Pebble Games and Linear Equations

Martin Grohe[1] and Martin Otto[*2]

1 **Humboldt-Universität zu Berlin, Germany**
grohe@informatik.hu-berlin.de
2 **Technische Universität Darmstadt, Germany**
otto@mathematik.tu-darmstadt.de

Abstract

We give a new, simplified and detailed account of the correspondence between levels of the Sherali–Adams relaxation of graph isomorphism and levels of pebble-game equivalence with counting (higher-dimensional Weisfeiler–Lehman colour refinement). The correspondence between basic colour refinement and fractional isomorphism, due to Ramana, Scheinerman and Ullman [18], is re-interpreted as the base level of Sherali–Adams and generalised to higher levels in this sense by Atserias and Maneva [1], who prove that the two resulting hierarchies interleave. In carrying this analysis further, we here give (a) a precise characterisation of the level k Sherali–Adams relaxation in terms of a modified counting pebble game; (b) a variant of the Sherali–Adams levels that precisely match the k-pebble counting game; (c) a proof that the interleaving between these two hierarchies is strict. We also investigate the variation based on boolean arithmetic instead of real/rational arithmetic and obtain analogous correspondences and separations for plain k-pebble equivalence (without counting). Our results are driven by considerably simplified accounts of the underlying combinatorics and linear algebra.

1998 ACM Subject Classification F.4.1, G.2.2

Keywords and phrases Finite model theory, finite variable logics, graph isomorphism, Weisfeiler-Lehman algorithm, linear programming, Sherali–Adams hierarchy

Digital Object Identifier 10.4230/LIPIcs.CSL.2012.289

1 Introduction

We study a surprising connection between equivalence in finite variable logics and a linear programming approach to the graph isomorphism problem. This connection has recently been uncovered by Atserias and Maneva [1], building on earlier work of Ramana, Scheinerman and Ullman [18] that just concerns the 2-variable case.

Finite variable logics play a central role in finite model theory. Most important for this paper are finite variable logics with counting, which have been specifically studied in connection with the question for a logical characterisation of polynomial time and in connection with the graph isomorphism problem (e.g. [4, 8, 9, 13, 14, 17]). Equivalence in finite variable logics can be characterised in terms of simple combinatorial games known as pebble games. Specifically, C^k-equivalence can be characterised by the bijective k-pebble game introduced by Hella [11]. Cai, Fürer and Immerman [4] observed that C^k-equivalence exactly corresponds to indistinguishability by the k-dimensional Weisfeiler-Lehman (WL) algorithm,[1]

* This Martin gratefully acknowledges the other Martin's academic hospitality at HU Berlin during his sabbatical in the winter 2011/12.

[1] The dimensions of the WL algorithm are counted differently in the literature; what we call "k-dimensional" here is sometimes called "$(k-1)$-dimensional".

a combinatorial graph isomorphism algorithm introduced by Babai, who attributed it to work of Weisfeiler and Lehman in the 1970s. The 2-dimensional version of the WL algorithm precisely corresponds to an even simpler isomorphism algorithm known as colour refinement.

The isomorphisms between two graphs can be described by the integral solutions of a system of linear equations. If we have two graphs with adjacency matrices A and B, then each isomorphism from the first to the second corresponds to a permutation matrix X such that $X^tAX = B$, or equivalently

$$AX = XB. \tag{1}$$

If we view the entries of X as variables, this equation corresponds to a system of linear equations. We can add inequalities that force X to be a permutation matrix and obtain a system ISO of linear equations and inequalities whose integral solutions correspond to the isomorphisms between the two graphs. In particular, the system ISO has an integral solution if, and only if, the two graphs are isomorphic.

What happens if we drop the integrality constraints, that is, we admit arbitrary real solutions of the system ISO? We can ask for doubly stochastic matrices X satisfying equation (1). (A real matrix is *doubly stochastic* if its entries are non-negative and all row sums and column sums are one.) Ramana, Scheinerman and Ullman [18] proved a beautiful result that establishes a connection between linear algebra and logic: the system ISO has a real solution if, and only if, the colour refinement algorithm does not distinguish the two graphs with adjacency matrices A and B. Recall that the latter is equivalent to the two graphs being C^2-equivalent.

To bridge the gap between integer linear programs and their LP-relaxations, researchers in combinatorial optimisation often add additional constraints to the linear programs to bring them closer to their integer counterparts. The Sherali–Adams hierarchy [21] of relaxations gives a systematic way of doing this. For every integer linear program IL in n variables and every positive integer k, there is a *rank-k Sherali–Adams relaxation* IL(k) of IL, such that IL(1) is the standard LP-relaxation of IL where all integrality constraints are dropped and IL(n) is equivalent to IL. There is a considerable body of research studying the strength of the various levels of this and related hierarchies (e.g. [2, 3, 5, 16, 20, 19]).

Quite surprisingly, Atserias and Maneva [1] were able to lift the Ramana–Scheinerman–Ullman result, which we may now restate as an equivalence between ISO(1) and C^2-equivalence, to a close correspondence between the higher levels of the Sherali–Adams hierarchy for ISO and the logics C^k. They proved for every $k \geq 2$:

1. if ISO(k) has a (real) solution, then the two graphs are C^k-equivalent;
2. if the two graphs are C^k-equivalent, then ISO($k-1$) has a solution.

Atserias and Maneva used this results to transfer results about the logics C^k to the world of polyhedral combinatorics and combinatorial optimisation, and conversely, results about the Sherali–Adams hierarchy to logic.

Atserias and Maneva [1] left open the question whether the interleaving between the levels of the Sherali–Adams hierarchy and the finite-variable-logic hierarchy is strict or whether either the correspondence between C^k-equivalence and ISO(k) or the correspondence between C^k-equivalence and ISO($k-1$) is exact. Note that for $k = 2$ the correspondence between C^k-equivalence and ISO($k-1$) is exact by the Ramana–Scheinerman–Ullman theorem. We prove that for all $k \geq 3$ the interleaving is strict. However, we can prove an exact correspondence between ISO($k-1$) and a variant of the bijective k-pebble game that characterises C^k-equivalence. This variant, which we call the weak bijective k-pebble game, is actually equivalent to a game called ($k-1$)-sliding game by Atserias and Maneva.

Maybe most importantly, we prove that a natural combination of equalities from ISO(k) and ISO($k-1$) gives a linear program ISO($k-1/2$) that characterises C^k-equivalence exactly.

To obtain these results, we give simple new proofs of the theorems of Ramana, Scheinerman and Ullman and of Atserias and Maneva. Whereas the previous proofs use two non-trivial results from linear algebra, the Perron–Frobenius Theorem (about the eigenvalues of positive matrices) and the Birkhoff–von Neumann Theorem (stating that every doubly stochastic matrix is a convex combination of permutation matrices), our proofs only use elementary linear algebra. This makes them more transparent and less mysterious (at least to us).

In fact, the linear algebra we use is so simple that much of it can be carried out not only over the field of real numbers, but over arbitrary semirings. By using similar algebraic arguments over the boolean semiring (with disjunction as addition and conjunction as multiplication), we obtain analogous results to those for C^k-equivalence for the ordinary k-variable logic L^k, characterising L^k-equivalence, i.e., k-pebble game equivalence without counting, by systems of 'linear' equations over the boolean semiring.

For the ease of presentation, we have decided to present our results only for undirected simple graphs. It is easy to extend all results to relational structures with at most binary relations. Atserias and Maneva did this for their results, and for ours the extension works analogously. An extension to structures with relations of higher arities also seems possible, but is more complicated and comes at the price of loosing some of the elegance of the results.

Due to space limitations, we have to omit many details and proofs in this conference version of the paper. They can be found in the full version of the paper [10]. The present version of the paper contains a fairly complete account of our proof of the Ramana–Scheinerman–Ullman theorem, including the linear algebra that is also underlying they higher-dimensional results. Most proofs regarding the correspondence between the Sherali–Adams hierarchy and C^k-equivalence are omitted.

2 Finite variable logics and pebble games

We assume the reader is familiar with the basics of first-order logic FO. We almost exclusively consider first-order logic over finite graphs, which we view as finite relational structures with one binary relation. We assume graphs to be undirected and loop-free. For every positive integer k, we let L^k be the fragment of FO consisting of all formulae that contain at most k distinct variables. We let C^k be the extension of L^k by *counting quantifiers* $\exists^{\geq n}$, where$\exists^{\geq n} x\,\varphi$ means that there are at least n elements x such that φ is satisfied. L^k-*equivalence* of structures $\mathcal{A}, \mathcal{B}$ is denoted by $\mathcal{A} \equiv^k_{\mathsf{L}} \mathcal{B}$ and C^k-*equivalence* by $\mathcal{A} \equiv^k_{\mathsf{C}} \mathcal{B}$. Both equivalences can be characterised in terms of pebble games. We briefly sketch the *bijective k-pebble game* [11] that characterises C^k-equivalence. The game is played by two players on a pair $\mathcal{A}, \mathcal{B}$ of structures. A *play* of the game consists of a (possibly infinite) sequence of *rounds*. In each round, player **I** picks up one of his pebbles, and player **II** picks up her corresponding pebble. Then player **II** chooses a bijection f between $\mathcal{A}$ and $\mathcal{B}$ (if no such bijection exists, that is, if the structures have different cardinalities, player **II** immediately looses). Then player **I** places his pebble on an element a of $\mathcal{A}$, and player **II** places her pebble on $f(a)$. Note that after each round r there is a subset $p \subseteq \mathcal{A} \times \mathcal{B}$ consisting of the at most k pairs of elements on which the pairs of corresponding pebbles are placed. We call p the *position* after round r. Player **I** wins the play if every position that occurs is a local isomorphism, that is, a local mapping from $\mathcal{A}$ to $\mathcal{B}$ that is injective and preserves membership and non-membership in all relations (adjacency and non-adjacecny if $\mathcal{A}$ and $\mathcal{B}$ are graphs). Then $\mathcal{A} \equiv^k_{\mathsf{C}} \mathcal{B}$ if, and only if, player **II** has a winning strategy for the game.

C^k-equivalence also corresponds to a simple combinatorial algorithm for graph isomorphism testing known as the Weisfeiler-Lehman algorithm.

We refer the reader to the textbooks [6, 7, 12, 15] and the monograph [17] for a more thorough exposition of the material sketched here.

3 Basic combinatorics and linear algebra

We consider matrices with entries in $\mathbb{B} = \{0,1\}$, $\mathbb{Q}$ or $\mathbb{R}$. A matrix $X \in \mathbb{R}^{m,n}$ with m rows and n columns has entry X_{ij} in row $i \in [m] = \{1,\dots,m\}$ and column $j \in [n] = \{1,\dots,n\}$. We write E_n for the n-dimensional unit matrix.

We write $X \geqslant 0$ to say that (the real or rational) matrix X has only non-negative entries, and $X > 0$ to say that all entries are strictly positive. We also speak of *non-negative* or *strictly positive matrices* in this sense. For a boolean matrix, strict positivity, $X > 0$ means that all entries are 1. A square $n \times n$-matrix is *doubly stochastic* if its entries are non-negative and if the sum of entries across every row and column is 1. *Permutation matrices* are doubly stochastic matrices over $\{0,1\}$, with precisely one 1 in every row and in every column.

It will be useful to have the shorthand notation $X_{D_1 D_2} = 0$ for the assertion that $X_{d_1 d_2} = 0$ for all $d_1 \in D_1$, $d_2 \in D_2$.

3.1 Decomposition into irreducible blocks

With $X \in \mathbb{R}^{n,n}$ associate the directed graph $\mathrm{G}(X) := ([n], \{(i,j) \colon X_{ij} \neq 0\})$. The strongly connected components of $\mathrm{G}(X)$ induce a partition of the set $[n] = \{1,\dots,n\}$ of rows/columns of X. X is called *irreducible* if this partition has just the set $[n]$ itself.

Note that X is irreducible iff $P^t X P$ is irreducible for every permutation matrix P.

▶ **Observation 3.1.** *Let $X \in \mathbb{R}^{n,n} \geqslant 0$ with strictly positive diagonal entries. If X is irreducible, then all powers X^ℓ for $\ell \geqslant n-1$ have non-zero entries throughout. Moreover, if X is irreducible, then so is X^ℓ for all $\ell \geqslant 1$.*

Let us call two matrices $Z, Z' \in \mathbb{R}^{n,n}$ *permutation-similar* or *S_n-similar*, $Z \sim_{S_n} Z'$, if $Z' = P^t Z P$ for some permutation matrix P, i.e., if one is obtained from the other by a coherent permutation of rows and columns.

▶ **Lemma 3.2.** *Every symmetric $Z \in \mathbb{R}^{n,n} \geqslant 0$ is permutation-similar to some block diagonal matrix $\mathrm{diag}(Z_1,\dots,Z_s)$ with irreducible blocks $Z_i \in \mathbb{R}^{n_i,n_i}$.*

The permutation matrix P corresponding to the row- and column-permutation $p \in S_n$ that puts Z into block diagonal form $P^t Z P = \mathrm{diag}(Z_1,\dots,Z_s)$ with irreducible blocks, is unique up to an outer permutation that re-arranges the block intervals $([k_i+1, k_i+n_i])_{1 \leqslant i \leqslant s}$ where $k_i = \sum_{j<i} n_j$, and a product of inner permutations within each one of these s blocks.

The underlying partition $[n] = \dot{\bigcup}_{1 \leqslant i \leqslant s} D_i$ where $D_i := p([k_i+1, k_i+n_i])$ for $k_i = \sum_{j<i} n_j$, is uniquely determined by Z.[2]

In the following we refer to the *partition induced by a symmetric matrix Z*.

[2] Here we regard two partitions as identical if they have the same partition sets, i.e., we ignore their indexing/enumeration.

▶ **Observation 3.3.** *In the situation of Lemma 3.2, the partition* $[n] = \dot{\bigcup}_i D_i$ *induced by the symmetric matrix* Z *is the partition of* $[n]$ *into the vertex sets of the connected components of* $G(Z)$*. Then, for every pair* $i \neq j$, $Z_{D_i D_j} = 0$*, while all the minors* $Z_{D_i D_i}$ *are irreducible.*[3]

If, moreover, Z *has strictly positive diagonal entries, then the partition induced by* Z *is the same as that induced by* Z^ℓ*, for any* $\ell \geqslant 1$*; for* $\ell \geqslant n-1$*, the diagonal blocks* $(Z^\ell)_{D_i D_i}$ *have non-zero entries throughout:* $(Z^\ell)_{D_i D_i} > 0$ *.*

The last assertion says that for a symmetric $n \times n$ matrix Z with non-negative entries and no zeroes on the diagonal, all powers Z^ℓ for $\ell \geqslant n-1$ are *good symmetric* in the sense of the following definition.

▶ **Definition 3.4.** Let $Z \geqslant 0$ be symmetric with strictly positive diagonal. Then Z is called *good symmetric* if w.r.t. the partition $[n] = \dot{\bigcup}_i D_i$ induced by Z, all $Z_{D_i D_i} > 0$.

More generally, a not necessarily symmetric matrix $X \geqslant 0$ without null rows or columns is *good* if $Z = XX^t$ and $Z' = X^t X$ are good in the above sense.

The importance of this notion lies in the fact that, as observed above, for an arbitrary symmetric $n \times n$ matrix $Z \geqslant 0$ without zeroes on the diagonal, the partition induced by Z is the same as that induced by the good symmetric matrix $\hat{Z} := Z^{n-1}$; and, as for any good matrix, this partition can simply be read off from $\hat{Z}$: $i, j \in [n]$ are in the same partition set if, and only if, $\hat{Z}_{ij} \neq 0$.

▶ **Definition 3.5.** Consider partitions $[n] = \dot{\bigcup}_{i \in I} D_i$ and $[m] = \dot{\bigcup}_{i \in I} D'_i$ of the sets $[n]$ and $[m]$ with the same number of partition sets. We say that these two partitions are X*-related* for some matrix $X \in \mathbb{R}^{n,m}$ if

(i) $X \geqslant 0$ has no null rows or columns, and
(ii) $X_{D_i D'_j} = 0$ for every pair of distinct indices $i, j \in I$.

Note that partitions that are X-related are X^t-related in the opposite direction. More importantly, each one of the X/X^t-related partitions can be recovered from the other one through X according to

$$
\begin{aligned}
D'_i &= \{d' \in [m] \colon X_{dd'} > 0 \text{ for some } d \in D_i\}, \\
D_i &= \{d \in [n] \colon X_{dd'} > 0 \text{ for some } d' \in D'_i\}.
\end{aligned}
$$

For a more algebraic treatment, we associate with the partition sets D_i of a partition $[n] = \dot{\bigcup}_{i \in I} D_i$ the *characteristic vectors* $\mathbf{d}_i$ with entries 1 and 0 according to whether the corresponding component belongs to D_i:

$$
\mathbf{d}_i = \textstyle\sum_{d \in D_i} \mathbf{e}_d,
$$

where $\mathbf{e}_d$ is the d-th standard basis vector. In terms of these characteristic vectors $\mathbf{d}_i$ for $[n] = \dot{\bigcup}_{i \in I} D_i$ and $\mathbf{d}'_i$ for $[m] = \dot{\bigcup}_{i \in I} D'_i$, the X/X^t-relatedness of these partitions means that

$$
\begin{aligned}
D'_i &= \{d' \in [m] \colon (X^t \mathbf{d}_i)_{d'} > 0\}, \\
D_i &= \{d \in [n] \colon (X \mathbf{d}'_i)_d > 0\}.
\end{aligned}
$$

[3] Note that this does not depend on the enumeration of the partition set D_i, because irreducibility is invariant under permutation-similarity.

▶ **Lemma 3.6.** *If two partitions* $[n] = \dot{\bigcup}_{i\in I} D_i$ *and* $[n] = \dot{\bigcup}_{i\in I} D'_i$ *of the same set* $[n]$ *are* X*-related for some doubly stochastic matrix* $X \in \mathbb{R}^{n,n}$*, then* $|D_i| = |D'_i|$ *for all* $i \in I$*, and for the characteristic vectors* $\mathbf{d}_i$ *and* $\mathbf{d}'_i$ *of the partition sets* D_i *and* D'_i

$$\mathbf{d}_i = X\mathbf{d}'_i \quad \text{and} \quad \mathbf{d}'_i = X^t\mathbf{d}_i.$$

Proof. Observe that for all $d \in [n]$ we have $0 \le (X\mathbf{d}'_i)_d = \sum_{d'\in D'_i} X_{dd'} \le 1$. It follows immediately from the definition of X-relatedness that $(X\mathbf{d}'_i)_d = 0$ for all $d \notin D_i$. Therefore,

$$|D_i| \geqslant \sum_{d\in D_i} (X\mathbf{d}'_i)_d = \sum_{d\in[n]} (X\mathbf{d}'_i)_d = \sum_{d'\in D'_i} \sum_{d\in[n]} X_{dd'} = |D'_i|.$$

Similarly, $0 \le (X^t\mathbf{d}_i)_{d'} \le 1$ for $d' \in [n]$, and $|D'_i| \ge \sum_{d'\in D'_i} (X^t\mathbf{d}_i)_{d'} = |D_i|$. Together, we obtain

$$|D_i| = \sum_{d\in D_i} (X\mathbf{d}'_i)_d = |D'_i| = \sum_{d'\in D'_i} (X^t\mathbf{d}_i)_{d'}.$$

As all summands are bounded by 1, this implies $(X\mathbf{d}'_i)_d = 1$ for all $d \in D_i$ and $(X^t\mathbf{d}_i)_{d'} = 1$ for all $d' \in D_i$. ◀

▶ **Lemma 3.7.** *Let* $X \geqslant 0$ *be an* $m\times n$ *matrix without null rows or columns. Then the* $m\times m$ *matrix* $Z := XX^t$ *and the* $n\times n$ *matrix* $Z' := X^tX$ *are symmetric with positive entries on their diagonals. Moreover, the (unique) partitions of* $[m]$ *and* $[n]$ *that are induced by* Z *and* Z'*, respectively, are* X/X^t*-related.*[4]

Proof. It is obvious that Z and Z' are symmetric with positive diagonal entries. Let partitions $[m] = \dot{\bigcup}_{i\in I} D_i$ and $[n] = \dot{\bigcup}_{i\in I'} D'_i$ be obtained from decompositions of Z and Z' into irreducible blocks. We need to show that the non-zero entries in X give rise to a coherent bijection between the index sets I and I' of the two partitions, in the sense that partition sets D_i and D'_j are related if, and only if, some pair of members $d \in D_i$ and $d' \in D'_j$ have a positive entry $X_{dd'}$. Then a re-numbering of one of these partitions will make them X-related in the sense of Definition 3.5. Recall from Observation 3.3 that the D_i are the vertex sets of the connected components of $G(XX^t)$ on $[m]$, while the D'_i the are the vertex sets of the connected components of $G(X^tX)$ on $[n]$.

Consider the uniformly directed bipartite graph $G(X)$ on $[m] \dot\cup [n]$ with an edge from $i \in [m]$ to $j \in [n]$ if $X_{ij} > 0$. In light of the symmetry of the whole situation w.r.t. X and X^t, it just remains to argue for instance that no $i \in [m]$ can have edges into two distinct sets of the partition $[n] = \dot{\bigcup}_{i\in I'} D'_i$. But any two target nodes of edges from one and the same $i \in [n]$ are in the same connected component of $G(X^tX)$, hence in the same partition set. ◀

In the situation of Lemma 3.7, powers of Z induce the same partitions as Z, and the partitions induced by $(Z^\ell X)(Z^\ell X)^t = Z^{2\ell+1}$ are X/X^t-related as well as $Z^\ell X/X^tZ^\ell$-related, for all $\ell \geqslant 1$.

For $\ell \geqslant n/2 - 1$, the matrix $Z^\ell X$ has no null rows or columns: else $Z^\ell X(Z^\ell X)^t = Z^{2\ell+1}$ would have to have a zero entry on the diagonal, contradicting the fact that this symmetric matrix is good symmetric in the sense of Definition 3.4. The same reasoning shows that $Z^\ell X$ is itself good in the sense of Definition 3.4.

[4] As X/X^t-relatedness refers to partitions presented with an indexing of the partition sets, we need to allow a suitable re-indexing for at least one of them, so as to match the other one.

▶ **Corollary 3.8.** *Let $X \geqslant 0$ be an $m \times n$ matrix without null rows or columns, $Z = XX^t$, $Z' = X^tX$ the associated symmetric matrices with non-zero entries on the diagonal. Then for $\ell \geqslant m-1$, the matrix $\hat{X} := Z^\ell X = X(Z')^\ell$ and its transpose $\hat{X}^t = X^t Z^\ell = (Z')^\ell X^t$ are good and relate the partitions $[m] = \dot{\bigcup}_i D_i$ and $[n] = \dot{\bigcup}_i D'_i$ induced by Z and Z', respectively.*[4] *Moreover,*

(i) $\hat{X}_{D_i D'_i} > 0$ *for all i, and*

(ii) $\hat{X}_{D_i D'_j} = 0$ *for all $i \neq j$.*

Aside: boolean vs. real arithmetic

Looking at matrices with $\{0,1\}$-entries, we may not only treat them as matrices over $\mathbb{R}$ as we have done so far, but also over other fields, or as matrices over the boolean semiring $\mathbb{B} = \{0,1\}$ with the logical operations of $\vee$ for addition and $\wedge$ for multiplication. Though not even forming a ring, boolean arithmetic yields a very natural interpretation in the context where we associate non-negative entries with edges, as we did in passage from X to $G(X)$. The 'normalisation map' $\chi \colon \mathbb{R}_{\geqslant 0} \to \{0,1\}$, $x \mapsto 1$ iff $x > 0$, relates the arithmetic of reals $x, y \geqslant 0$ to boolean arithmetic in

$$\chi(x+y) = \chi(x) \vee \chi(y) \quad \text{and} \quad \chi(xy) = \chi(x) \wedge \chi(y).$$

This is the 'logical' arithmetic that supports, for instance, arguments used in Observation 3.1: for any real $n \times n$ matrix $X \geqslant 0$, $(XX)_{ij} = \sum_k X_{ik}X_{kj} \neq 0$ iff there is at least one $k \in [n]$ for which $X_{ik} \neq 0$ and $X_{kj} \neq 0$ iff $\bigvee_{k\in[n]}(\chi(X_{ik}) \wedge \chi(X_{kj})) = 1$. It is no surprise, therefore, that several of the considerations apparently presented for real non-negative matrices above, have immediate analogues for boolean arithmetic – in fact, one could argue, that the boolean interpretation is closer to the combinatorial essence. We briefly sum up these analogues with a view to their use in the analysis of L^k-equivalence, while the real versions are related to C^k-equivalence. The boolean analogue of a doubly stochastic matrix with non-negative real entries is a matrix without null rows or columns.

Also note that the definitions of irreducibility and X-relatedness are applicable to boolean matrices without any changes. Observations 3.1 and 3.3 go through (as just indicated), and so does Lemma 3.2. For Lemma 3.6, one may look at X-related partitions of sets $[m]$ and $[n]$, where not necessarily $n = m$, by any boolean matrix X without null rows or columns, and obtains the relationship between the characteristic vectors as stated there, now in terms of boolean arithmetic – but of course we do not get any numerical equalities between the sizes of the partition sets. Lemma 3.7, finally, applies to boolean arithmetic, exactly as stated.

▶ **Lemma 3.9.** *In the sense of boolean arithmetic for matrices with entries in $\mathbb{B} = \{0,1\}$:*

(a) *Any symmetric $Z \in \mathbb{B}^{n,n}$ induces a unique partition of $[n]$ for which the diagonal minors induced by the partition sets are irreducible and the remaining blocks null; $d, d' \in [n]$ are in the same partition set if, and only if, in the sense of boolean arithmetic $(Z^\ell)_{dd'} = 1$ for any/all $\ell \geqslant n-1$.*

(b) *If two partitions (not necessarily of the same set) with the same number of partition sets are related by some boolean matrix $X \in \mathbb{B}^{m,n}$, then the characteristic vectors $(\mathbf{d}_i)_{i\in I}$ and $(\mathbf{d}'_i)_{i\in I}$ of the partitions are related by $\mathbf{d}_i = X\mathbf{d}'_i$ and $\mathbf{d}'_i = X^t\mathbf{d}_i$ in the sense of boolean arithmetic.*

(c) *For any matrix $X \in \mathbb{B}^{m,n}$ without null rows or columns, the symmetric boolean matrices $Z = XX^t$ and $Z' = X^tX$ have diagonal entries 1 and induce partitions that are X/X^t-related, and agree with the partitions induced by higher powers of Z and Z' or on the basis of $Z^\ell X$ and $X(Z')^\ell$ for any $\ell \in \mathbb{N}$. For $\ell \geqslant m-1, n-1$, the partition blocks in Z*

and Z' have entries 1 throughout, and $Z^\ell X$ and $X(Z')^\ell$ have entries 1 in all positions relating elements from matching partition sets.

▶ **Observation 3.10.** *For a symmetric boolean matrix $Z \in \mathbb{B}^{n,n}$ with $Z_{dd} = 1$ for all $d \in [n]$, the characteristic vectors $\mathbf{d}_i$ of the partition $[n] = \dot{\bigcup}_{i\in I} D_i$ induced by Z satisfy the following 'eigenvector' equation in terms of boolean arithmetic:*

$$Z\mathbf{d}_i = \mathbf{d}_i \quad \text{(boolean)}, \quad \text{for all } i \in I.$$

3.2 Eigenvalues and -vectors

▶ **Lemma 3.11.** *If $Z \in \mathbb{R}^{n,n}$ is doubly stochastic, then it has eigenvalue 1. If Z is doubly stochastic and irreducible with strictly positive diagonal entries, then the eigenspace for eigenvalue 1 has dimension 1 and is spanned by the vector $\mathbf{d} := (1, \ldots, 1)^t$.*

Proof. It is obvious that $\mathbf{d}$ is an eigenvector of Z with eigenvalue 1. The eigenspace with eigenvalue 1 is contained in that of Z^{n-1}, which has entries strictly between 0 and 1 throughout if Z is irreducible with strictly positive diagonal, by Observation 3.1. For 1-dimensionality observe that all entries of $Z^{n-1}\mathbf{v}$ are convex combinations of the entries of $\mathbf{v}$ with coefficients strictly between 0 and 1. ◀

▶ **Corollary 3.12.** **(a)** *Let $Z \in \mathbb{R}^{n,n}$ be doubly stochastic with positive diagonal, and $[n] = \dot{\bigcup}_i D_i$ a partition with $Z_{D_iD_j} = 0$ for $i \neq j$ and such that the minors $Z_{D_iD_i}$ are irreducible for all i. Then the eigenspace for eigenvalue 1 of Z is the direct sum of the 1-dimensional subspaces spanned by the characteristic vectors $\mathbf{d}_i$ of the partition sets D_i.*

(b) *If $Z = X^tX \in \mathbb{R}^{n,n}$ for some doubly stochastic matrix X, then the eigenspace for eigenvalue 1 is the direct sum of the spans of the characteristic vectors $\mathbf{d}_i$ from the unique partition $[n] = \dot{\bigcup}_i D_i$ of $[n]$ induced by Z according to Lemma 3.2.*

3.3 Stable partitions

▶ **Definition 3.13.** Let $A \in \mathbb{R}^{n,n}$ and $[n] = \dot{\bigcup}_{i\in I} D_i$ a partition. We call this partition a *stable partition for A* if there are numbers $(s_{ij})_{i,j\in I}$ and $(t_{ij})_{i,j\in I}$ such that for all $i, j \in I$:

$$d \in D_i \quad \Rightarrow \quad \sum_{d'\in D_j} A_{dd'} = s_{ij} \quad \text{and} \quad \sum_{d'\in D_j} A_{d'd} = t_{ij}.$$

If there are s_{ij} such that $\sum_{d'\in D_j} A_{dd'} = s_{ij}$ for all $d \in D_i$, we call the partition *row-stable*; similarly, for t_{ij} such that $\sum_{d'\in D_j} A_{d'd} = t_{ij}$ for all $d \in D_i$, *column-stable.*

For symmetric A, column- and row-stability are equivalent (with $t_{ij} = s_{ij}$).

Note that the row and column sums in the definition are the D_i-components of $A\mathbf{d}_j$ and of $\mathbf{d}_j^tA = (A^t\mathbf{d}_j)^t$, respectively. So, for instance, row stability precisely says that for all i the vector $A\mathbf{d}_i$ is in the span of the vectors $\mathbf{d}_j$.

▶ **Lemma 3.14.** *Let $A \in \mathbb{R}^{n,n}$ commute with some symmetric matrix of the form $Z = XX^t \in \mathbb{R}^{n,n}$ for some doubly stochastic $X \in \mathbb{R}^{n,n}$. Then the partition $[n] = \dot{\bigcup}_i D_i$ of $[n]$ induced by Z according to Lemma 3.2 is stable for A.*

Proof. Using the characteristic vectors $\mathbf{d}_i$ of the partition sets again, we have $ZA\mathbf{d}_i = AZ\mathbf{d}_i = A\mathbf{d}_i$, and thus $A\mathbf{d}_i$ is an eigenvector of Z with eigenvalue 1. Hence by Corollary 3.12, it is in the span of the vectors $\mathbf{d}_j$, and this means that the partition is row stable. Column stability is established similarly. ◀

▶ **Corollary 3.15.** *Let A commute with $Z = XX^t$ and B commute with $Z' = X^tX$, where X is doubly stochastic (cf. Lemma 3.14). Then the partitions induced by Z and Z', which are X-related by Lemma 3.7, are stable for A and B, respectively.*

Aside: boolean arithmetic

We give a separate elementary proof of the analogue of Lemma 3.14 for boolean arithmetic. Here the definition of a *boolean* stable partition is this natural analogue of Definition 3.13.

▶ **Definition 3.16.** A partition $[n] = \dot{\bigcup}_{i\in I} D_i$ is *boolean stable* for $A \in \mathbb{B}^{n,n}$ if, in the sense of boolean arithmetic, $\sum_{d'\in D_j} A_{dd'}$ and $\sum_{d'\in D_j} A_{d'd}$ only depend on the set D_i for which $d \in D_i$.

Note that boolean stability implies that, for the characteristic vectors $\mathbf{d}_i$ of the partition, $(A\mathbf{d}_j)_d = \sum_{d'\in D_j} A_{dd'}$ is the same for all $d \in D_i$, so that also here $A\mathbf{d}_j$ is a boolean linear combination of the characteristic vectors $\mathbf{d}_i$.

▶ **Lemma 3.17.** *Let $A \in \mathbb{B}^{n,n}$ commute, in the sense of boolean arithmetic, with some symmetric matrix of the form $Z = XX^t \in \mathbb{B}^{n,n}$ with entries $Z_{dd} = 1$ for all $d \in [n]$. Then the partition $[n] = \dot{\bigcup}_i D_i$ induced by Z according to Lemma 3.9 is boolean stable for A.*

4 Fractional isomorphism

4.1 C^2-equivalence and linear equations

The *adjacency matrix* of graph $\mathcal{A}$ is the square matrix A with rows and columns indexed by vertices of $\mathcal{A}$ and entries $A_{aa'} = 1$ if aa' is an edge of $\mathcal{A}$ and $A_{aa'} = 0$ otherwise. By our assumption that graphs are undirected and simple, A is a symmetric square matrix with null diagonal. It will be convenient to assume that our graphs always have an initial segment $[n]$ of the positive integers as their vertex set. Then the adjacency matrices are in $\mathbb{B}^{n,n} \subseteq \mathbb{R}^{n,n}$. Throughout this subsection, we assume that $\mathcal{A}$ and $\mathcal{B}$ are graphs with vertex set $[n]$ and with adjacency matrices A, B, respectively. It will be notationally suggestive to denote typical indices of matrices $a, a', \ldots \in [n]$ when they are to be interpreted as vertices of $\mathcal{A}$, and $b, b', \ldots \in [n]$ when they are to be interpreted as vertices of $\mathcal{B}$.

Recall (from the discussion in the introduction) that two graphs $\mathcal{A}, \mathcal{B}$ are isomorphic if, and only if, there is a permutation matrix X such that $AX = XB$. We can rewrite this as the following integer linear program in the variables X_{ab} for $a, b \in [n]$.

ISO
$$\sum_{b'\in[n]} X_{ab'} = \sum_{a'\in[n]} X_{a'b} = 1,$$
$$\sum_{a'\in[n]} A_{aa'}X_{a'b} = \sum_{b'\in[n]} X_{ab'}B_{b'b},$$
$$X_{ab} \geq 0 \qquad \text{for all } a, b \in [n].$$

Then $\mathcal{A}$ and $\mathcal{B}$ are isomorphic if, and only if, ISO has an integer solution.

▶ **Definition 4.1.** Two graphs $\mathcal{A}, \mathcal{B}$ are *fractionally isomorphic*, $\mathcal{A} \approx \mathcal{B}$, if, and only if, the system ISO has a real solution.

So graphs are fractionally isomorphic if, and only if, there is a doubly stochastic matrix X such that $AX = XA$. Note that fractionally isomorphic graphs necessarily have the same number of vertices (this will be different for the boolean analogue, which cannot count).

A *stable partition* of the vertex set of an undirected graph is a stable partition $[n] = \dot{\bigcup}_{i \in I} D_i$ for its adjacency matrix in the sense of Definition 3.13. The characteristic parameters for a stable partition $[n] = \dot{\bigcup}_{i \in I} D_i$ for A are the numbers $s_{ij} = s_{ij}^A$ such that $s_{ij} = \sum_{d' \in D_j} A_{dd'}$ for all $d \in D_i$. (As A is symmetric, the parameters t_{ij} of Definition 3.13 are equal to the s_{ij}.) We call two stable partitions $\dot{\bigcup}_{i \in I} D_i$ for a matrix A and $\dot{\bigcup}_{i \in J} D'_i$ for a matrix B *equivalent* if $I = J$ and $|D_i| = |D'_i|$ for all $i \in I$ and $s_{ij}^A = s_{ij}^B$ and for all $i, j \in I$.

▶ **Lemma 4.2.** *$\mathcal{A}$ and $\mathcal{B}$ are C^2-equivalent if, and only if, there are equivalent stable partitions $\dot{\bigcup}_{i \in I} D_i$ for A and $\dot{\bigcup}_{i \in I} D'_i$ for B.*

Proof sketch. The partition of the elements $\mathcal{A}$ and $\mathcal{B}$ according to their C^2-type yields equivalent stable partitions (two elements have the same C^2-type if they satisfy the same C^2-formulae with one free variable). For the converse, it can be shown that equivalent stable partitions give player **II** a winning strategy in the bijective 2-pebble game. ◀

▶ **Theorem 4.3** (Ramana–Scheinerman–Ullman)**.** *Two graphs are C^2-equivalent if, and only if, they are fractionally isomorphic.*

Proof. In view of Lemma 4.2, it suffices to prove that $\mathcal{A}$ and $\mathcal{B}$ have equivalent stable partitions if, and only if, they are fractionally isomorphic.

For the forward direction, suppose that we have equivalent stable partitions $\dot{\bigcup}_{i \in I} D_i$ for A and $\dot{\bigcup}_{i \in J} D'_i$ for B. For all $a \in D_i, b \in D'_j$ we let $X_{ab} := \delta(i,j)/n_i$, where $n_i := |D_i| = |D'_i|$. (Here and elsewhere we use Kronecker's δ function defined by $\delta(i,j) = 1$ if $i = j$ and $\delta(i,j) = 0$ otherwise.) An easy calculation shows that this defines a doubly stochastic matrix X with $AX = XB$, that is, a solution for ISO.

For the converse implication, suppose that X is a doubly stochastic matrix such that $AX = XB$. Since A and B are symmetric, also $X^t A = B X^t$, which implies that A commutes with $Z := XX^t$ and B with $Z' := X^t X$.

From Lemma 3.14 and Corollary 3.15, the partitions $[n] = \dot{\bigcup}_{i \in I} D_i$ and $[n] = \dot{\bigcup}_{i \in I} D'_i$ that are induced by the symmetric matrices Z and Z' are X-related and stable for A and for B, respectively. We need to show that $|D_i| = |D'_i|$ and that the partitions also agree w.r.t. the parameters s_{ij}.

By Lemma 3.6 we have $|D_i| = |D'_i|$ and $\mathbf{d}_i = X\mathbf{d}'_i$ and $\mathbf{d}'_i = X^t \mathbf{d}_i$, where $\mathbf{d}_i$ and $\mathbf{d}'_i$ for $i \in I$ are the characteristic vectors of the two partitions. Thus for all $i, j \in I$,

$$(\mathbf{d}'_i)^t B \mathbf{d}'_j = (X^t \mathbf{d}_i)^t B X^t \mathbf{d}_j = \mathbf{d}_i^t X B X^t \mathbf{d}_j = \mathbf{d}_i^t A X X^t \mathbf{d}_j = \mathbf{d}_i^t A Z \mathbf{d}_j = \mathbf{d}_i^t A \mathbf{d}_j,$$

where the last equality follows from the fact that $\mathbf{d}_j$ is an eigenvector of Z with eigenvalue 1 by Corollary 3.12. Note that $\mathbf{d}_i^t A \mathbf{d}_j$ is the number of edges of $\mathcal{A}$ from D_i to D_j. By stability of the partition, we have $s_{ij}^A = \mathbf{d}_i^t A \mathbf{d}_j / |D_i|$ and similarly $s_{ij}^B = (\mathbf{d}'_i)^t B \mathbf{d}'_j / |D'_i|$, so that $s_{ij}^A = s_{ij}^B$. ◀

4.2 L^2-equivalence and boolean linear equations

W.r.t. an adjacency matrix $A \in \mathbb{B}^{n,n}$, a boolean stable partition $[n] = \dot{\bigcup}_{i \in I} D_i$ has as parameters just the boolean values ι_{ij}^A defined by $\iota_{ij}^A = 0$ if $A_{D_i D_j} = 0$ and $\iota_{ij}^A = 1$ otherwise. Boolean (row-)stability of the partition for A implies that $\iota_{ij}^A = 1$ if, and only if, for each individual $d \in D_i$ there is at least one $d' \in D_j$ such that $A_{dd'} = 1$.

To capture the situation of 2-pebble game equivalence, though, we now need to work with similar partitions that are stable both w.r.t. A and w.r.t. to the adjacency matrix A^c of the complement of the graph with adjacency matrix A. Here the complement of a graph

$\mathcal{A}$ is the graph $\mathcal{A}^c$ with the same vertex set as $\mathcal{A}$ obtained by replacing edges by non-edges and vice versa. Hence $A^c_{aa'} = 1$ if $A_{aa'} = 0$ and $a \neq a'$, and $A^c_{aa'} = 0$ otherwise. While a partition in the sense of real arithmetic is stable for A if, and only if, it is stable for A^c, this is no longer the case for boolean arithmetic. Let us call a partition that is boolean stable for both A and A^c, *boolean bi-stable* for A.

Then the following captures the situation of two graphs that are 2-pebble game equivalent. We note that 2-pebble equivalence is a very rough notion of equivalence, if we look at just simple undirected graphs – but the concepts explored here do have natural extensions to coloured, directed graphs, and form the basis for the analysis of k-pebble equivalence, which is non-trivial even for simple undirected graphs.

L^2-equivalence of two graphs does not imply that the graphs have the same size. In the following, we always assume that $\mathcal{A}, \mathcal{B}$ are graphs with vertex sets $[m], [n]$ respectively and that $A \in \mathbb{B}^{m,m}$ and $b \in \mathbb{B}^{n,n}$ are their adjacency matrices. We call two bi-stable partitions $[m] = \dot{\bigcup}_{i \in I} D_i$ for A (and A^c) and $[n] = \dot{\bigcup}_{i \in J} D'_i$ for B (and B^c) *b-equivalent* if $I = J$ and $\iota^A_{ij} = \iota^B_{ij}$ and $\iota^{A^c}_{ij} = \iota^{B^c}_{ij}$ and for all $i, j \in I$. Note that b-equivalence does not imply $|D_i| = |D'_i|$.

▶ **Lemma 4.4.** *$\mathcal{A}$ and $\mathcal{B}$ are L^2-equivalent if, and only if, there are b-equivalent bi-stable partitions $[m] = \dot{\bigcup}_{i \in I} D_i$ for A and $[n] = \dot{\bigcup}_{i \in J} D'_i$ for B.*

▶ **Definition 4.5.** $\mathcal{A}$ and $\mathcal{B}$ are *boolean isomorphic*, $\mathcal{A} \approx_{\mathrm{bool}} \mathcal{B}$, if there is some boolean matrix X without null rows or columns such that $AX = XB$ and $A^cX = XB^c$.

▶ **Theorem 4.6.** *Two graphs are L^2-equivalent if, and only if, they are boolean isomorphic.*

5 Relaxations in the style of Sherali–Adams

In this section we refine the connection between the Sherali–Adams hierarchy of LP relaxation of the integer linear program ISO to equivalence in the finite variable counting logics. Throughout this section, our parameter $k \geqslant 2$ is the number of variables available in the logics C^k or L^k. As before, $\mathcal{A}$ and $\mathcal{B}$ are graphs with vertex sets $[m]$ and $[n]$, respectively, and A and B are their adjacency matrices.

The *level-$(k-1)$ Sherali–Adams relaxation* of the integer linear program ISO is the following linear program in the variables X_p for all $p \subseteq [m] \times [n]$ of size $|p| < k$. We write $p^\frown ab$ for the extension of p by the pair (a, b) (which need not be a proper extension).

ISO$(k-1)$	$X_\emptyset = 1$ and $X_p = \sum_{b'} X_{p^\frown ab'} = \sum_{a'} X_{p^\frown a'b}$ for $\ell := |p| + 1 < k, a \in [m], b \in [n]$	} $\textsc{Cont}(\ell)$ for $\ell < k$
	$\sum_{a'} A_{aa'} X_{p^\frown a'b} = \sum_{b'} X_{p^\frown ab'} B_{b'b}$ for $\ell := |p| + 1 < k, a \in [m], b \in [n]$	} $\textsc{Comp}(\ell)$ for $\ell < k$
	$X_p \geq 0$ for $|p| \leq k - 1$	

We call the equations $\textsc{Comp}(\ell)$ for $1 \leq \ell < k$ *comaptibility equations* and the equations $\textsc{Cont}(\ell)$ for $0 \leq \ell < k$ *continuity equations*, where we let $\textsc{Cont}(0)$ be the equation $X_\emptyset = 1$. We will also consider these equations independently of ISO$(k-1)$, as in the next lemma.

▶ **Lemma 5.1.** *If $\mathcal{A} \equiv_{\mathsf{C}}^k \mathcal{B}$, then there is a non-negative solution (X_p) for the combination of the continuity equations* $\text{Cont}(\ell)$ *of levels $\ell \leqslant k$ (!) with the compatibility equations* $\text{Comp}(\ell)$ *of levels $\ell < k$.*

Proof. For tuples $\mathbf{a} \in [m]^\ell$ and $\mathbf{b} \in [n]^\ell$ of length $\ell \leq k$, we write $\text{tp}(\mathbf{a}) = \text{tp}(\mathbf{b})$ if $\mathcal{A}, \mathbf{a}$ and $\mathcal{B}, \mathbf{b}$ satisfy the same C^k-formulae $\varphi(\mathbf{x})$ (equality of C^k-types). We write $p = \mathbf{ab}$ to indicate that p consist of the pairs $a_i b_i$ of corresponding entries in these tuples.

To define the solution, we let $X_\emptyset := 1$. For $p = \mathbf{ab}$, we let $X_p = 0$ if $\text{tp}(\mathbf{a}) \neq \text{tp}(\mathbf{b})$, and we let $X_p := 1/\#_{\mathbf{b}'}(\text{tp}(\mathbf{a}) = \text{tp}(\mathbf{b}'))$ otherwise. Tedious but straightforward calculations show that this indeed defines a solution of the desired equations. ◀

Thus in particular, if $\mathcal{A}$ and $\mathcal{B}$ are C^k-equivalent then the system $\text{ISO}(k-1)$ has a solution. Unfortunately, the converse does not hold (as we will see later). The solvability of $\text{ISO}(k-1)$ only implies a weaker equivalence between $\mathcal{A}$ and $\mathcal{B}$, which we call $\mathsf{C}^{<k}$-equivalence. It is defined in terms of a game, the *weak bijective k-pebble game* on $\mathcal{A}, \mathcal{B}$. The game is played by two players. Positions of the game are sets $p \subseteq [m] \times [n]$ of size $|p| \leq k-1$, and the initial position is $\emptyset$. A single round of the game, starting in position p, is played as follows.

1. If $|p| = k-1$, player **I** selects a pair $ab \in p$. If $|p| < k-1$, he omits this step.
2. Player **II** selects a bijection between $[m]$ and $[n]$. If no such bijection exists, i.e., if $m \neq n$, the game ends and player **II** loses.
3. Player **I** chooses a pair $a'b'$ from this bijection.
4. If $p^+ := p^\frown a'b'$ is a local isomorphism then the new position is

$$p' := \begin{cases} (p \backslash ab)^\frown a'b' & \text{if } |p| = k-1, \\ p^\frown a'b' & \text{if } |p| < k-1. \end{cases}$$

Otherwise, the play ends and player **II** loses.

Player **II** wins a play if it lasts forever. Structure $\mathcal{A}$ and $\mathcal{B}$ are $\mathsf{C}^{<k}$-*equivalent*, $\mathcal{A} \equiv_{\mathsf{C}}^{<k} \mathcal{B}$, if player **II** has a winning strategy for the game.

Note that the weak bijective k-pebble game requires more of the second player than the bijective $(k-1)$-pebble game, because p^+ rather than just p' is required to be a local isomorphism. On the other hand, it requires less than the bijective k-pebble game: the bijective k-pebble game precisely requires the second player to choose the bijection without prior knowledge of the pair ab that will be removed from the position. A strategy for player **II** in the weak version is good for the usual version if it is fully symmetric or uniform w.r.t. the pebble pair that is going to be removed. However, this is only relevant if $k \geq 3$. The weak bijective 2-pebble game and the bijective 2-pebble game are essentially the same.

The core of the proof of the following is analogous to that of Theorem 4.3.

▶ **Theorem 5.2.** *$\mathcal{A} \equiv_{\mathsf{C}}^{<k} \mathcal{B}$ if, and only if,* $\text{ISO}(k-1)$ *has a solution.*

▶ Remark 5.3. The weak bijective k-pebble game is equivalent to a bisimulation-like game with $k-1$ pebbles where in each round the first player may slide a pebble along an edge of one of the graphs and the second player has to respond by sliding the corresponding pebble along an edge of the other graph. In this version, the game coresponds to the $(k-1)$-pebble sliding game introduced by Atserias and Maneva [1].

We will see in the next section that $\mathsf{C}^{<k}$-equivalence neither coincides with C^{k-1}-equivalence nor with C^k-equivalence. Thus it remains to give a characterisation of C^k-equivalence. By the previous theorem and the observation that C^k-equivalence is situated between $\mathsf{C}^{<k}$-equivalence and $\mathsf{C}^{<k+1}$-equivalence, we know that we need a linear program

that is "between" ISO$(k-1)$ and ISO(k). Surprisingly, we obtain such a linear program by combining the two in a very simple way: we take the continuity equations from ISO(k) and the compatibility equations from ISO$(k-1)$. Thus the resulting linear program, which we call ISO$(k-1/2)$, has variables X_p for all $p \subseteq [m] \times [n]$ of size $|p| \leq k$ and consists of the equations $\textsc{Cont}(\ell)$ for $\ell \leq k$ and the equations $\textsc{Comp}(\ell)$ for $\ell \leq k-1$, together with the non-negativity constraints $X_p \geq 0$. So Lemma 5.1 proves one implication of the theorem.

▶ **Theorem 5.4.** *$\mathcal{A} \equiv^k_{\mathsf{C}} \mathcal{B}$ if, and only if,* ISO$(k-1/2)$ *has a solution.*

5.1 Boolean arithmetic and L^k-equivalence

We saw in Section 4.2 that equations, which are direct consequences of the basic continuity and compatibility equations w.r.t. the adjacency matrices A and B, may carry independent weight in their boolean interpretation. This is no surprise, because the boolean reading is much weaker, especially due to the absorptive nature of $\vee$, which unlike $+$ does not allow for inversion. $AX = XB$ for doubly stochastic X and $A, B \in \mathbb{B}^{n,n}$ implies $A^c X = XB^c$.

We now augment the boolean requirements by corresponding boolean equations that express

(a) compatibility also w.r.t. A^c and B^c, as in boolean fractional isomorphism,

(b) the new constraint $X_p = 0$ whenever p is not a local bijection.

In the presence of the continuity equations, which force monotonicity, it suffices for (b) to stipulate $X_{aa'bb'} = 0$ for all $a, a' \in [m]$, $b, b' \in [n]$ such that *not* $a = a' \Leftrightarrow b = b'$. This is captured by the constraint $\textsc{Match}(2)$ below. So we now use the following boolean version of the Sherali–Adams hierarchy ISO$(k-1)$ and ISO$(k-1/2)$ for $k \geq 2$.

B-ISO$(k-1)$

$X_\emptyset = 1$ and $X_p = \sum_{b'} X_{p^\frown ab'} = \sum_{a'} X_{p^\frown a'b}$ for $\lvert p\rvert < k, a \in [m], b \in [n]$	$\textsc{Cont}(\ell)$ for $\ell < k$
$X_{ab^\frown ab'} = 0 = X_{ab^\frown a'b}$ for $a \neq a' \in [m], b \neq b' \in [n]$	$\textsc{Match}(2)$
$\sum_{a'} A_{aa'} X_{p^\frown a'b} = \sum_{b'} X_{p^\frown ab'} B_{b'b}$ for $\lvert p\rvert < k-1, a \in [m], b \in [n]$	$\textsc{Comp}(\ell)$ for $\ell < k$
$\sum_{a'} A^c_{aa'} X_{p^\frown a'b} = \sum_{b'} X_{p^\frown ab'} B^c_{b'b}$ for $\lvert p\rvert < k-1, a \in [m], b \in [n]$	$\textsc{Comp}(\ell)^c$ for $\ell < k$

For B-ISO$(k-1/2)$ we require $\textsc{Cont}(\ell)$ for all $\ell \leqslant k$, i.e., additionally for $\ell = k$.

▶ Remark 5.5. B-ISO$(k-1)$ and B-ISO$(k-1/2)$ are systems of boolean equations, and the reader may wonder whether they can be solved efficiently. At first sight, it may seem NP-complete to solve such systems (just like boolean satisfiability). However, our systems consist of "linear" equations of the forms $\sum_{i \in I} X_i = \sum_{j \in J} X_j$ and $\sum_{i \in I} X_i = 0$ (which is actually a special case of the first for $J = \emptyset$) and $\sum_{i \in I} X_i = 1$. It is an easy exercise to prove that such systems of linear boolean equations can be solved in polynomial time.

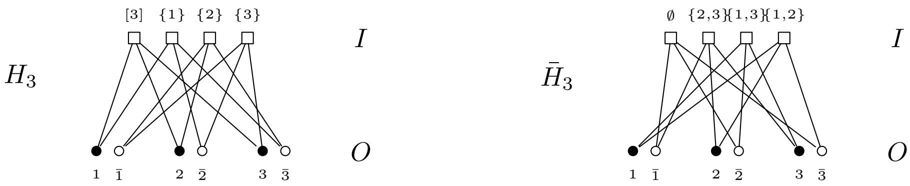

Figure 1 The Cai–Fürer–Immerman gadgets.

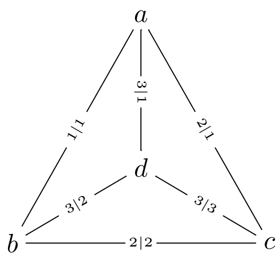

Figure 2 Structure $\mathcal{A}$.

We define a *weak k-pebble* game as a straightforward adaptation of the weak bijective k-pebble game to the setting without counting, and we denote weak k-pebble equivalence as in $\mathcal{A} \equiv_{\mathsf{L}}^{<k} \mathcal{B}$.

▶ **Theorem 5.6.** *W.r.t. boolean arithmetic:*
(a) B-ISO$(k-1)$ *has a solution if, and only if,* $\mathcal{A} \equiv_{\mathsf{L}}^{<k} \mathcal{B}$.
(b) B-ISO$(k-1/2)$ *has a solution if, and only if,* $\mathcal{A} \equiv_{\mathsf{L}}^{k} \mathcal{B}$.

6 The gap

Based on a construction due to Cai, Fürer, and Immerman [4], for $k \geqslant 3$ we construct graphs showing that $\mathcal{A} \equiv_{\mathsf{C}}^{<k} \mathcal{B} \not\Rightarrow \mathcal{A} \equiv_{\mathsf{C}}^{k} \mathcal{B}$, and that $\mathcal{A} \equiv_{\mathsf{C}}^{k-1} \mathcal{B} \not\Rightarrow \mathcal{A} \equiv_{\mathsf{C}}^{<k} \mathcal{B}$.

▶ **Example 6.1.** For every $k \geqslant 3$, there are graphs $\mathcal{A}$ and $\mathcal{B}$ such that $\mathcal{A} \equiv_{\mathsf{C}}^{k-1} \mathcal{B}$ but $\mathcal{A} \not\equiv_{\mathsf{C}}^{<k} \mathcal{B}$.

We describe the graphs $\mathcal{A}$ and $\mathcal{B}$ for $k = 4$; the adaptation of the construction to other k is straightforward. The graphs are the straight and the twisted version of the Cai–Fürer–Immerman companions of the 4-clique.

We use copies of the standard degree 3 gadget H_3 and its dual $\bar{H}_3$ shown in Figure 1. We think of these as coloured graphs where the colours distinguish inner vertices (marked I) as well as outer vertices (marked O) as well as the three pairs of outer vertices. This is without loss of generality, since we may eliminate colours, e.g., by attaching simple, disjoint paths of different lengths to the members of each group of vertices. The non-trivial automorphisms of this decorated variant of H_3 and $\bar{H}_3$ precisely allow for simultaneous swaps within exactly two pairs of outer vertices.

Let $\mathcal{A}$ consist of four decorated copies of H_3, copies a, b, c, d say, that are linked by edges in corresponding outer nodes as shown in Figure 2. $\mathcal{B}$ consists of three decorated copies of H_3 (labelled a, b, c) and one of $\bar{H}_3$ (labelled d), and linked in the same manner.

It can be shown that player **I** has a winning strategy in the weak bijective 4-pebble game on $\mathcal{A}, \mathcal{B}$, whereas player **II** has a winning strategy in the bijective 3-pebble game.

▶ **Example 6.2.** For every $k \geqslant 3$, there are graphs $\mathcal{A}$ and $\mathcal{B}$ such that $\mathcal{A} \equiv_{\mathsf{C}}^{<k} \mathcal{B}$ but $\mathcal{A} \not\equiv_{\mathsf{C}}^{k} \mathcal{B}$.

We describe the graphs for $k = 3$. We use variants of $\mathcal{A}$ and $\mathcal{B}$ as in the last example, but with one marked inner node: in both $\mathcal{A}$ and $\mathcal{B}$ we mark the inner node $(a, [3])$ by a new colour (which can be eliminated by attaching a path of some characteristic length, as observed above). We denote these modified structures as $\mathcal{A}_*$ and $\mathcal{B}_*$.

Then it can be shown that player **I** has a winning strategy in the bijective 3-pebble game on $\mathcal{A}_*, \mathcal{B}_*$, whereas player **II** has a winning strategy in the weak bijective 3-pebble game.

References

1 A. Atserias and E. Maneva. Sherali–Adams relaxations and indistinguishability in counting logics. In *Innovations in Theoretical Computer Science (ITCS)*, 2012.

2 D. Bienstock and N. Ozbay. Tree-width and the Sherali-Adams operator. *Discrete Optimization*, 1:13–21, 2004.

3 J. Buresh-Oppenheim, N. Galesi, S. Hoory, A. Magen, and T. Pitassi. Rank bounds and integrality gaps for cutting planes procedures. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 318–327, 2003.

4 J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12:389–410, 1992.

5 M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for Sherali-Adams relaxations. In *Proceedings of the 41st ACM Symposium on Theory of Computing*, pages 283–292, 2009.

6 H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer-Verlag, 2nd edition, 1999.

7 E. Grädel, P.G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M.Y. Vardi, Y. Venema, and S. Weinstein. *Finite Model Theory and Its Applications.* Springer-Verlag, 2007.

8 E. Grädel and M. Otto. Inductive definability with counting on finite structures. In E. Börger, G. Jäger, H. Kleine Büning, S. Martini, and M.M. Richter, editors, *Computer Science Logic, 6th Workshop, CSL '92, San Miniato 1992, Selected Papers*, volume 702 of *Lecture Notes in Computer Science*, pages 231–247. Springer-Verlag, 1993.

9 M. Grohe. Fixed-point definability and polynomial time on graphs with excluded minors. In *Proceedings of the 25th IEEE Symposium on Logic in Computer Science*, 2010.

10 M. Grohe and M. Otto. Pebble games and linear equations, 2012. Full version of this paper. Available on arXiv, arXiv:1204.1990.

11 L. Hella. Logical hierarchies in PTIME. In *Proceedings of the 6th IEEE Symposium on Logic in Computer Science*, pages 360–368, 1992.

12 N. Immerman. *Descriptive Complexity.* Springer-Verlag, 1999.

13 N. Immerman and E. Lander. Describing graphs: A first-order approach to graph canonization. In A. Selman, editor, *Complexity theory retrospective*, pages 59–81. Springer-Verlag, 1990.

14 B. Laubner. Capturing polynomial time on interval graphs. In *Proceedings of the 25th IEEE Symposium on Logic in Computer Science*, pages 199–208, 2010.

15 L. Libkin. *Elements of Finite Model Theory.* Springer-Verlag, 2004.

16 C. Mathieu and A. Sinclair. Sherali-Adams relaxations of the matching polytope. In *Proceedings of the 41st ACM Symposium on Theory of Computing*, pages 293–302, 2009.

17 M. Otto. *Bounded variable logics and counting – A study in finite models*, volume 9 of *Lecture Notes in Logic.* Springer-Verlag, 1997.

18 M. Ramana, E. Scheinerman, and D. Ullman. Fractional isomorphism of graphs. *Discrete Mathematics*, 132:247–265, 1994.

19 G. Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 593–602, 2008.

20 G. Schoenebeck, L. Trevisan, and M. Tulsiani. Tight integrality gaps for Lovász-Schrijver LP relaxations of vertex cover and max cut. In *Proceedings of the 39th ACM Symposium on Theory of Computing*, pages 302–310, 2007.

21 H.D. Sherali and W.P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3:411, 1990.