $\mathbf{v}_\mathbf{x} = \mathbf{x}$. Applying Proposition 43, we deduce that $f(\mathbf{x}) \in R^\mathbf{A}$, whence it follows that f is indeed a homomorphism. □

8 The semi-direct product of minions

Is it possible for multiple linear minions to "join forces", to obtain a new linear minion corresponding to a stronger relaxation? The natural way to do so is to take as the elements of the new linear minion block matrices, whose blocks are the elements of the original minions. Let $\mathscr{M}$ and $\mathscr{N}$ be two linear minions that we wish to "merge". A zero row in a matrix of $\mathscr{M}$ corresponds to zero weight assigned to the variable associated with the row by the relaxation given by $\mathscr{M}$. Ideally, we would like to preserve this information when we run the relaxation given by the second minion $\mathscr{N}$. In other words, we require that zero rows in $\mathscr{M}$ should be associated with zero rows in $\mathscr{N}$. For this to make sense (i.e., for the resulting object to be a linear minion), we need to assume that $\mathscr{M}$ is conic. Under this assumption, it turns out that the new linear minion is conic, too. Therefore, this construction yields a method to transform a linear minion into a conic one, by taking its product with a fixed conic minion (for instance, $\mathscr{Q}_{\text{conv}}$). Equivalently, the semi-direct product provides a way to turn a hierarchy of linear tests into a more powerful hierarchy of conic tests – which enjoys the appealing properties described in Section 7.

Proposition 45. *Let $\mathscr{M}$ be a conic minion with semiring $\mathcal{S}$ and depth d, let $\mathscr{N}$ be a linear minion with semiring $\mathcal{S}$ and depth d', and consider, for each $L \in \mathbb{N}$, the set $(\mathscr{M} \ltimes \mathscr{N})^{(L)} = \{[\begin{array}{cc} M & N \end{array}] : M \in \mathscr{M}^{(L)}, N \in \mathscr{N}^{(L)},$ and $N^T\mathbf{e}_i = \mathbf{0}_{d'}$ for any $i \in [L]$ such that $M^T\mathbf{e}_i = \mathbf{0}_d\}$. Then $\mathscr{M} \ltimes \mathscr{N} = \bigsqcup_{L\in\mathbb{N}} (\mathscr{M} \ltimes \mathscr{N})^{(L)}$ is a conic minion with semiring $\mathcal{S}$ and depth $d + d'$.*

Definition 46. Let $\mathscr{M}$ and $\mathscr{N}$ be a conic minion and a linear minion, respectively, over the same semiring. The *semi-direct product* of $\mathscr{M}$ and $\mathscr{N}$ is the conic minion $\mathscr{M} \ltimes \mathscr{N}$ described in Proposition 45.

Remark 47. If the minions $\mathscr{M}$ and $\mathscr{N}$ in Definition 46 have different semirings $\mathcal{S}$ and $\mathcal{S}'$, we cannot in general use the definition to build their semi-direct product. However, it is immediate to check that a linear minion over $\mathcal{S}$ is also a linear minion over any semiring of which $\mathcal{S}$ is a sub-semiring. Hence, if $\mathcal{S}$ is a sub-semiring of $\mathcal{S}'$ (or vice-versa), $\mathscr{M} \ltimes \mathscr{N}$ is well defined (see Example 48 below). In general, however, we might not be able to find a common semiring of which $\mathcal{S}$ and $\mathcal{S}'$ are both sub-semirings. In particular, it is not true in general that the direct sum of semirings admits homomorphic injections from the components, see [51]. To be able to define $\mathscr{M} \ltimes \mathscr{N}$ also in this case, we would need to redefine linear minions by allowing each of the d columns of the matrices in a linear minion of depth d to contain entries from a possibly different semiring. In this way, the requirement in Definition 46 can be circumvented – which, for example, makes it possible to define the minion $\mathscr{H} \ltimes \mathscr{Z}_{\text{aff}}$ (see Remark 59).

Example 48. Since $\mathbb{Z}$ is a sub-semiring of $\mathbb{Q}$, we can view $\mathscr{Q}_{\text{conv}}$ and $\mathscr{Z}_{\text{aff}}$ as a conic minion and a linear minion over the same semiring $\mathbb{Q}$, respectively. It is easy to check that their semi-direct product $\mathscr{Q}_{\text{conv}} \ltimes \mathscr{Z}_{\text{aff}}$ is precisely the minion $\mathscr{M}_{\text{BA}}$ from [27].

Proof of Proposition 45. We start by showing that $\mathscr{M} \ltimes \mathscr{N}$ is a linear minion of depth $d+d'$. Notice that each set $(\mathscr{M} \ltimes \mathscr{N})^{(L)}$ consists of $L \times (d+d')$ matrices having entries in $\mathcal{S}$. Given $\pi : [L] \to [L']$ and $[\begin{array}{cc} M & N \end{array}] \in (\mathscr{M} \ltimes \mathscr{N})^{(L)}$, we claim that $P[\begin{array}{cc} M & N \end{array}] = [\begin{array}{cc} PM & PN \end{array}]$ belongs to $(\mathscr{M} \ltimes \mathscr{N})^{(L')}$, where P is the $L' \times L$ matrix corresponding to π as per Definition 26. First, since $\mathscr{M}$ and $\mathscr{N}$ are both linear minions, we have that $PM = M_{/\pi} \in \mathscr{M}^{(L')}$ and $PN = N_{/\pi} \in \mathscr{N}^{(L')}$. Let $j \in [L']$ be such that $(PM)^T\mathbf{e}_j = \mathbf{0}_d$. We find

$$\mathbf{0}_d = M^T P^T \mathbf{e}_j = \sum_{i\in\pi^{-1}(j)} M^T\mathbf{e}_i.$$

Using that $\mathscr{M}$ is conic, we obtain $M^T\mathbf{e}_i = \mathbf{0}_d$ for each $i \in \pi^{-1}(j)$. By the definition of $(\mathscr{M} \ltimes \mathscr{N})^{(L)}$, this means that $N^T\mathbf{e}_i = \mathbf{0}_{d'}$ for each $i \in \pi^{-1}(j)$. Therefore,

$$\mathbf{0}_{d'} = \sum_{i\in\pi^{-1}(j)} N^T\mathbf{e}_i = N^T P^T \mathbf{e}_j = (PN)^T\mathbf{e}_j,$$

which proves the claim. It follows that $\mathscr{M} \ltimes \mathscr{N}$ is indeed a linear minion.

To show that $\mathscr{M} \ltimes \mathscr{N}$ is conic, we first note that no element of $\mathscr{M} \ltimes \mathscr{N}$ is the all-zero matrix, since $\mathscr{M}$ is conic. If $\sum_{i \in V} [\ M \quad N\ ]^T \mathbf{e}_i = \mathbf{0}_{d+d'}$ for some $L \in \mathbb{N}$, $[\ M \quad N\ ] \in (\mathscr{M} \ltimes \mathscr{N})^{(L)}$, and $V \subseteq [L]$, we find in particular that $\sum_{i \in V} M^T \mathbf{e}_i = \mathbf{0}_d$, which implies that $M^T \mathbf{e}_i = \mathbf{0}_d$ for each $i \in V$ by the fact that $\mathscr{M}$ is conic. By the definition of $(\mathscr{M} \ltimes \mathscr{N})^{(L)}$, this yields $N^T \mathbf{e}_i = \mathbf{0}_{d'}$ for each $i \in V$, so

$$[\ M \quad N\ ]^T \mathbf{e}_i = \begin{bmatrix} M^T \mathbf{e}_i \\ N^T \mathbf{e}_i \end{bmatrix} = \begin{bmatrix} \mathbf{0}_d \\ \mathbf{0}_{d'} \end{bmatrix} = \mathbf{0}_{d+d'}$$

for each $i \in V$, as required. □

The next result – crucial for the characterisation of the BA^k hierarchy in Theorem 25, cf. the proof of Proposition 57 – shows that homomorphisms corresponding to the semi-direct product of two minions factor into homomorphisms corresponding to the components.

Proposition 49. *Let $\mathscr{M}$ be a conic minion with semiring $\mathcal{S}$ and depth d, let $\mathscr{N}$ be a linear minion with semiring $\mathcal{S}$ and depth d', let $k \in \mathbb{N}$, and let $\mathbf{X}, \mathbf{A}$ be k-enhanced σ-structures such that $k \geq \operatorname{armax}(\sigma)$. Then there exists a homomorphism $\vartheta : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M} \ltimes \mathscr{N}}(\mathbf{A}^{ⓚ})$ if and only if there exist homomorphisms $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ and $\zeta : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{N}}(\mathbf{A}^{ⓚ})$ such that, for any $\mathbf{x} \in X^k$ and $\mathbf{a} \in A^k$, $E_{\mathbf{a}} * \xi(\mathbf{x}) = \mathbf{0}_d$ implies $E_{\mathbf{a}} * \zeta(\mathbf{x}) = \mathbf{0}_{d'}$.*

Proof. To prove the "if" part, take two homomorphisms ξ and ζ as in the statement of the proposition, and consider the map

$$\begin{aligned} \vartheta : X^k &\to (\mathscr{M} \ltimes \mathscr{N})^{(n^k)} \\ \mathbf{x} &\mapsto [\ \xi(\mathbf{x}) \quad \zeta(\mathbf{x})\ ]. \end{aligned}$$

Observe that ϑ is well defined since we are assuming that $E_{\mathbf{a}} * \xi(\mathbf{x}) = \mathbf{0}_d$ implies $E_{\mathbf{a}} * \zeta(\mathbf{x}) = \mathbf{0}_{d'}$ for any $\mathbf{x} \in X^k$ and $\mathbf{a} \in A^k$. We claim that ϑ yields a homomorphism from $\mathbf{X}^{ⓚ}$ to $\mathbb{F}_{\mathscr{M} \ltimes \mathscr{N}}(\mathbf{A}^{ⓚ})$. To this end, take $R \in \sigma$ of arity r and $\mathbf{x} \in R^{\mathbf{X}}$, so $\mathbf{x}^{ⓚ} \in R^{\mathbf{X}^{ⓚ}}$. We need to show that $\vartheta(\mathbf{x}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{M} \ltimes \mathscr{N}}(\mathbf{A}^{ⓚ})}$; equivalently, we need to find some $W \in (\mathscr{M} \ltimes \mathscr{N})^{(m)}$ such that $\vartheta(\mathbf{x_i}) = W_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$, where $m = |R^{\mathbf{A}}|$. Using that ξ is a homomorphism, we have that $\xi(\mathbf{x}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})}$, so there exists $Q \in \mathscr{M}^{(m)}$ for which $\xi(\mathbf{x_i}) = Q_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$. Similarly, using that ζ is a homomorphism, we can find $Z \in \mathscr{N}^{(m)}$ for which $\zeta(\mathbf{x_i}) = Z_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$. The crucial part is to show that $[\ Q \quad Z\ ] \in (\mathscr{M} \ltimes \mathscr{N})^{(m)}$. To this end, take $\mathbf{a} \in R^{\mathbf{A}}$ such that $E_{\mathbf{a}} * Q = \mathbf{0}_d$; we need to prove that $E_{\mathbf{a}} * Z = \mathbf{0}_{d'}$. Using the assumption that $k \geq r$, let us pick the tuple $\mathbf{j} = (1, 2, \ldots, r, 1, 1, \ldots, 1) \in [r]^k$. Notice that this choice guarantees that $\{\mathbf{b} \in R^{\mathbf{A}} : \mathbf{b_j} = \mathbf{a_j}\} = \{\mathbf{a}\}$. Hence,

$$E_{\mathbf{a_j}} * \xi(\mathbf{x_j}) = E_{\mathbf{a_j}} * Q_{/\pi_{\mathbf{j}}} = E_{\mathbf{a_j}} * P_{\mathbf{j}} * Q = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b_j} = \mathbf{a_j}}} E_{\mathbf{b}} * Q = E_{\mathbf{a}} * Q,$$

where the third equality follows from Lemma 31. Similarly, $E_{\mathbf{a_j}} * \zeta(\mathbf{x_j}) = E_{\mathbf{a}} * Z$. Then, from our assumption $E_{\mathbf{a}} * Q = \mathbf{0}_d$ it follows that $E_{\mathbf{a_j}} * \xi(\mathbf{x_j}) = \mathbf{0}_d$. Using the hypothesis of the proposition, we deduce that $E_{\mathbf{a_j}} * \zeta(\mathbf{x_j}) = \mathbf{0}_{d'}$, and we thus conclude that $E_{\mathbf{a}} * Z = \mathbf{0}_{d'}$, as wanted. Call $W = [\ Q \quad Z\ ]$. For each $\mathbf{i} \in [r]^k$, we find

$$\vartheta(\mathbf{x_i}) = [\ \xi(\mathbf{x_i}) \quad \zeta(\mathbf{x_i})\ ] = [\ Q_{/\pi_\mathbf{i}} \quad Z_{/\pi_\mathbf{i}}\ ] = [\ P_\mathbf{i} \overset{1}{*} Q \quad P_\mathbf{i} \overset{1}{*} Z\ ] = P_\mathbf{i} \overset{1}{*} [\ Q \quad Z\ ] = W_{/\pi_\mathbf{i}},$$

as required. This concludes the proof that ϑ is a homomorphism.

Conversely, let $\vartheta : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M} \ltimes \mathscr{N}}(\mathbf{A}^{ⓚ})$ be a homomorphism. For each $\mathbf{x} \in X^k$, write $\vartheta(\mathbf{x}) \in (\mathscr{M} \ltimes \mathscr{N})^{(n^k)}$ as $\vartheta(\mathbf{x}) = [\ M_{(\mathbf{x})} \quad N_{(\mathbf{x})}\ ]$, where $M_{(\mathbf{x})} \in \mathscr{M}^{(n^k)}$ and $N_{(\mathbf{x})} \in \mathscr{N}^{(n^k)}$. Using the same argument as in the previous part of the proof, we check that the assignment $\mathbf{x} \mapsto M_{(\mathbf{x})}$ (resp. $\mathbf{x} \mapsto N_{(\mathbf{x})}$) yields a homomorphism from $\mathbf{X}^{ⓚ}$ to $\mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ (resp. to $\mathbb{F}_{\mathscr{N}}(\mathbf{A}^{ⓚ})$), and that the implication $E_{\mathbf{a}} * \xi(\mathbf{x}) = \mathbf{0}_d \;\Rightarrow\; E_{\mathbf{a}} * \zeta(\mathbf{x}) = \mathbf{0}_{d'}$ is met for each $\mathbf{x} \in X^k$ and $\mathbf{a} \in A^k$. □

Remark 50. In recent work [55], Dalmau and Opršal studied reductions between PCSPs. For a minion $\mathscr{M}$, they construct a new minion $\omega(\mathscr{M})$ that they use to characterise the applicability of arc-consistency reductions. If $\mathscr{M}$ is linear, $\omega(\mathscr{M})$ coincides with the semi-direct product between $\mathscr{H}$ and $\mathscr{M}$ (cf. Remark 47). It is not hard to show that a linear minion $\mathscr{M}$ satisfies $\mathscr{M} \to \mathscr{H} \ltimes \mathscr{M}$ if and only if it is homomorphically equivalent to a conic minion. Using this fact, it can be shown[15] that the k-consistency reductions from [55] and the k-th level of a hierarchy of minion tests as defined in this paper are equivalent for conic minions. We also point out that the definition of semi-direct product of minions can be extended to abstract minions via the notion of essential coordinates (see Remark 41). Indeed, we can define $(\mathscr{M} \ltimes \mathscr{N})^{(L)}$ as the set of pairs (M, N) such that $M \in \mathscr{M}^{(L)}$, $N \in \mathscr{N}^{(L)}$, and any essential coordinate for N is essential for M, with minor maps defined coordinate-wise. Using essentially the same argument as in the proof of Proposition 45, we can verify that the set $\mathscr{M} \ltimes \mathscr{N} = \bigsqcup_{L \in \mathbb{N}} (\mathscr{M} \ltimes \mathscr{N})^{(L)}$ is closed under minor maps provided that $\mathscr{M}$ satisfies the abstract conicity properties of Remark 41.

9 A proof of Theorem 25

In this section, we prove Theorem 25 using the machinery developed in Sections 5, 6, 7, and 8.

Theorem (Theorem 25 restated). *If $k \in \mathbb{N}$ is at least the maximum arity of the template,*

- $\mathrm{BW}^k = \mathrm{Test}^k_{\mathscr{H}}$
- $\mathrm{SA}^k = \mathrm{Test}^k_{\mathscr{Q}_{\mathrm{conv}}}$
- $\mathrm{AIP}^k = \mathrm{Test}^k_{\mathscr{Z}_{\mathrm{aff}}}$
- $\mathrm{SoS}^k = \mathrm{Test}^k_{\mathscr{S}}$
- $\mathrm{BA}^k = \mathrm{Test}^k_{\mathscr{M}_{\mathrm{BA}}}$.

The five parts of the theorem will be formally stated and proved separately, in Propositions 51, 53, 55, 56, and 57. The statements in Propositions 51 and 56, concerning SA^k and SoS^k, are actually slightly stronger than Theorem 25, as they do not require that the level k of the hierarchy be at least the maximum arity of the template (cf. Remark 58). We start with SA^k, whose proof is slightly simpler than (and provides intuition for) the proof for BW^k.

[15] Personal communication with Jakub Opršal.

Proposition 51. *Let $k \in \mathbb{N}$ and let $\mathbf{X}, \mathbf{A}$ be k-enhanced σ-structures such that $k \geq \min(2, \operatorname{armax}(\sigma))$. Then* $\operatorname{SA}^k(\mathbf{X}, \mathbf{A}) = \operatorname{Test}^k_{\mathscr{Q}_{\text{conv}}}(\mathbf{X}, \mathbf{A})$.

Proof. Suppose $\operatorname{SA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$ and let the rational numbers $\lambda_{R,\mathbf{x},\mathbf{a}}$ witness it, for $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$. Consider the map $\xi : X^k \to \mathcal{T}^{n \cdot \mathbf{1}_k}(\mathbb{Q})$ defined by $E_{\mathbf{a}} * \xi(\mathbf{x}) = \lambda_{R_k,\mathbf{x},\mathbf{a}}$ for any $\mathbf{x} \in X^k$, $\mathbf{a} \in A^k$. We claim that ξ yields a homomorphism from $\mathbf{X}^{Ⓚ}$ to $\mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{Ⓚ})$. Notice first that, for any $\mathbf{x} \in X^k$, $\xi(\mathbf{x})$ is an entrywise nonnegative tensor in the space $\mathcal{T}^{n \cdot \mathbf{1}_k}(\mathbb{Q})$ (which can be identified with $\mathcal{T}^{n \cdot \mathbf{1}_k, 1}(\mathbb{Q})$). Moreover, using ♣1, we find

$$\sum_{\mathbf{a} \in A^k} E_{\mathbf{a}} * \xi(\mathbf{x}) = \sum_{\mathbf{a} \in R_k^{\mathbf{A}}} \lambda_{R_k,\mathbf{x},\mathbf{a}} = 1.$$

It follows that $\xi(\mathbf{x}) \in \mathscr{Q}_{\text{conv}}{}^{(n^k)}$. We now prove that ξ yields a homomorphism from $\mathbf{X}^{Ⓚ}$ to $\mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{Ⓚ})$. Take a symbol $R \in \sigma$ of arity r and a tuple $\mathbf{x} \in R^{\mathbf{X}}$, so that $\mathbf{x}^{Ⓚ} \in R^{\mathbf{X}^{Ⓚ}}$. We need to show that $\xi(\mathbf{x}^{Ⓚ}) \in R^{\mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{Ⓚ})}$. Equivalently, we seek some vector $\mathbf{q} \in \mathscr{Q}_{\text{conv}}{}^{(|R^{\mathbf{A}}|)}$ such that $\xi(\mathbf{x}_{\mathbf{i}}) = \mathbf{q}_{/\pi_{\mathbf{i}}}$ for any $\mathbf{i} \in [r]^k$. Consider the vector $\mathbf{q} \in \mathcal{T}^{|R^{\mathbf{A}}|}(\mathbb{Q})$ defined by $E_{\mathbf{a}} * \mathbf{q} = \lambda_{R,\mathbf{x},\mathbf{a}}$ for any $\mathbf{a} \in R^{\mathbf{A}}$. Similarly as before, ♣1 implies that $\mathbf{q} \in \mathscr{Q}_{\text{conv}}{}^{(|R^{\mathbf{A}}|)}$. For $\mathbf{i} \in [r]^k$ and $\mathbf{a} \in A^k$, we have

$$E_{\mathbf{a}} * \xi(\mathbf{x}_{\mathbf{i}}) = \lambda_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{a}} = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a}}} \lambda_{R,\mathbf{x},\mathbf{b}} = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a}}} E_{\mathbf{b}} * \mathbf{q} = E_{\mathbf{a}} * P_{\mathbf{i}} * \mathbf{q} = E_{\mathbf{a}} * \mathbf{q}_{/\pi_{\mathbf{i}}},$$

where the second equality is ♣2 and the fourth follows from Lemma 31. We deduce that ξ is a homomorphism, as claimed, which means that $\operatorname{Test}^k_{\mathscr{Q}_{\text{conv}}}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$.

Conversely, suppose that ξ is a homomorphism from $\mathbf{X}^{Ⓚ}$ to $\mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{Ⓚ})$ witnessing that $\operatorname{Test}^k_{\mathscr{Q}_{\text{conv}}}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. We associate with any pair $(R, \mathbf{x})$ such that $R \in \sigma$ and $\mathbf{x} \in R^{\mathbf{X}}$ a vector $\mathbf{q}_{R,\mathbf{x}} \in \mathscr{Q}_{\text{conv}}{}^{(|R^{\mathbf{A}}|)}$ defined as follows. Using that $\mathbf{x}^{Ⓚ} \in R^{\mathbf{X}^{Ⓚ}}$ and ξ is a homomorphism, we deduce that $\xi(\mathbf{x}^{Ⓚ}) \in R^{\mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{Ⓚ})}$ – i.e., there exists $\mathbf{q} \in \mathscr{Q}_{\text{conv}}{}^{(|R^{\mathbf{A}}|)}$ such that $\xi(\mathbf{x}_{\mathbf{i}}) = \mathbf{q}_{/\pi_{\mathbf{i}}} = P_{\mathbf{i}} * \mathbf{q}$ for each $\mathbf{i} \in [r]^k$, where r is the arity of R. We set $\mathbf{q}_{R,\mathbf{x}} = \mathbf{q}$. We now build a solution to $\operatorname{SA}^k(\mathbf{X}, \mathbf{A})$ as follows: For any $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$, we set $\lambda_{R,\mathbf{x},\mathbf{a}} = E_{\mathbf{a}} * \mathbf{q}_{R,\mathbf{x}}$. Notice that each $\lambda_{R,\mathbf{x},\mathbf{a}}$ is a rational number in the interval $[0, 1]$. Moreover, for $R \in \sigma$ and $\mathbf{x} \in R^{\mathbf{X}}$, we have

$$\sum_{\mathbf{a} \in R^{\mathbf{A}}} \lambda_{R,\mathbf{x},\mathbf{a}} = \sum_{\mathbf{a} \in R^{\mathbf{A}}} E_{\mathbf{a}} * \mathbf{q}_{R,\mathbf{x}} = 1$$

since $\mathbf{q}_{R,\mathbf{x}} \in \mathscr{Q}_{\text{conv}}$, thus yielding ♣1. Observe that, for any $\mathbf{y} \in X^k$, we have $\mathbf{q}_{R_k,\mathbf{y}} = \xi(\mathbf{y})$. Indeed, letting $\mathbf{j} = (1, \dots, k) \in [k]^k$, we have

$$\mathbf{q}_{R_k,\mathbf{y}} = \Pi_{\mathbf{j}} * \mathbf{q}_{R_k,\mathbf{y}} = P_{\mathbf{j}} * \mathbf{q}_{R_k,\mathbf{y}} = \mathbf{q}_{R_k,\mathbf{y}_{/\pi_{\mathbf{j}}}} = \xi(\mathbf{y}_{\mathbf{j}}) = \xi(\mathbf{y}), \tag{9}$$

where the first equality follows from the fact that the contraction by $\Pi_{\mathbf{j}}$ acts as the identity (cf. the proof of Proposition 37) and the second from Remark 36. Then, ♣2 follows by noticing that, for $\mathbf{i} \in [r]^k$ and $\mathbf{b} \in A^k$,

$$\sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ \mathbf{a}_{\mathbf{i}} = \mathbf{b}}} \lambda_{R,\mathbf{x},\mathbf{a}} = \sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ \mathbf{a}_{\mathbf{i}} = \mathbf{b}}} E_{\mathbf{a}} * \mathbf{q}_{R,\mathbf{x}} = E_{\mathbf{b}} * P_{\mathbf{i}} * \mathbf{q}_{R,\mathbf{x}} = E_{\mathbf{b}} * \xi(\mathbf{x}_{\mathbf{i}}) = E_{\mathbf{b}} * \mathbf{q}_{R_k,\mathbf{x}_{\mathbf{i}}} = \lambda_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{b}},$$

where the second and fourth equalities follow from Lemma 31 and (9), respectively. Recall from Proposition 39 that $\mathscr{Q}_{\text{conv}}$ is a conic minion. Using either Lemma 34 or Lemma 42 (depending on whether $k \geq \operatorname{armax}(\sigma)$ or $k \geq 2$), if $\mathbf{a} \in R^{\mathbf{A}}$ is such that $\mathbf{x} \not\prec \mathbf{a}$, we obtain

$$\lambda_{R,\mathbf{x},\mathbf{a}} = E_{\mathbf{a}} * \mathbf{q}_{R,\mathbf{x}} = 0,$$

thus showing that ♣3 is satisfied, too. It follows that $\mathrm{SA}^k(\mathbf{X}, \mathbf{A}) = \mathrm{YES}$, as required. □

Remark 52. Observe that condition ♣3 is not used in the first implication of the proof of Proposition 51, showing that $\mathrm{SA}^k(\mathbf{X}, \mathbf{A}) = \mathrm{YES}$ implies $\mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{ⓚ})$. Since, however, ♣3 is implied by $\mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{Q}_{\text{conv}}}(\mathbf{A}^{ⓚ})$ (as showed in the second part of the proof), it follows that ♣3 is redundant in the system (♣) for SA^k when $k \geq \min(2, \operatorname{armax}(\sigma))$. This fact can be easily seen from Lemma 34 and Lemma 42, which guarantee that the $\mathbf{a}$-th entry of the witness that a homomorphism $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}}(\mathbf{A}^{ⓚ})$ (i.e., a hierarchy solution) preserves a tuple $\mathbf{x} \in R^{\mathbf{X}}$ is zero whenever $\mathbf{x} \not\prec \mathbf{a}$—which precisely corresponds to condition ♣3. Similarly, the proofs of Propositions 55 and 57 show that ♣3 is redundant for both AIP^k and BA^k, while the proof of Proposition 56 shows that ♠4 is redundant for SoS^k. We do not omit such conditions from the definition of the hierarchies in Section 2.3 since this is the way they are commonly described in the literature, with variables corresponding to sets $V \subseteq X$ and partial assignments from V to A.

Proposition 53. *Let $k \in \mathbb{N}$ and let $\mathbf{X}, \mathbf{A}$ be k-enhanced σ-structures such that $k \geq \operatorname{armax}(\sigma)$. Then $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \mathrm{Test}^k_{\mathscr{H}}(\mathbf{X}, \mathbf{A})$.*

Proof. Given two sets S, T, an integer $p \in \mathbb{N}$, and two tuples $\mathbf{s} = (s_1, \dots, s_p) \in S^p, \mathbf{t} = (t_1, \dots, t_p) \in T^p$ such that $\mathbf{s} \prec \mathbf{t}$, we shall consider the function $f_{\mathbf{s},\mathbf{t}} : \{\mathbf{s}\} \to T$ defined by $f_{\mathbf{s},\mathbf{t}}(s_\alpha) = t_\alpha$ for each $\alpha \in [p]$. Also, we denote by $\epsilon : \emptyset \to \mathbf{A}$ the empty mapping.

Suppose $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \mathrm{YES}$ and let $\mathcal{F}$ be a nonempty collection of partial homomorphisms from $\mathbf{X}$ to $\mathbf{A}$ witnessing it. Recall from Section 2.1 that the space of tensors $\mathcal{T}^{n\cdot\mathbf{1}_k,1}(\{0,1\})$ can be identified with $\mathcal{T}^{n\cdot\mathbf{1}_k}(\{0,1\})$. Define the map $\xi : X^k \to \mathcal{T}^{n\cdot\mathbf{1}_k}(\{0,1\})$ by setting, for $\mathbf{x} \in X^k$ and $\mathbf{a} \in A^k$, $E_{\mathbf{a}} * \xi(\mathbf{x}) = 1$ if $\mathbf{x} \prec \mathbf{a}$ and $f_{\mathbf{x},\mathbf{a}} \in \mathcal{F}$, $E_{\mathbf{a}} * \xi(\mathbf{x}) = 0$ otherwise. We claim that ξ yields a homomorphism from $\mathbf{X}^{ⓚ}$ to $\mathbb{F}_{\mathscr{H}}(\mathbf{A}^{ⓚ})$. Take $R \in \sigma$ of arity r and $\mathbf{y} \in R^{\mathbf{X}}$, so $\mathbf{y}^{ⓚ} \in R^{\mathbf{X}^{ⓚ}}$. We need to show that $\xi(\mathbf{y}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{H}}(\mathbf{A}^{ⓚ})}$. Since $k \geq r$, we can write $\mathbf{y} = \mathbf{x}_{\mathbf{i}}$ for some $\mathbf{x} \in X^k$, $\mathbf{i} \in [k]^r$. Given $\mathbf{a} \in R^{\mathbf{A}}$, consider the set $B_{\mathbf{a}} = \{\mathbf{b} \in A^k : \mathbf{b}_{\mathbf{i}} = \mathbf{a} \text{ and } \mathbf{x} \prec \mathbf{b}\}$. We define a vector $\mathbf{q} \in \mathcal{T}^{|R^{\mathbf{A}}|}(\{0,1\})$ by letting, for each $\mathbf{a} \in R^{\mathbf{A}}$, the $\mathbf{a}$-th entry of $\mathbf{q}$ be 1 if $f_{\mathbf{x},\mathbf{b}} \in \mathcal{F}$ for some $\mathbf{b} \in B_{\mathbf{a}}$, 0 otherwise. We now show that $\mathbf{q} \in \mathscr{H}^{(|R^{\mathbf{A}}|)}$; i.e., that $\mathbf{q}$ is not identically zero. Observe first that, since $\mathcal{F}$ is nonempty and closed under restrictions, it contains the empty mapping ϵ. Applying the extension property to ϵ, we find that there exists some $f \in \mathcal{F}$ whose domain is $\{\mathbf{x}\}$ – that is, there exists some $\mathbf{c} \in A^k$ such that $\mathbf{x} \prec \mathbf{c}$ and $f_{\mathbf{x},\mathbf{c}} = f \in \mathcal{F}$. Notice that $\mathbf{y} \in R^{\mathbf{X}} \cap \{\mathbf{x}\}^r = R^{\mathbf{X}[\{\mathbf{x}\}]}$ (where we recall that $\mathbf{X}[\{\mathbf{x}\}]$ is the substructure of $\mathbf{X}$ induced by $\{\mathbf{x}\}$). Using that $f_{\mathbf{x},\mathbf{c}}$ is a partial homomorphism, we obtain $\mathbf{c}_{\mathbf{i}} = f_{\mathbf{x},\mathbf{c}}(\mathbf{x}_{\mathbf{i}}) = f_{\mathbf{x},\mathbf{c}}(\mathbf{y}) \in R^{\mathbf{A}}$. We then conclude that $\mathbf{e}_{\mathbf{c}_{\mathbf{i}}} * \mathbf{q} = 1$, so $\mathbf{q} \in \mathscr{H}^{(|R^{\mathbf{A}}|)}$, as required. If we manage to show that $\mathbf{q}_{/\pi_{\boldsymbol{\ell}}} = \xi(\mathbf{y}_{\boldsymbol{\ell}})$ for any $\boldsymbol{\ell} \in [r]^k$, we can conclude that $\xi(\mathbf{y}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{H}}(\mathbf{A}^{ⓚ})}$, thus proving the claim. Recall from Proposition 28 that $\mathscr{H}$ is a linear minion on the semiring $(\{0,1\}, \vee, \wedge)$. For $\mathbf{a} \in A^k$, using Lemma 31, we find

$$E_{\mathbf{a}} * \mathbf{q}_{/\pi_{\boldsymbol{\ell}}} = E_{\mathbf{a}} * P_{\boldsymbol{\ell}} * \mathbf{q} = \sum_{\substack{\mathbf{c} \in R^{\mathbf{A}} \\ \mathbf{c}_{\boldsymbol{\ell}} = \mathbf{a}}} E_{\mathbf{c}} * \mathbf{q} = \bigvee_{\substack{\mathbf{c} \in R^{\mathbf{A}} \\ \mathbf{c}_{\boldsymbol{\ell}} = \mathbf{a}}} E_{\mathbf{c}} * \mathbf{q}. \tag{10}$$

It follows that the expression in (10) equals 1 if

$$\exists \mathbf{c} \in R^{\mathbf{A}} \text{ s.t. } \mathbf{c}_{\boldsymbol{\ell}} = \mathbf{a} \text{ and } f_{\mathbf{x},\mathbf{b}} \in \mathcal{F} \text{ for some } \mathbf{b} \in B_{\mathbf{c}}, \tag{$\star$}$$

0 otherwise. On the other hand, $E_{\mathbf{a}} * \xi(\mathbf{y}_{\boldsymbol{\ell}})$ equals 1 if

$$\mathbf{y}_{\boldsymbol{\ell}} \prec \mathbf{a} \text{ and } f_{\mathbf{y}_{\boldsymbol{\ell}},\mathbf{a}} \in \mathcal{F}, \tag{$\bullet$}$$

0 otherwise. We now show that the conditions ($\star$) and ($\bullet$) are equivalent, which concludes the proof of the claim. Suppose that ($\star$) holds. Since $\mathbf{b} \in B_{\mathbf{c}}$, we have $\mathbf{x} \prec \mathbf{b}$, which yields $\mathbf{x}_{\mathbf{i}_{\boldsymbol{\ell}}} \prec \mathbf{b}_{\mathbf{i}_{\boldsymbol{\ell}}}$ as "$\prec$" is preserved under projections. Using the restriction property applied to $f_{\mathbf{x},\mathbf{b}}$, we find that $f_{\mathbf{x}_{\mathbf{i}_{\boldsymbol{\ell}}},\mathbf{b}_{\mathbf{i}_{\boldsymbol{\ell}}}} \in \mathcal{F}$. Then, ($\bullet$) follows by observing that $\mathbf{x}_{\mathbf{i}_{\boldsymbol{\ell}}} = \mathbf{y}_{\boldsymbol{\ell}}$ and, since $\mathbf{b} \in B_{\mathbf{c}}$ and $\mathbf{c}_{\boldsymbol{\ell}} = \mathbf{a}$, $\mathbf{b}_{\mathbf{i}_{\boldsymbol{\ell}}} = \mathbf{a}$. Suppose now that ($\bullet$) holds. Using the extension property applied to $f_{\mathbf{y}_{\boldsymbol{\ell}},\mathbf{a}} = f_{\mathbf{x}_{\mathbf{i}_{\boldsymbol{\ell}}},\mathbf{a}}$ we find that $f_{\mathbf{x},\mathbf{b}} \in \mathcal{F}$ for some $\mathbf{b} \in A^k$ such that $\mathbf{x} \prec \mathbf{b}$ and $\mathbf{b}_{\mathbf{i}_{\boldsymbol{\ell}}} = \mathbf{a}$. Since $f_{\mathbf{x},\mathbf{b}}$ is a partial homomorphism from $\mathbf{X}$ to $\mathbf{A}$, $R^{\mathbf{A}} \ni f_{\mathbf{x},\mathbf{b}}(\mathbf{y}) = f_{\mathbf{x},\mathbf{b}}(\mathbf{x}_{\mathbf{i}}) = \mathbf{b}_{\mathbf{i}}$. Calling $\mathbf{c} = \mathbf{b}_{\mathbf{i}}$, we obtain ($\star$).

Conversely, suppose that $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{H}}(\mathbf{A}^{ⓚ})$ is a homomorphism witnessing that $\text{Test}^k_{\mathscr{H}}(\mathbf{X}, \mathbf{A})$ accepts, and consider the collection $\mathcal{F} = \{f_{\mathbf{x},\mathbf{a}} : \mathbf{x} \in X^k, \mathbf{a} \in \text{supp}(\xi(\mathbf{x}))\} \cup \{\epsilon\}$. Notice that $\mathcal{F}$ is well defined by virtue of Lemma 33, as $\mathbf{a} \in \text{supp}(\xi(\mathbf{x}))$ implies that $\mathbf{x} \prec \mathbf{a}$, and it is nonempty. We claim that any function $f_{\mathbf{x},\mathbf{a}} \in \mathcal{F}$ is a partial homomorphism from $\mathbf{X}$ to $\mathbf{A}$. (Notice that ϵ is trivially a partial homomorphism.) Indeed, given $R \in \sigma$ of arity r and $\mathbf{y} \in R^{\mathbf{X}[\{\mathbf{x}\}]} = R^{\mathbf{X}} \cap \{\mathbf{x}\}^r$, we can write $\mathbf{y} = \mathbf{x}_{\mathbf{i}}$ for some $\mathbf{i} \in [k]^r$. Then, $f_{\mathbf{x},\mathbf{a}}(\mathbf{y}) = f_{\mathbf{x},\mathbf{a}}(\mathbf{x}_{\mathbf{i}}) = \mathbf{a}_{\mathbf{i}} \in R^{\mathbf{A}}$, where we have used Proposition 43 (which applies to $\mathscr{H}$ since, by Proposition 39, $\mathscr{H}$ is a conic minion). To show that $\mathcal{F}$ is closed under restrictions, take $f \in \mathcal{F}$ and $V \subseteq \text{dom}(f)$; we need to show that $f|_V \in \mathcal{F}$. The cases $f = \epsilon$ or $V = \emptyset$ are trivial, so we can assume $f = f_{\mathbf{x},\mathbf{a}}$ (which means that $\text{dom}(f) = \{\mathbf{x}\}$) and write $V = \{\mathbf{x}_{\boldsymbol{\ell}}\}$ for some $\boldsymbol{\ell} \in [k]^k$. Observe that $f_{\mathbf{x},\mathbf{a}}|_V = f_{\mathbf{x}_{\boldsymbol{\ell}},\mathbf{a}_{\boldsymbol{\ell}}}$. We claim that $E_{\mathbf{a}_{\boldsymbol{\ell}}} * \xi(\mathbf{x}_{\boldsymbol{\ell}}) = 1$. Otherwise, using Proposition 37 and Lemma 35, we would have

$$0 = E_{\mathbf{a}_{\boldsymbol{\ell}}} * \xi(\mathbf{x}_{\boldsymbol{\ell}}) = E_{\mathbf{a}_{\boldsymbol{\ell}}} * (\Pi_{\boldsymbol{\ell}} \overset{k}{*} \xi(\mathbf{x})) = E_{\mathbf{a}_{\boldsymbol{\ell}}} * \Pi_{\boldsymbol{\ell}} * \xi(\mathbf{x}) = \bigvee_{\substack{\mathbf{b} \in A^k \\ \mathbf{b}_{\boldsymbol{\ell}} = \mathbf{a}_{\boldsymbol{\ell}}}} E_{\mathbf{b}} * \xi(\mathbf{x}).$$

This would imply that $E_{\mathbf{b}} * \xi(\mathbf{x}) = 0$ whenever $\mathbf{b} \in A^k$ is such that $\mathbf{b}_{\boldsymbol{\ell}} = \mathbf{a}_{\boldsymbol{\ell}}$; in particular, $E_{\mathbf{a}} * \xi(\mathbf{x}) = 0$, a contradiction. So $E_{\mathbf{a}_{\boldsymbol{\ell}}} * \xi(\mathbf{x}_{\boldsymbol{\ell}}) = 1$, as claimed, and it follows that $f_{\mathbf{x}_{\boldsymbol{\ell}},\mathbf{a}_{\boldsymbol{\ell}}} \in \mathcal{F}$. We now claim that $\mathcal{F}$ has the extension property up to k. Take $f \in \mathcal{F}$ and $V \subseteq X$ such that $|V| \leq k$ and $\text{dom}(f) \subseteq V$; we need to show that there exists $g \in \mathcal{F}$ such that g extends f and $\text{dom}(g) = V$. If $f = \epsilon$ and $V = \emptyset$, the claim is trivial; if $f = \epsilon$ and $V \neq \emptyset$, we can write $V = \{\mathbf{x}\}$ for some $\mathbf{x} \in X^k$, and the claim follows by noticing that, by the definition of $\mathscr{H}$, $\text{supp}(\xi(\mathbf{x})) \neq \emptyset$. Therefore, we can assume that $f \neq \epsilon$, so $f = f_{\mathbf{x},\mathbf{a}}$ for some $\mathbf{x} \in X^k$, $\mathbf{a} \in \text{supp}(\xi(\mathbf{x}))$. Since $V \neq \emptyset$, we can write $V = \{\mathbf{y}\}$ for some $\mathbf{y} \in X^k$. Then, $\text{dom}(f) \subseteq V$ becomes $\{\mathbf{x}\} \subseteq \{\mathbf{y}\}$, so $\mathbf{x} = \mathbf{y}_{\boldsymbol{\ell}}$ for some $\boldsymbol{\ell} \in [k]^k$. Using Proposition 37 and Lemma 35, we find

$$1 = E_{\mathbf{a}} * \xi(\mathbf{x}) = E_{\mathbf{a}} * \xi(\mathbf{y}_{\boldsymbol{\ell}}) = E_{\mathbf{a}} * \Pi_{\boldsymbol{\ell}} * \xi(\mathbf{y}) = \bigvee_{\substack{\mathbf{b} \in A^k \\ \mathbf{b}_{\boldsymbol{\ell}} = \mathbf{a}}} E_{\mathbf{b}} * \xi(\mathbf{y}),$$

which implies that $E_{\mathbf{b}} * \xi(\mathbf{y}) = 1$ for some $\mathbf{b} \in A^k$ such that $\mathbf{b}_{\boldsymbol{\ell}} = \mathbf{a}$. It follows that $f_{\mathbf{y},\mathbf{b}} \in \mathcal{F}$. Notice that $\text{dom}(f_{\mathbf{y},\mathbf{b}}) = \{\mathbf{y}\} = V$, and $f_{\mathbf{y},\mathbf{b}}|_{\{\mathbf{x}\}} = f_{\mathbf{x},\mathbf{a}}$, so the claim is true. Hence, $\mathcal{F}$ witnesses that $\text{BW}^k(\mathbf{X}, \mathbf{A}) = \text{YES}$. □

Remark 54. Recall from Remark 4 that, unlike the other hierarchies in Theorem 25, the BW^k hierarchy does not require that the structures $\mathbf{X}$ and $\mathbf{A}$ to which it is applied should be k-enhanced. In particular, the proof of Proposition 53 actually establishes the following, slightly stronger statement: *Let $k \in \mathbb{N}$ and let $\mathbf{X}, \mathbf{A}$ be σ-structures such that $k \geq \mathrm{armax}(\sigma)$. Then $\mathrm{BW}^k(\mathbf{X}, \mathbf{A}) = \mathrm{Test}^k_{\mathscr{H}}(\tilde{\mathbf{X}}, \tilde{\mathbf{A}})$, where $\tilde{\mathbf{X}}$ (resp. $\tilde{\mathbf{A}}$) is the k-enhanced version of $\mathbf{X}$ (resp. $\mathbf{A}$).*

Proposition 55. *Let $k \in \mathbb{N}$ and let $\mathbf{X}, \mathbf{A}$ be k-enhanced σ-structures such that $k \geq \mathrm{armax}(\sigma)$. Then $\mathrm{AIP}^k(\mathbf{X}, \mathbf{A}) = \mathrm{Test}^k_{\mathscr{Z}_{\text{aff}}}(\mathbf{X}, \mathbf{A})$.*

Proof. The proof is analogous to that of Proposition 51, the only difference being that Lemma 42 cannot be applied in this case since $\mathscr{Z}_{\text{aff}}$ is not a conic minion (cf. Proposition 39). As a consequence, unlike in Proposition 51, we now need to assume that $k \geq \mathrm{armax}(\sigma)$. $\square$

The proof of Theorem 25 for SoS^k, given in Proposition 56 below, follows the same scheme as that of Proposition 51. There is, however, one additional complication due to the fact that the objects in the minion $\mathscr{S}$ are matrices having infinitely many columns. We deal with this technical issue through the orthonormalisation argument we already used for the proof of Proposition 13: We find an orthonormal basis for the finitely generated vector space defined as the sum of the row-spaces of the matrices in $\mathscr{S}$ appearing as images of a given homomorphism.

Proposition 56. *Let $k \in \mathbb{N}$ and let $\mathbf{X}, \mathbf{A}$ be k-enhanced σ-structures such that $k \geq \min(2, \mathrm{armax}(\sigma))$. Then $\mathrm{SoS}^k(\mathbf{X}, \mathbf{A}) = \mathrm{Test}^k_{\mathscr{S}}(\mathbf{X}, \mathbf{A})$.*

Proof. Suppose that $\mathrm{SoS}^k(\mathbf{X}, \mathbf{A}) = \text{YES}$ and let the family of vectors $\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \in \mathbb{R}^\gamma$ witness it, where $\gamma = \sum_{R \in \sigma} |R^{\mathbf{X}}| \cdot |R^{\mathbf{A}}|$. Consider the map $\xi : X^k \to \mathcal{T}^{n \cdot \mathbf{1}_k, \aleph_0}(\mathbb{R})$ defined by

$$E_{\mathbf{a}} * \xi(\mathbf{x}) = \begin{bmatrix} \boldsymbol{\lambda}_{R_k,\mathbf{x},\mathbf{a}} \\ \mathbf{0}_{\aleph_0} \end{bmatrix} \qquad \mathbf{x} \in X^k, \mathbf{a} \in A^k.$$

We claim that ξ yields a homomorphism from $\mathbf{X}^{ⓚ}$ to $\mathbb{F}_{\mathscr{S}}(\mathbf{A}^{ⓚ})$. First of all, we need to show that $\xi(\mathbf{x}) \in \mathscr{S}^{(n^k)}$ for each $\mathbf{x} \in X^k$. The requirement (C1) is trivially satisfied since, by construction, the j-th entry of $E_{\mathbf{a}} * \xi(\mathbf{x})$ is zero whenever $j > \gamma$. Given $\mathbf{a}, \mathbf{a}' \in A^k$,

$$(E_{\mathbf{a}} * \xi(\mathbf{x}))^T (E_{\mathbf{a}'} * \xi(\mathbf{x})) = \boldsymbol{\lambda}_{R_k,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R_k,\mathbf{x},\mathbf{a}'}.$$

If $\mathbf{a} \neq \mathbf{a}'$, this quantity is zero by ♠2, so (C2) is satisfied. Finally,

$$\sum_{\mathbf{a} \in A^k} (E_{\mathbf{a}} * \xi(\mathbf{x}))^T (E_{\mathbf{a}} * \xi(\mathbf{x})) = \sum_{\mathbf{a} \in A^k} \|\boldsymbol{\lambda}_{R_k,\mathbf{x},\mathbf{a}}\|^2 = 1$$

by ♠1, so (C3) is also satisfied. Therefore, $\xi(\mathbf{x}) \in \mathscr{S}^{(n^k)}$. To show that ξ is in fact a homomorphism, take $R \in \sigma$ of arity r and $\mathbf{x} \in R^{\mathbf{X}}$, so $\mathbf{x}^{ⓚ} \in R^{\mathbf{X}^{ⓚ}}$. We need to show that $\xi(\mathbf{x}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{S}}(\mathbf{A}^{ⓚ})}$. Consider the matrix $Q \in \mathcal{T}^{|R^{\mathbf{A}}|, \aleph_0}(\mathbb{R})$ defined by

$$Q^T \mathbf{e}_{\mathbf{a}} = \begin{bmatrix} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \\ \mathbf{0}_{\aleph_0} \end{bmatrix} \qquad \mathbf{a} \in R^{\mathbf{A}}.$$

Using the same arguments as above, we check that Q satisfies (C1) and that $\mathbf{e}_{\mathbf{a}}^T QQ^T\mathbf{e}_{\mathbf{a}'} = \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'}$, so (C2) follows from ♠2 and (C3) from ♠1. Therefore, $Q \in \mathscr{S}^{(|R^{\mathbf{A}}|)}$. We now claim that $\xi(\mathbf{x}_{\mathbf{i}}) = Q_{/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$. Indeed, observe that, for each $\mathbf{a} \in A^k$,

$$E_{\mathbf{a}} * Q_{/\pi_{\mathbf{i}}} = E_{\mathbf{a}} * P_{\mathbf{i}} * Q = \Big(\sum_{\substack{\mathbf{b}\in R^{\mathbf{A}}\\ \mathbf{b}_{\mathbf{i}}=\mathbf{a}}} E_{\mathbf{b}}\Big) * Q = Q^T\Big(\sum_{\substack{\mathbf{b}\in R^{\mathbf{A}}\\ \mathbf{b}_{\mathbf{i}}=\mathbf{a}}} \mathbf{e}_{\mathbf{b}}\Big) = \sum_{\substack{\mathbf{b}\in R^{\mathbf{A}}\\ \mathbf{b}_{\mathbf{i}}=\mathbf{a}}} Q^T\mathbf{e}_{\mathbf{b}}$$

$$= \begin{bmatrix} \sum_{\substack{\mathbf{b}\in R^{\mathbf{A}}\\ \mathbf{b}_{\mathbf{i}}=\mathbf{a}}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{b}} \\ \mathbf{0}_{\aleph_0} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\lambda}_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{a}} \\ \mathbf{0}_{\aleph_0} \end{bmatrix} = E_{\mathbf{a}} * \xi(\mathbf{x}_{\mathbf{i}}),$$

where the second and sixth equalities are obtained using Lemma 31 and ♠3, respectively. It follows that the claim is true, so $\xi(\mathbf{x}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{S}}(\mathbf{A}^{ⓚ})}$, which concludes the proof that ξ is a homomorphism and that $\text{Test}^k_{\mathscr{S}}(\mathbf{X},\mathbf{A}) = \text{YES}$.

Conversely, let $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{S}}(\mathbf{A}^{ⓚ})$ be a homomorphism witnessing that $\text{Test}^k_{\mathscr{S}}(\mathbf{X},\mathbf{A}) = \text{YES}$. Take $R \in \sigma$ of arity r and $\mathbf{x} \in R^{\mathbf{X}}$. We have that $\mathbf{x}^{ⓚ} \in R^{\mathbf{X}^{ⓚ}}$, so $\xi(\mathbf{x}^{ⓚ}) \in R^{\mathbb{F}_{\mathscr{S}}(\mathbf{A}^{ⓚ})}$ since ξ is a homomorphism. As a consequence, we can fix a matrix $Q_{R,\mathbf{x}} \in \mathscr{S}^{(|R^{\mathbf{A}}|)}$ satisfying $\xi(\mathbf{x}_{\mathbf{i}}) = Q_{R,\mathbf{x}/\pi_{\mathbf{i}}}$ for each $\mathbf{i} \in [r]^k$. Consider the set $S = \{Q^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} : R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}\}$ and the vector space $\mathcal{U} = \text{span}(S) \subseteq \mathbb{R}^{\aleph_0}$, and observe that $\dim(\mathcal{U}) \le |S| \le \sum_{R\in\sigma} |R^{\mathbf{X}}| \cdot |R^{\mathbf{A}}| = \gamma$. Choose a vector space $\mathcal{V}$ of dimension γ such that $\mathcal{U} \subseteq \mathcal{V} \subseteq \mathbb{R}^{\aleph_0}$. Using the Gram–Schmidt process, we find a projection matrix $Z \in \mathbb{R}^{\aleph_0,\gamma}$ such that $Z^TZ = I_\gamma$ and $ZZ^T\mathbf{v} = \mathbf{v}$ for any $\mathbf{v} \in \mathcal{V}$. Consider the family of vectors

$$\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} = Z^TQ^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} \in \mathbb{R}^\gamma \qquad R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}. \tag{11}$$

We claim that (11) witnesses that $\text{SoS}^k(\mathbf{X},\mathbf{A}) = \text{YES}$. Take $R \in \sigma$ of arity r and $\mathbf{x} \in R^{\mathbf{X}}$. Recall from Proposition 39 that $\mathscr{S}$ is a conic minion. Using either Lemma 34 or Lemma 42 (depending on whether $k \ge \text{armax}(\sigma)$ or $k \ge 2$), given $\mathbf{a} \in R^{\mathbf{A}}$ such that $\mathbf{x} \not\prec \mathbf{a}$, we find

$$\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} = Z^TQ^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} = Z^T\mathbf{0}_{\aleph_0} = \mathbf{0}_\gamma,$$

so ♠4 holds. ♠1 follows from

$$\sum_{\mathbf{a}\in R^{\mathbf{A}}} \|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \sum_{\mathbf{a}\in R^{\mathbf{A}}} \mathbf{e}_{\mathbf{a}}^TQ_{R,\mathbf{x}}ZZ^TQ^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} = \sum_{\mathbf{a}\in R^{\mathbf{A}}} \mathbf{e}_{\mathbf{a}}^TQ_{R,\mathbf{x}}Q^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} = \text{tr}(Q_{R,\mathbf{x}}Q^T_{R,\mathbf{x}}) = 1,$$

where the second equality is true since $Q^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} \in S \subseteq \mathcal{U} \subseteq \mathcal{V}$ and the fourth follows from (C3). Similarly, using (C2), we find that, if $\mathbf{a} \ne \mathbf{a}' \in R^{\mathbf{A}}$,

$$\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'} = \mathbf{e}_{\mathbf{a}}^TQ_{R,\mathbf{x}}ZZ^TQ^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}'} = \mathbf{e}_{\mathbf{a}}^TQ_{R,\mathbf{x}}Q^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}'} = 0,$$

so ♠2 holds. We now show that $Q_{R_k,\mathbf{y}} = \xi(\mathbf{y})$ for any $\mathbf{y} \in X^k$. Indeed, using the same argument as in (9), letting $\mathbf{j} = (1,\dots,k) \in [k]^k$, we have

$$Q_{R_k,\mathbf{y}} = \Pi_{\mathbf{j}} \overset{k}{*} Q_{R_k,\mathbf{y}} = Q_{R_k,\mathbf{y}/\pi_{\mathbf{j}}} = \xi(\mathbf{y}_{\mathbf{j}}) = \xi(\mathbf{y}). \tag{12}$$

Given $\mathbf{i} \in [r]^k$ and $\mathbf{b} \in A^k = R_k^{\mathbf{A}}$, we have

$$\sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ \mathbf{a}_{\mathbf{i}}=\mathbf{b}}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} = \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ \mathbf{a}_{\mathbf{i}}=\mathbf{b}}} Z^TQ^T_{R,\mathbf{x}}\mathbf{e}_{\mathbf{a}} = \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ \mathbf{a}_{\mathbf{i}}=\mathbf{b}}} E_{\mathbf{a}} * Q_{R,\mathbf{x}} * Z = E_{\mathbf{b}} * P_{\mathbf{i}} * Q_{R,\mathbf{x}} * Z = E_{\mathbf{b}} * Q_{R,\mathbf{x}/\pi_{\mathbf{i}}} * Z$$

$$= E_{\mathbf{b}} * \xi(\mathbf{x}_{\mathbf{i}}) * Z = E_{\mathbf{b}} * Q_{R_k,\mathbf{x}_{\mathbf{i}}} * Z = Z^TQ^T_{R_k,\mathbf{x}_{\mathbf{i}}}\mathbf{e}_{\mathbf{b}} = \boldsymbol{\lambda}_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{b}},$$

where the third and sixth equalities follow from Lemma 31 and (12), respectively. This shows that ♠3 holds, too, so that (11) yields a solution for $\mathrm{SoS}^k(\mathbf{X}, \mathbf{A})$, as claimed. □

Proposition 57. *Let $k \in \mathbb{N}$ and let $\mathbf{X}, \mathbf{A}$ be k-enhanced σ-structures such that $k \geq \operatorname{armax}(\sigma)$. Then $\mathrm{BA}^k(\mathbf{X}, \mathbf{A}) = \mathrm{Test}^k_{\mathscr{M}_{\mathrm{BA}}}(\mathbf{X}, \mathbf{A})$.*

Proof. Recall from Section 2.3 that $\mathrm{BA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$ is equivalent to the existence of a rational nonnegative solution (denoted by the superscript (B)) and an integer solution (denoted by the superscript (A)) to the system (♣), such that

$$\lambda^{(\mathrm{B})}_{R,\mathbf{x},\mathbf{a}} = 0 \quad \Rightarrow \quad \lambda^{(\mathrm{A})}_{R,\mathbf{x},\mathbf{a}} = 0 \tag{13}$$

for each $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$. Note that requiring (13) for each $R \in \sigma$ is equivalent to only requiring it for $R = R_k$. Indeed, take some $R \in \sigma$ of arity r, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$, and consider the tuple $\mathbf{i} = (1, 2, \ldots, r, 1, 1, \ldots, 1) \in [r]^k$, which is well defined as $k \geq r$. Noting that $\{\mathbf{b} \in R^{\mathbf{A}} : \mathbf{b}_{\mathbf{i}} = \mathbf{a}_{\mathbf{i}}\} = \{\mathbf{a}\}$, we find from ♣2 that

$$\lambda^{(\mathrm{B})}_{R,\mathbf{x},\mathbf{a}} = \sum_{\substack{\mathbf{b} \in R^{\mathbf{A}} \\ \mathbf{b}_{\mathbf{i}} = \mathbf{a}_{\mathbf{i}}}} \lambda^{(\mathrm{B})}_{R,\mathbf{x},\mathbf{b}} = \lambda^{(\mathrm{B})}_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{a}_{\mathbf{i}}}$$

and, similarly, $\lambda^{(\mathrm{A})}_{R,\mathbf{x},\mathbf{a}} = \lambda^{(\mathrm{A})}_{R_k,\mathbf{x}_{\mathbf{i}},\mathbf{a}_{\mathbf{i}}}$. As a consequence, if (13) holds for R_k, it also holds for R. Therefore, it follows from Propositions 51 and 55 that $\mathrm{BA}^k(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$ is equivalent to the existence of homomorphisms $\xi : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{Q}_{\mathrm{conv}}}(\mathbf{A}^{ⓚ})$ and $\zeta : \mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{Z}_{\mathrm{aff}}}(\mathbf{A}^{ⓚ})$ such that $\operatorname{supp}(\zeta(\mathbf{x})) \subseteq \operatorname{supp}(\xi(\mathbf{x}))$ for each $\mathbf{x} \in X^k$. By virtue of Proposition 49, this happens precisely when $\mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{Q}_{\mathrm{conv}} \ltimes \mathscr{Z}_{\mathrm{aff}}}(\mathbf{A}^{ⓚ})$. Since $\mathscr{M}_{\mathrm{BA}} = \mathscr{Q}_{\mathrm{conv}} \ltimes \mathscr{Z}_{\mathrm{aff}}$ (cf. Example 48), this is equivalent to $\mathbf{X}^{ⓚ} \to \mathbb{F}_{\mathscr{M}_{\mathrm{BA}}}(\mathbf{A}^{ⓚ})$; i.e., to $\mathrm{Test}^k_{\mathscr{M}_{\mathrm{BA}}}(\mathbf{X}, \mathbf{A}) = \textsc{Yes}$. □

Remark 58. The characterisations of SA^k and SoS^k in Propositions 51 and 56 hold for any higher level than the first, unlike the characterisation of AIP^k in Proposition 55. This is due to the fact that $\mathscr{Q}_{\mathrm{conv}}$ and $\mathscr{S}$ are conic minions, so Lemma 42 applies, while $\mathscr{Z}_{\mathrm{aff}}$ is not. As for BW^k, Proposition 53 requires $k \geq \operatorname{armax}(\sigma)$ even if $\mathscr{H}$ is a conic minion. The reason for this lies in the definition of the bounded width hierarchy. Essentially, any constraint whose scope has more than k distinct variables does not appear among the constraints of the partial homomorphisms witnessing acceptance of BW^k, while it does appear in the requirements of SA^k and SoS^k. Finally, assuming $k \geq \operatorname{armax}(\sigma)$ is also required in the characterisation of BA^k in Proposition 57, in order to make use of Proposition 49 and of the characterisation of AIP^k.

Remark 59. As it was shown in Section 7, hierarchies of relaxations built on conic minions (such as BW^k, SA^k, SoS^k, and BA^k) are "sound in the limit", in that their k-th level correctly classifies instances $\mathbf{X}$ with $|X| \leq k$ (cf. Proposition 44). This is not the case for the non-conic hierarchy AIP^k, as it was established in the follow-up work [41]. In [18], a stronger affine hierarchy was proposed, which – contrary to AIP^k – requires that the variables in the relaxation should be partial homomorphisms and is thus sound in the limit. By virtue of Proposition 43, this requirement can be captured by taking the semi-direct product of any conic minion and $\mathscr{Z}_{\mathrm{aff}}$. In particular, it follows that the hierarchy in [18] is not stronger than the hierarchy built on the minion $\mathscr{H} \ltimes \mathscr{Z}_{\mathrm{aff}}$ (cf. Remark 47). In recent work [46], a different algorithm for (P)CSPs has been proposed. The relationship of [46] with our work is an interesting direction for future research.

Acknowledgements

The authors are grateful to Marcin Kozik, Jakub Opršal, and Caterina Viola for useful discussions on SDP relaxations. We also thank the anonymous reviewers of this paper, and its extended abstract [44].

A Notes on relaxations and hierarchies

In this appendix, we discuss some basic properties and alternative formulations of the relaxations, and hierarchies thereof, presented in Section 2.3.

A.1 $\mathrm{SA}^{\mathbf{k}}$

The hierarchy defining SA^k given by the system (♣) slightly differs from the one described in [36]. For completeness, we report below the hierarchy in [36] and show that it is equivalent to the one adopted in this work.

Given two σ-structures $\mathbf{X}, \mathbf{A}$, introduce a variable $\mu_V(f)$ for every subset $V \subseteq X$ with $1 \leq |V| \leq k$ and every function $f : V \to A$, and a variable $\mu_{R,\mathbf{x}}(f)$ for every $R \in \sigma$, every $\mathbf{x} \in R^{\mathbf{X}}$, and every $f : \{\mathbf{x}\} \to A$. The k-th level of the hierarchy defined in [36] is given by the following constraints:

$$
\begin{array}{lll}
(\heartsuit 1) & \displaystyle\sum_{f:V\to A} \mu_V(f) = 1 & V \subseteq X \text{ such that } 1 \leq |V| \leq k \\
(\heartsuit 2) & \displaystyle\mu_U(f) = \sum_{g:V\to A,\, g|_U=f} \mu_V(g) & U \subseteq V \subseteq X \text{ such that } 1 \leq |V| \leq k, U \neq \emptyset, f : U \to A \\
(\heartsuit 3) & \displaystyle\mu_U(f) = \sum_{g:\{\mathbf{x}\}\to A,\, g|_U=f} \mu_{R,\mathbf{x}}(g) & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, U \subseteq \{\mathbf{x}\} \text{ such that } 1 \leq |U| \leq k, f : U \to A \\
(\heartsuit 4) & \mu_{R,\mathbf{x}}(f) = 0 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, f : \{\mathbf{x}\} \to A \text{ such that } f(\mathbf{x}) \notin R^{\mathbf{A}}.
\end{array}
\qquad (\heartsuit)
$$

Lemma 60. *Let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be two σ-structures, and let $\tilde{\mathbf{X}}$ (resp. $\tilde{\mathbf{A}}$) be the structure obtained from $\mathbf{X}$ (resp. $\mathbf{A}$) by adding the relation $R_k^{\tilde{\mathbf{X}}} = X^k$ (resp. $R_k^{\tilde{\mathbf{A}}} = A^k$). Then the system (♥) applied to $\mathbf{X}$ and $\mathbf{A}$ is equivalent to the system (♣) applied to $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{A}}$.*

Proof. Let λ be a solution to (♣) applied to $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{A}}$. Given $V \subseteq X$ with $1 \leq |V| \leq k$ and $f : V \to A$, let $\mathbf{x} \in X^k$ be such that $V = \{\mathbf{x}\}$ and set $\mu_V(f) = \lambda_{R_k,\mathbf{x},f(\mathbf{x})}$. We claim that this assignment does not depend on the choice of $\mathbf{x}$; i.e., we claim that $\lambda_{R_k,\mathbf{x},f(\mathbf{x})} = \lambda_{R_k,\mathbf{y},f(\mathbf{y})}$ whenever $\mathbf{x}, \mathbf{y} \in X^k$ are such that $\{\mathbf{x}\} = \{\mathbf{y}\}$. The latter condition implies that $\mathbf{x} = \mathbf{y_i}$ and $\mathbf{y} = \mathbf{x_j}$ for some $\mathbf{i}, \mathbf{j} \in [k]^k$. Using ♣2 and ♣3, we find

$$
\lambda_{R_k,\mathbf{y},f(\mathbf{y})} = \lambda_{R_k,\mathbf{x_j},f(\mathbf{x_j})} = \sum_{\substack{\mathbf{a}\in A^k \\ \mathbf{a_j}=f(\mathbf{x_j})}} \lambda_{R_k,\mathbf{x},\mathbf{a}} = \sum_{\substack{\mathbf{a}\in A^k \\ \mathbf{a_j}=f(\mathbf{x_j}) \\ \mathbf{x}\prec\mathbf{a}}} \lambda_{R_k,\mathbf{x},\mathbf{a}} = \lambda_{R_k,\mathbf{x},f(\mathbf{x})} + \sum_{\substack{\mathbf{a}\in A^k \\ \mathbf{a_j}=f(\mathbf{x_j}) \\ \mathbf{x}\prec\mathbf{a} \\ \mathbf{a}\neq f(\mathbf{x})}} \lambda_{R_k,\mathbf{x},\mathbf{a}}.
$$

The claim then follows if we show that there is no $\mathbf{a} \in A^k$ such that $\mathbf{a_j} = f(\mathbf{x_j})$, $\mathbf{x} \prec \mathbf{a}$, and $\mathbf{a} \neq f(\mathbf{x})$. If such $\mathbf{a}$ exists, using that $\mathbf{x} = \mathbf{x_{j_i}}$, we find that for some $p \in [k]$

$$
a_p \neq f(x_p) = f(x_{j_{i_p}}) = a_{j_{i_p}}.
$$

Since $\mathbf{x} \prec \mathbf{a}$, this implies that $x_p \neq x_{j_{i_p}}$, a contradiction. Therefore, the claim is true. Additionally, given $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $f : \{\mathbf{x}\} \to A$, we set $\mu_{R,\mathbf{x}}(f) = \lambda_{R,\mathbf{x},f(\mathbf{x})}$ if $f(\mathbf{x}) \in R^{\mathbf{A}}$, $\mu_{R,\mathbf{x}}(f) = 0$ otherwise. It is straightforward to check that μ satisfies all constraints in the system (♥) applied to $\mathbf{X}$ and $\mathbf{A}$.

Conversely, let μ be a solution to (♥) applied to $\mathbf{X}$ and $\mathbf{A}$. As in the proof of Proposition 53, given two sets S, T, an integer $p \in \mathbb{N}$, and two tuples $\mathbf{s} \in S^p, \mathbf{t} \in T^p$ such that $\mathbf{s} \prec \mathbf{t}$, we define the map $f_{\mathbf{s},\mathbf{t}} : \{\mathbf{s}\} \to T$ by $f_{\mathbf{s},\mathbf{t}}(s_\alpha) = t_\alpha$ for each $\alpha \in [p]$. For every $R \in \sigma$, $\mathbf{x} \in R^{\mathbf{X}}$, and $\mathbf{a} \in R^{\mathbf{A}}$, we set $\lambda_{R,\mathbf{x},\mathbf{a}} = \mu_{R,\mathbf{x}}(f_{\mathbf{x},\mathbf{a}})$ if $\mathbf{x} \prec \mathbf{a}$, $\lambda_{R,\mathbf{x},\mathbf{a}} = 0$ otherwise. Additionally, for every $\mathbf{x} \in X^k = R_k^{\tilde{\mathbf{X}}}$ and $\mathbf{a} \in A^k = R_k^{\tilde{\mathbf{A}}}$, we set $\lambda_{R_k,\mathbf{x},\mathbf{a}} = \mu_{\{\mathbf{x}\}}(f_{\mathbf{x},\mathbf{a}})$ if $\mathbf{x} \prec \mathbf{a}$, $\lambda_{R_k,\mathbf{x},\mathbf{a}} = 0$ otherwise. It is easily verified that λ yields a solution to (♣) applied to $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{A}}$. □

We also note that [5] has yet another definition of the Sherali–Adams hierarchy. However, it was shown in [36, Appendix A] that the hierarchy given in [5] interleaves with the one in [36] and, by virtue of Lemma 60, with the hierarchy used in this work. In particular, the class of PCSPs solved by constant levels of the hierarchy is the same for all definitions.

A.2 SDP

The relaxation defined by (♦) is not in semidefinite programming form, because of the constraint ♦4. However, it can be easily translated into a semidefinite program by introducing γ additional variables $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_\gamma$ taking values in $\mathbb{R}^\gamma$, and requiring that the following constraints are met:

$$
\begin{array}{lll}
(\blacklozenge 4') & \boldsymbol{\mu}_p \cdot \boldsymbol{\mu}_q = \delta_{p,q} & p, q \in [\gamma] \\
(\blacklozenge 4'') & \displaystyle\sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ a_i = a}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\mu}_p = \boldsymbol{\lambda}_{x_i,a} \cdot \boldsymbol{\mu}_p & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, a \in A, i \in [\mathrm{ar}(R)], p \in [\gamma]
\end{array}
$$

where $\delta_{p,q}$ is the Kronecker delta. One easily checks that the requirements ♦4′ and ♦4″ are together equivalent to the requirement ♦4, and they are expressed in semidefinite programming form. Next, we give a proof for the following basic result.

Proposition (Proposition 16 restated)**.** *Let* $\mathbf{X}, \mathbf{A}$ *be two* σ*-structures. The system* (♦) *implies the following facts:*

$$
\begin{array}{lll}
(i) & \displaystyle\|\sum_{a \in A} \boldsymbol{\lambda}_{x,a}\|^2 = 1 & x \in X; \\
(ii) & \displaystyle\sum_{\mathbf{a} \in R^{\mathbf{A}}} \|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \|\sum_{\mathbf{a} \in R^{\mathbf{A}}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = 1 & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}; \\
(iii) & \displaystyle\sum_{\substack{\mathbf{a} \in R^{\mathbf{A}} \\ a_i = a,\, a_j = a'}} \|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \boldsymbol{\lambda}_{x_i,a} \cdot \boldsymbol{\lambda}_{x_j,a'} & R \in \sigma, \mathbf{x} \in R^{\mathbf{X}}, a, a' \in A, i, j \in [\mathrm{ar}(R)].
\end{array}
$$

If, in addition, $\mathbf{X}$ *and* $\mathbf{A}$ *are* 2*-enhanced,*

$$
\begin{array}{lll}
(iv) & \displaystyle\sum_{a \in A} \boldsymbol{\lambda}_{x,a} = \sum_{a \in A} \boldsymbol{\lambda}_{x',a} & x, x' \in X.
\end{array}
$$

Proof. (*i*) We have

$$
\|\sum_{a \in A} \boldsymbol{\lambda}_{x,a}\|^2 = \left(\sum_{a \in A} \boldsymbol{\lambda}_{x,a}\right) \cdot \left(\sum_{a' \in A} \boldsymbol{\lambda}_{x,a'}\right) = \sum_{a,a' \in A} \boldsymbol{\lambda}_{x,a} \cdot \boldsymbol{\lambda}_{x,a'} = \sum_{a \in A} \|\boldsymbol{\lambda}_{x,a}\|^2 = 1,
$$

where the third equality comes from ♦2 and the fourth from ♦1.

(*ii*) We have

$$\sum_{\mathbf{a}\in R^{\mathbf{A}}} \|\lambda_{R,\mathbf{x},\mathbf{a}}\|^2 = \sum_{\mathbf{a},\mathbf{a}'\in R^{\mathbf{A}}} \lambda_{R,\mathbf{x},\mathbf{a}} \cdot \lambda_{R,\mathbf{x},\mathbf{a}'} = \left(\sum_{\mathbf{a}\in R^{\mathbf{A}}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\right) \cdot \left(\sum_{\mathbf{a}'\in R^{\mathbf{A}}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'}\right)$$
$$= \|\sum_{\mathbf{a}\in R^{\mathbf{A}}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \|\sum_{a\in A}\sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_1=a}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2 = \|\sum_{a\in A} \boldsymbol{\lambda}_{x_1,a}\|^2 = 1,$$

where the first equality comes from ♦3, the fifth from ♦4, and the sixth from part (*i*) of this proposition.

(*iii*) We have

$$\boldsymbol{\lambda}_{x_i,a} \cdot \boldsymbol{\lambda}_{x_j,a'} = \left(\sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_i=a}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\right) \cdot \left(\sum_{\substack{\mathbf{a}'\in R^{\mathbf{A}}\\ a'_j=a'}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'}\right) = \sum_{\substack{\mathbf{a},\mathbf{a}'\in R^{\mathbf{A}}\\ a_i=a,\ a'_j=a'}} \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}} \cdot \boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}'}$$
$$= \sum_{\substack{\mathbf{a}\in R^{\mathbf{A}}\\ a_i=a,\ a_j=a'}} \|\boldsymbol{\lambda}_{R,\mathbf{x},\mathbf{a}}\|^2,$$

where the first equality comes from ♦4 and the third from ♦3.

(*iv*) If $\mathbf{X}$ and $\mathbf{A}$ are 2-enhanced, we have

$$\sum_{a\in A} \boldsymbol{\lambda}_{x,a} = \sum_{a\in A}\sum_{\substack{\mathbf{a}\in R_2^{\mathbf{A}}\\ a_1=a}} \boldsymbol{\lambda}_{R_2,(x,x'),\mathbf{a}} = \sum_{\mathbf{a}\in R_2^{\mathbf{A}}} \boldsymbol{\lambda}_{R_2,(x,x'),\mathbf{a}} = \sum_{a\in A}\sum_{\substack{\mathbf{a}\in R_2^{\mathbf{A}}\\ a_2=a}} \boldsymbol{\lambda}_{R_2,(x,x'),\mathbf{a}} = \sum_{a\in A} \boldsymbol{\lambda}_{x',a},$$

where the first and fourth equalities come from ♦4. □

We point out that slightly different versions of the "standard SDP relaxation" appeared in the literature on CSPs, some of which use parts (*i*) through (*iii*) of Proposition 16 as constraints defining the relaxation. In particular, certain versions require that the scalar products $\boldsymbol{\lambda}_{x,a} \cdot \boldsymbol{\lambda}_{y,b}$ should be nonnegative for all choices of $x, y \in X$ and $a, b \in A$. For example, this is the case of the SDP relaxation used in [14]. It follows from Proposition 61, proved in Appendix A.3, that one can enforce nonnegativity of the scalar products by taking the second level of the SoS hierarchy of the SDP relaxation as defined in this work.

A.3 $\mathrm{SoS}^{\mathbf{k}}$

First, we note that the relaxation defined by (♠) can be easily translated into a semidefinite program through the procedure described at the beginning of Appendix A.2. Next, we show that the $2k$-th level of the SoS hierarchy enforces additional constraints – in particular, nonnegativity of the scalar products of the SoS vectors – on the vectors corresponding to the k-th level.

Proposition 61. *Let $k \in \mathbb{N}$, let $\mathbf{X}, \mathbf{A}$ be $2k$-enhanced σ-structures, suppose that $\mathrm{SoS}^{2k}(\mathbf{X}, \mathbf{A}) =$ YES, and let $\boldsymbol{\lambda}$ denote a solution. Then $\boldsymbol{\lambda}$ satisfies the following additional constraints:*

(i)	$\boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{x}),(\mathbf{a},\mathbf{a})} \cdot \boldsymbol{\lambda}_{R_{2k},(\mathbf{y},\mathbf{y}),(\mathbf{b},\mathbf{b})} \geq 0$	$\mathbf{x}, \mathbf{y} \in X^k, \mathbf{a}, \mathbf{b} \in A^k$
(ii)	$\boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{x}),(\mathbf{a},\mathbf{a})} \cdot \boldsymbol{\lambda}_{R_{2k},(\mathbf{y},\mathbf{y}),(\mathbf{b},\mathbf{b})} = 0$	$\mathbf{x}, \mathbf{y} \in X^k, \mathbf{a}, \mathbf{b} \in A^k$, $\mathbf{a}_\mathbf{i} \neq \mathbf{b}_\mathbf{j}$ *for some* $\mathbf{i}, \mathbf{j} \in [k]^k$ *such that* $\mathbf{x}_\mathbf{i} = \mathbf{y}_\mathbf{j}$
(iii)	$\boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{x}),(\mathbf{a},\mathbf{a})} \cdot \boldsymbol{\lambda}_{R_{2k},(\mathbf{y},\mathbf{y}),(\mathbf{b},\mathbf{b})} = \boldsymbol{\lambda}_{R_{2k},(\hat{\mathbf{x}},\hat{\mathbf{x}}),(\hat{\mathbf{a}},\hat{\mathbf{a}})} \cdot \boldsymbol{\lambda}_{R_{2k},(\hat{\mathbf{y}},\hat{\mathbf{y}}),(\hat{\mathbf{b}},\hat{\mathbf{b}})}$	$\mathbf{x}, \hat{\mathbf{x}}, \mathbf{y}, \hat{\mathbf{y}} \in X^k, \mathbf{a}, \hat{\mathbf{a}}, \mathbf{b}, \hat{\mathbf{b}} \in A^k$, $(\hat{\mathbf{x}}, \hat{\mathbf{y}})_{\boldsymbol{\ell}} = (\mathbf{x}, \mathbf{y}), (\hat{\mathbf{a}}, \hat{\mathbf{b}})_{\boldsymbol{\ell}} = (\mathbf{a}, \mathbf{b})$ *for some* $\boldsymbol{\ell} \in [2k]^{2k}$ *such that* $\lvert\{\boldsymbol{\ell}\}\rvert = 2k$.

Proof. Observe that, for $\mathbf{x}, \mathbf{y} \in X^k$ and $\mathbf{a}, \mathbf{b} \in A^k$,

$$\begin{aligned}
\boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{x}),(\mathbf{a},\mathbf{a})} \cdot \boldsymbol{\lambda}_{R_{2k},(\mathbf{y},\mathbf{y}),(\mathbf{b},\mathbf{b})} &= \left(\sum_{\mathbf{c} \in A^k} \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{a},\mathbf{c})}\right) \cdot \left(\sum_{\mathbf{c}' \in A^k} \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{c}',\mathbf{b})}\right) \\
&= \sum_{\mathbf{c},\mathbf{c}' \in A^k} \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{a},\mathbf{c})} \cdot \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{c}',\mathbf{b})} \\
&= \lVert \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{a},\mathbf{b})} \rVert^2, \qquad (14)
\end{aligned}$$

where the first and third equalities come from ♠3 and ♠2, respectively. Hence, (i) holds. If, in addition, $\mathbf{a}_\mathbf{i} \neq \mathbf{b}_\mathbf{j}$ for some $\mathbf{i}, \mathbf{j} \in [k]^k$ such that $\mathbf{x}_\mathbf{i} = \mathbf{y}_\mathbf{j}$, we deduce that $(\mathbf{x}_\mathbf{i}, \mathbf{y}_\mathbf{j}) \not\preceq (\mathbf{a}_\mathbf{i}, \mathbf{b}_\mathbf{j})$ and, therefore,

$$\begin{aligned}
0 = \lVert \boldsymbol{\lambda}_{R_{2k},(\mathbf{x}_\mathbf{i},\mathbf{y}_\mathbf{j}),(\mathbf{a}_\mathbf{i},\mathbf{b}_\mathbf{j})} \rVert^2 &= \lVert \sum_{\substack{(\mathbf{c},\mathbf{d}) \in A^{2k} \\ \mathbf{c}_\mathbf{i} = \mathbf{a}_\mathbf{i}, \mathbf{d}_\mathbf{j} = \mathbf{b}_\mathbf{j}}} \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{c},\mathbf{d})} \rVert^2 \\
&= \sum_{\substack{(\mathbf{c},\mathbf{d}) \in A^{2k} \\ \mathbf{c}_\mathbf{i} = \mathbf{a}_\mathbf{i}, \mathbf{d}_\mathbf{j} = \mathbf{b}_\mathbf{j}}} \lVert \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{c},\mathbf{d})} \rVert^2 \geq \lVert \boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{a},\mathbf{b})} \rVert^2, \qquad (15)
\end{aligned}$$

where the first, second, and third equalities come from ♠4, ♠3, and ♠2, respectively. Combining (14) and (15), we obtain (ii). Suppose now that $\mathbf{x}, \hat{\mathbf{x}}, \mathbf{y}, \hat{\mathbf{y}} \in X^k$ and $\mathbf{a}, \hat{\mathbf{a}}, \mathbf{b}, \hat{\mathbf{b}} \in A^k$ are such that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})_{\boldsymbol{\ell}} = (\mathbf{x}, \mathbf{y})$ and $(\hat{\mathbf{a}}, \hat{\mathbf{b}})_{\boldsymbol{\ell}} = (\mathbf{a}, \mathbf{b})$ for some $\boldsymbol{\ell} \in [2k]^{2k}$ such that $\lvert\{\boldsymbol{\ell}\}\rvert = 2k$. Using ♠3, we find

$$\boldsymbol{\lambda}_{R_{2k},(\mathbf{x},\mathbf{y}),(\mathbf{a},\mathbf{b})} = \boldsymbol{\lambda}_{R_{2k},(\hat{\mathbf{x}},\hat{\mathbf{y}})_{\boldsymbol{\ell}},(\hat{\mathbf{a}},\hat{\mathbf{b}})_{\boldsymbol{\ell}}} = \sum_{\substack{(\mathbf{c},\mathbf{d}) \in A^{2k} \\ (\mathbf{c},\mathbf{d})_{\boldsymbol{\ell}} = (\hat{\mathbf{a}},\hat{\mathbf{b}})_{\boldsymbol{\ell}}}} \boldsymbol{\lambda}_{R_{2k},(\hat{\mathbf{x}},\hat{\mathbf{y}}),(\mathbf{c},\mathbf{d})} = \boldsymbol{\lambda}_{R_{2k},(\hat{\mathbf{x}},\hat{\mathbf{y}}),(\hat{\mathbf{a}},\hat{\mathbf{b}})},$$

where the last equality is due to the fact that $\lvert\{\boldsymbol{\ell}\}\rvert = 2k$. Hence, (iii) follows from (14). □

We observe that the relaxation in (♠) is formally different from the one described in [99]. However, it can be shown that the $2k$-th level of the hierarchy as defined here is at least as tight as the k-th level of the hierarchy as defined in [99]. Let us denote the two relaxations by SoS and SoS′, respectively. First of all, each variable in SoS′ corresponds to a subset S of X and an assignment $f : S \to A$, while in SoS the variables correspond to pairs of tuples $\mathbf{x} \in R^{\mathbf{X}}, \mathbf{a} \in R^{\mathbf{A}}$. This is an inessential difference, as one can check through the same argument used to prove Lemma 60 – in particular, ♠4 ensures that the only variables having nonzero weight are those corresponding to well-defined assignments (cf. Footnote 7). The k-th level of SoS′ contains constraints that, in our language, are expressed as ♠4, ♠1, and parts $(i), (ii), (iii)$ of Proposition 61. By virtue of Proposition 61, therefore, any solution $\boldsymbol{\lambda}$ to the $2k$-th level of SoS yields a solution to the k-th level of SoS′ – which means that the $2k$-th level of SoS is at least as tight as the k-th level of SoS′.

References

[1] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009. doi:10.1145/1502793.1502794.

[2] Kristina Asimi and Libor Barto. Finitely tractable promise constraint satisfaction problems. In *Proc. 46th International Symposium on Mathematical Foundations of Computer Science (MFCS'21)*, volume 202 of *LIPIcs*, pages 11:1–11:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. arXiv:2010.04618, doi:10.4230/LIPIcs.MFCS.2021.11.

[3] Albert Atserias, Andrei A. Bulatov, and Víctor Dalmau. On the Power of k-Consistency. In *Proc. 34th International Colloquium on Automata, Languages and Programming (ICALP'07)*, volume 4596 of *Lecture Notes in Computer Science*, pages 279–290. Springer, 2007. doi:10.1007/978-3-540-73420-8_26.

[4] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *J. Comput. Syst. Sci.*, 74(3):323–334, 2008. doi:10.1016/j.jcss.2007.06.025.

[5] Albert Atserias and Víctor Dalmau. Promise Constraint Satisfaction and Width. In *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*, pages 1129–1153, 2022. arXiv:2107.05886, doi:10.1137/1.9781611977073.48.

[6] Albert Atserias and Elitza N. Maneva. Sherali-Adams Relaxations and Indistinguishability in Counting Logics. *SIAM J. Comput.*, 42(1):112–137, 2013. doi:10.1137/120867834.

[7] Per Austrin. Towards Sharp Inapproximability for Any 2-CSP. *SIAM J. Comput.*, 39(6):2430–2463, 2010. doi:10.1137/070711670.

[8] Per Austrin, Amey Bhangale, and Aditya Potukuchi. Improved inapproximability of rainbow coloring. In *Proc. 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*, pages 1479–1495, 2020. arXiv:1810.02784, doi:10.1137/1.9781611975994.90.

[9] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$-Sat is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017. doi:10.1137/15M1006507.

[10] Libor Barto. The collapse of the bounded width hierarchy. *J. Log. Comput.*, 26(3):923–943, 2016. doi:10.1093/logcom/exu070.

[11] Libor Barto, Diego Battistelli, and Kevin M. Berg. Symmetric Promise Constraint Satisfaction Problems: Beyond the Boolean Case. In *Proc. 38th International Symposium on Theoretical Aspects of Computer Science (STACS'21)*, volume 187 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. arXiv:2010.04623, doi:10.4230/LIPIcs.STACS.2021.10.

[12] Libor Barto, Jakub Bulín, Andrei A. Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–28:66, 2021. arXiv:1811.00970, doi:10.1145/3457606.

[13] Libor Barto and Marcin Kozik. Constraint Satisfaction Problems Solvable by Local Consistency Methods. *J. ACM*, 61(1), 2014. Article No. 3. doi:10.1145/2556646.

[14] Libor Barto and Marcin Kozik. Robustly solvable constraint satisfaction problems. *SIAM J. Comput.*, 45(4):1646–1669, 2016. arXiv:1512.01157, doi:10.1137/130915479.

[15] Libor Barto and Marcin Kozik. Combinatorial Gap Theorem and Reductions between Promise CSPs. In *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*, pages 1204–1220, 2022. arXiv:2107.09423, doi:10.1137/1.9781611977073.50.

[16] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In Andrei Krokhin and Stanislav Živný, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2017. doi:10.4230/DFU.Vol7.15301.1.

[17] Libor Barto, Jakub Opršal, and Michael Pinsker. The wonderland of reflections. *Isr. J. Math*, 223(1):363–398, Feb 2018. arXiv:1510.04521, doi:10.1007/s11856-017-1621-9.

[18] Christoph Berkholz and Martin Grohe. Linear Diophantine Equations, Group CSPs, and Graph Isomorphism. In *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'17)*, pages 327–339. SIAM, 2017. arXiv:1607.04287, doi:10.1137/1.9781611974782.21.

[19] Amey Bhangale and Subhash Khot. Optimal Inapproximability of Satisfiable k-LIN over Non-Abelian Groups. In *Proc. 53rd Annual ACM Symposium on Theory of Computing (STOC'21)*, pages 1615–1628. ACM, 2021. arXiv:2009.02815, doi:10.1145/3406325.3451003.

[20] Amey Bhangale, Subhash Khot, and Don Minzer. On Approximability of Satisfiable k-CSPs: I. In *Proc. 54th Annual ACM Symposium on Theory of Computing (STOC'22)*, pages 976–988. ACM, 2022. doi:10.1145/3519935.3520028.

[21] Joshua Brakensiek and Venkatesan Guruswami. New hardness results for graph and hypergraph colorings. In *Proc. 31st Conference on Computational Complexity (CCC'16)*, volume 50 of *LIPIcs*, pages 14:1–14:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.CCC.2016.14.

[22] Joshua Brakensiek and Venkatesan Guruswami. An algorithmic blend of LPs and ring equations for promise CSPs. In *Proc. 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*, pages 436–455, 2019. arXiv:1807.05194, doi:10.1137/1.9781611975482.28.

[23] Joshua Brakensiek and Venkatesan Guruswami. Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021. arXiv:1704.01937, doi:10.1137/19M128212X.

[24] Joshua Brakensiek and Venkatesan Guruswami. The quest for strong inapproximability results with perfect completeness. *ACM Trans. Algorithms*, 17(3):27:1–27:35, 2021. doi:10.1145/3459668.

[25] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. Conditional Dichotomy of Boolean Ordered Promise CSPs. *TheoretiCS*, 2, 2023. arXiv:2102.11854, doi:10.46298/theoretics.23.2.

[26] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. SDPs and Robust Satisfiability of Promise CSP. In *Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23)*, pages 609–622. ACM, 2023. arXiv:2211.08373, doi:10.1145/3564246.3585180.

[27] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic LP and affine relaxation for promise CSPs. *SIAM J. Comput.*, 49:1232–1248, 2020. arXiv:1907.04383, doi:10.1137/20M1312745.

[28] Alex Brandts, Marcin Wrochna, and Stanislav Živný. The complexity of promise SAT on non-Boolean domains. *ACM Trans. Comput. Theory*, 13(4):26:1–26:20, 2021. arXiv:1911.09065, doi:10.1145/3470867.

[29] Mark Braverman, Subhash Khot, Noam Lifshitz, and Dor Minzer. An Invariance Principle for the Multi-slice, with Applications. In *Proc. 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'21)*, pages 228–236. IEEE, 2021. arXiv:2110.10725, doi:10.1109/FOCS52979.2021.00030.

[30] Mark Braverman, Subhash Khot, and Dor Minzer. On rich 2-to-1 games. In *Proc. 12th Innovations in Theoretical Computer Science Conference (ITCS'21)*, volume 185 of *LIPIcs*, pages 27:1–27:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ITCS.2021.27.

[31] Andrei Bulatov. Bounded relational width. Unpublished manuscript, 2009. URL: https://www2.cs.sfu.ca/~abulatov/papers/relwidth.pdf.

[32] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005. doi:10.1137/S0097539700376676.

[33] Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *J. ACM*, 60(5):34, 2013. doi:10.1145/2528400.

[34] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *Proc. 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 319–330, 2017. arXiv:1703.03021, doi:10.1109/FOCS.2017.37.

[35] Andrei A. Bulatov, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, and Colin McQuillan. The expressibility of functions on the Boolean domain, with applications to Counting CSPs. *J. ACM*, 60(5):32, 2013. doi:10.1145/2528401.

[36] Silvia Butti and Víctor Dalmau. Fractional Homomorphism, Weisfeiler-Leman Invariance, and the Sherali-Adams Hierarchy for the Constraint Satisfaction Problem. In *Proc. 46th International Symposium on Mathematical Foundations of Computer Science (MFCS'21)*, volume 202 of *LIPIcs*, pages 27:1–27:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. arXiv:2107.02956, doi:10.4230/LIPIcs.MFCS.2021.27.

[37] Siu On Chan, James R. Lee, Prasad Raghavendra, and David Steurer. Approximate Constraint Satisfaction Requires Large LP Relaxations. *J. ACM*, 63(4):34:1–34:22, 2016. doi:10.1145/2811255.

[38] Siu On Chan and Hiu Tsun Ng. How random CSPs fool hierarchies: II. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1710–1721, 2025.

[39] Xi Chen, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. *J. Comput. Syst. Sci.*, 81(1):311–329, 2015. arXiv:1208.1783, doi:10.1016/j.jcss.2014.06.006.

[40] Lorenzo Ciardo and Stanislav Živný. The Sherali-Adams hierarchy for promise CSPs through tensors. 2022. arXiv:2203.02478.

[41] Lorenzo Ciardo and Stanislav Živný. Approximate graph colouring and crystals. In *Proc. 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23)*, pages 2256–2267, 2023. arXiv:2210.08293, doi:10.1137/1.9781611977554.ch86.

[42] Lorenzo Ciardo and Stanislav Živný. Approximate graph colouring and the hollow shadow. In *Proc. 55th Annual ACM Symposium on Theory of Computing (STOC'23)*. ACM, 2023. arXiv:2211.03168.

[43] Lorenzo Ciardo and Stanislav Živný. CLAP: A New Algorithm for Promise CSPs. *SIAM Journal on Computing*, 52(1):1–37, 2023. arXiv:2107.05018, doi:10.1137/22M1476435.

[44] Lorenzo Ciardo and Stanislav Živný. Hierarchies of minion tests for PCSPs through tensors. In *Proc. 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23)*, pages 568–580, 2023. arXiv:2207.02277, doi:10.1137/1.9781611977554.ch25.

[45] Lorenzo Ciardo and Stanislav Živný. The periodic structure of local consistency. *arXiv preprint arXiv:2406.19685*, 2024.

[46] Adam Ó Conghaile. Cohomology in Constraint Satisfaction and Structure Isomorphism. In *Proc. 47th International Symposium on Mathematical Foundations of Computer Science (MFCS'22)*, volume 241 of *LIPIcs*, pages 75:1–75:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. arXiv:2206.15253, doi:10.4230/LIPIcs.MFCS.2022.75.

[47] Nadia Creignou. A dichotomy theorem for maximum generalized satisfiability problems. *J. Comput. Syst. Sci.*, 51(3):511–522, 1995. doi:10.1006/jcss.1995.1087.

[48] Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems. *Inf. Comput.*, 125(1):1–12, 1996. doi:10.1006/inco.1996.0016.

[49] Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. *Complexity Classification of Boolean Constraint Satisfaction Problems*, volume 7 of *SIAM Monographs on Discrete Mathematics and Applications*. SIAM, 2001.

[50] Nadia Creignou, Phokion G. Kolaitis, and Heribert Vollmer, editors. *Complexity of Constraints - An Overview of Current Research Themes*, volume 5250 of *Lecture Notes in Computer Science*. Springer, 2008. doi:10.1007/978-3-540-92800-3.

[51] Louis Dale. Direct sums of semirings and the Krull-Schmidt theorem. *Kyungpook Mathematical Journal*, 17(2):135–141, 1977. URL: https://koreascience.kr/article/JAKO197725748112087.pdf.

[52] Víctor Dalmau, Marcin Kozik, Andrei A. Krokhin, Konstantin Makarychev, Yury Makarychev, and Jakub Opršal. Robust algorithms with polynomial loss for near-unanimity CSPs. *SIAM J. Comput.*, 48(6):1763–1795, 2019. arXiv:1607.04787, doi:10.1137/18M1163932.

[53] Víctor Dalmau and Andrei A. Krokhin. Robust Satisfiability for CSPs: Hardness and Algorithmic Results. *ACM Trans. Comput. Theory*, 5(4):15:1–15:25, 2013. doi:10.1145/2540090.

[54] Víctor Dalmau, Andrei A. Krokhin, and Rajsekar Manokaran. Towards a characterization of constant-factor approximable finite-valued CSPs. *J. Comput. Syst. Sci.*, 97:14–27, 2018. arXiv:1610.01019, doi:10.1016/j.jcss.2018.03.003.

[55] Víctor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. In *Proc. 39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'24)*, pages 29:1–29:15. ACM, 2024. arXiv:2301.05084, doi:10.1145/3661814.3662068.

[56] Vladimir G. Deineko, Peter Jonsson, Mikael Klasson, and Andrei A. Krokhin. The approximability of MAX CSP with fixed-value constraints. *J. ACM*, 55(4):16:1–16:37, 2008. doi:10.1145/1391289.1391290.

[57] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional Hardness for Approximate Coloring. *SIAM J. Comput.*, 39(3):843–873, 2009. doi:10.1137/07068062X.

[58] Irit Dinur, Oded Regev, and Clifford Smyth. The hardness of 3-uniform hypergraph coloring. *Comb.*, 25(5):519–535, September 2005. doi:10.1007/s00493-005-0032-4.

[59] Martin E. Dyer and David Richerby. An Effective Dichotomy for the Counting Constraint Satisfaction Problem. *SIAM J. Comput.*, 42(3):1245–1274, 2013. doi:10.1137/100811258.

[60] Alina Ene, Jan Vondrák, and Yi Wu. Local Distribution and the Symmetry Gap: Approximability of Multiway Partitioning Problems. In *Proc. 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'13)*, pages 306–325. SIAM, 2013. arXiv:1503.03905, doi:10.1137/1.9781611973105.23.

[61] Tomás Feder and Moshe Y. Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM J. Comput.*, 28(1):57–104, 1998. doi:10.1137/S0097539794266766.

[62] M. R. Garey and David S. Johnson. The complexity of near-optimal graph coloring. *J. ACM*, 23(1):43–49, 1976. doi:10.1145/321921.321926.

[63] Bernd Gärtner and Jiří Matoušek. *Approximation Algorithms and Semidefinite Programming*. Springer Science & Business Media, 2012.

[64] Mrinalkanti Ghosh and Madhur Tulsiani. From Weak to Strong Linear Programming Gaps for All Constraint Satisfaction Problems. *Theory Comput.*, 14(1):1–33, 2018. arXiv:1608.00497, doi:10.4086/toc.2018.v014a010.

[65] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1):1–24, 2007. doi:10.1145/1206035.1206036.

[66] Venkatesan Guruswami and Sanjeev Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM Journal on Discrete Mathematics*, 18(1):30–40, 2004. doi:10.1137/S0895480100376794.

[67] Venkatesan Guruswami and Sai Sandeep. d-To-1 Hardness of Coloring 3-Colorable Graphs with O(1) Colors. In *Proc. 47th International Colloquium on Automata, Languages, and Programming (ICALP'20)*, volume 168 of *LIPIcs*, pages 62:1–62:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.62.

[68] Pavol Hell, J Nešetřil, and Xuding Zhu. Duality and polynomial testing of tree homomorphisms. *Transactions of the American Mathematical Society*, 348(4):1281–1297, 1996. doi:10.1090/S0002-9947-96-01537-1.

[69] Wilfrid Hodges. *Model theory*. Cambridge University Press, 1993.

[70] Sangxia Huang. Improved hardness of approximating chromatic number. In *Proc. 16th International Workshop on Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques and the 17th International Workshop on Randomization and Computation (APPROX-RANDOM'13)*, pages 233–243. Springer, 2013. arXiv:1301.5216, doi:10.1007/978-3-642-40328-6_17.

[71] Anna Huber, Andrei A. Krokhin, and Robert Powell. Skew Bisubmodularity and Valued CSPs. *SIAM J. Comput.*, 43(3):1064–1084, 2014. doi:10.1137/120893549.

[72] Peter G. Jeavons, David A. Cohen, and Marc Gyssens. Closure Properties of Constraints. *J. ACM*, 44(4):527–548, 1997. doi:10.1145/263867.263489.

[73] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the chromatic number. *Comb.*, 20(3):393–415, 2000. doi:10.1007/s004930070013.

[74] Sanjeev Khanna, Madhu Sudan, Luca Trevisan, and David Williamson. The approximability of constraint satisfaction problems. *SIAM J. Comput.*, 30(6):1863–1920, 2000. doi:10.1137/S0097539799349948.

[75] Subhash Khot. Improved Inaproximability Results for MaxClique, Chromatic Number and Approximate Graph Coloring. In *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS'01)*, pages 600–609. IEEE Computer Society, 2001. doi:10.1109/SFCS.2001.959936.

[76] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. doi:10.1137/S0097539705447372.

[77] Phokion G. Kolaitis and Moshe Y. Vardi. On the expressive power of datalog: Tools and a case study. *J. Comput. Syst. Sci.*, 51(1):110–134, 1995. doi:10.1006/jcss.1995.1055.

[78] Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. *J. Comput. Syst. Sci.*, 61(2):302–332, 2000. doi:10.1006/jcss.2000.1713.

[79] Vladimir Kolmogorov, Johan Thapper, and Stanislav Živný. The power of linear programming for general-valued CSPs. *SIAM J. Comput.*, 44(1):1–36, 2015. arXiv:1311.4219, doi:10.1137/130945648.

[80] Pravesh K. Kothari, Raghu Meka, and Prasad Raghavendra. Approximating Rectangles by Juntas and Weakly Exponential Lower Bounds for LP Relaxations of CSPs. *SIAM J. Comput.*, 51(2):17–305, 2022. arXiv:1610.02704, doi:10.1137/17m1152966.

[81] Andrei Krokhin and Stanislav Živný, editors. *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/DFU.Vol7.15301.i.

[82] Andrei A. Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM Journal on Computing*, 52(1):37–79, 2023. arXiv:2003.11351, doi:10.1137/20M1378223.

[83] Gábor Kun, Ryan O'Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *Proc. 3rd Innovations in Theoretical Computer Science (ITCS'12)*, pages 484–495. ACM, 2012. doi:10.1145/2090236.2090274.

[84] Jean B. Lasserre. An explicit equivalent positive semidefinite program for nonlinear 0-1 programs. *SIAM J. Optim.*, 12(3):756–769, 2002. doi:10.1137/S1052623400380079.

[85] Monique Laurent. A Comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre Relaxations for 0-1 Programming. *Math. Oper. Res.*, 28(3):470–496, 2003. doi:10.1287/moor.28.3.470.16391.

[86] James R. Lee, Prasad Raghavendra, and David Steurer. Lower Bounds on the Size of Semidefinite Programming Relaxations. In *Proc. 47th Annual ACM on Symposium on Theory of Computing (STOC'15)*, pages 567–576. ACM, 2015. arXiv:1411.6317, doi:10.1145/2746539.2746599.

[87] Lek-Heng Lim. Tensors and hypermatrices. In Leslie Hogben, editor, *Handbook of Linear Algebra*, pages 231–260. CRC Press Boca Raton, 2nd edition, 2014.

[88] László Lovász and Alexander Schrijver. Cones of Matrices and Set-Functions and 0-1 Optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991. doi:10.1137/0801013.

[89] Alan K. Mackworth. Consistency in networks of relations. *Artif. Intell.*, 8(1):99–118, 1977. doi:10.1016/0004-3702(77)90007-8.

[90] Anatolii Malcev. Untersuchungen aus dem Gebiete der mathematischen Logik. *Matematicheskii Sbornik*, 43(3):323–336, 1936. URL: https://www.mathnet.ru/links/2090b73a4fdf4a07674132012862f576/sm5392.pdf.

[91] Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM*, 60(6), 2013. Article No. 42. arXiv:0911.0801, doi:10.1145/2535926.

[92] Antoine Mottet. Algebraic and algorithmic synergies between promise and infinite-domain CSPs. In *Proc. 40th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'25)*. ACM, 2025. arXiv:2501.13740.

[93] Tamio-Vesa Nakajima and Stanislav Živný. Linearly ordered colourings of hypergraphs. *ACM Trans. Comput. Theory*, 13(3–4), 2022. arXiv:2204.05628, doi:10.1145/3570909.

[94] Michael Pinsker, Jakub Rydval, Moritz Schöbi, and Christoph Spiess. Three meta-questions on infinite-domain constraint satisfaction problems. In *Proc. 50th International Symposium on Mathematical Foundations of Computer Science (MFCS'25)*, 2025. arXiv:2502.06621.

[95] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 245–254, 2008. doi:10.1145/1374376.1374414.

[96] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discret. Math.*, 3(3):411–430, 1990. doi:10.1137/0403036.

[97] Johan Thapper and Stanislav Živný. The complexity of finite-valued CSPs. *J. ACM*, 63(4):37:1–37:33, 2016. arXiv:1210.2987, doi:10.1145/2974019.

[98] Johan Thapper and Stanislav Živný. The power of Sherali-Adams relaxations for general-valued CSPs. *SIAM J. Comput.*, 46(4):1241–1279, 2017. arXiv:1606.02577, doi:10.1137/16M1079245.

[99] Madhur Tulsiani. CSP gaps and reductions in the Lasserre hierarchy. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC'09)*, pages 303–312. ACM, 2009. doi:10.1145/1536414.1536457.

[100] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020. arXiv:1704.01914, doi:10.1145/3402029.

[101] Uri Zwick. Finding almost-satisfying assignments. In *Proc. 13th Annual ACM Symposium on the Theory of Computing (STOC'98)*, pages 551–560. ACM, 1998. doi:10.1145/276698.276869.