

DISMANTLABILITY, CONNECTEDNESS, AND MIXING IN RELATIONAL STRUCTURES

RAIMUNDO BRICEÑO, ANDREI BULATOV, VÍCTOR DALMAU, AND BENOÎT LAROSE

ABSTRACT. The Constraint Satisfaction Problem (CSP) and its counting counterpart appears under different guises in many areas of mathematics, computer science, and elsewhere. Its structural and algorithmic properties have demonstrated to play a crucial role in many of those applications. For instance, in the decision CSPs, structural properties of the relational structures involved—like, for example, dismantlability—and their logical characterizations have been instrumental for determining the complexity and other properties of the problem. Topological properties of the solution set such as connectedness are related to the hardness of CSPs over random structures. Additionally, in approximate counting and statistical physics, where CSPs emerge in the form of spin systems, mixing properties and the uniqueness of Gibbs measures have been heavily exploited for approximating partition functions and free energy.

In spite of the great diversity of those features, there are some eerie similarities between them. These were observed and made more precise in the case of graph homomorphisms by Brightwell and Winkler, who showed that dismantlability of the target graph, connectedness of the set of homomorphisms, and good mixing properties of the corresponding spin system are all equivalent. In this paper we go a step further and demonstrate similar connections for arbitrary CSPs. This requires much deeper understanding of dismantling and the structure of the solution space in the case of relational structures, and new refined concepts of mixing introduced by Briceño. In addition, we develop properties related to the study of valid extensions of a given partially defined homomorphism, an approach that turns out to be novel even in the graph case. We also add to the mix the combinatorial property of finite duality and its logic counterpart, FO-definability, studied by Larose, Loten, and Tardif.

1. INTRODUCTION

The Constraint Satisfaction Problem (CSP) provides a powerful framework in a wide range of areas of mathematics, computer science, statistical physics, and elsewhere. The goal in a CSP is to find an assignment to variables from a given set that satisfies a number of given constraints. The counting version of the problem asks about the number of such assignments. The CSP however appears in different forms: as the standard one outlined above in AI and computer science [22], as the homomorphism problem in graph and model theory [26, 33], as conjunctive query evaluation in logic and database theory [38], as computing the partition function of a spin system in statistical physics [48] and related areas, like symbolic dynamics and coding [44, 46].

The CSP allows for many approaches of diverse nature, and every application field exploits some of its many facets: structural properties of constraints for complexity and algorithms, probabilistic properties and the topology of the solution space in Random CSP

2010 *Mathematics Subject Classification.* 08A70, 68Q87, 68R01 (Primary); 82B20, 05B45, 05C15 (Secondary).

Key words and phrases. Relational structure; constraint satisfaction problem; homomorphism; mixing properties; Gibbs measure.

and random structures, mixing properties in statistical physics and dynamical systems, decay of correlations and the uniqueness of probabilistic measures in approximate counting, and homomorphic duality and logical characterizations in model theory. In [13], it was observed that some of these properties are actually closely related, at least in the simple case of graph homomorphisms. In this paper we take this research direction a step further by extending Brightwell and Winkler’s results to the general CSP, and by refining and widening the range of the properties involved.

We start off with a brief introduction of the features of the CSP considered in this paper. Afterwards, we provide a detailed account of the necessary background and a description of our results. Every CSP involves a set of variables and a domain, a set of possible values for the variables. Assumptions about these two sets differ in different areas. The most studied case in combinatorics and complexity theory is when both sets are finite. However, many interesting problems such as scheduling and temporal and spatial reasoning involve infinite domains; see also the extensive literature on infinite CSPs (for example, [6] and the references therein). In other cases such as in statistical physics, it is natural to choose the set of variables to be infinite (a lattice, for example). Then, it is also natural to study probability distributions over such assignments—where *Gibbs measures* and the problem of their (*non-*)*uniqueness* appear naturally [30]—and also study quantities such as *entropy* and *free energy* [5, 9].

Following [26], CSPs can be formulated as the problem of deciding the existence of a homomorphism from a finite relational structure \mathbb{G} to a target relational structure \mathbb{H} , where \mathbb{G} and \mathbb{H} encode the variables and the values of the CSP. The complexity of this problem, especially the case when \mathbb{H} is a fixed finite relational structure, has received a lot of attention, culminating with the proof of the Feder-Vardi conjecture [15, 56], which asserts that every CSP is either in PTIME or NP-complete. In the present paper we focus as well on the case when \mathbb{H} is finite, although our main focus is not algorithmic but rather structural. In particular, we are interested in studying the space $\text{Hom}(\mathbb{G}, \mathbb{H})$ of homomorphisms from \mathbb{G} to \mathbb{H} . Furthermore, following [13], we consider homomorphisms from both finite and infinite relational structures \mathbb{G} (although [13] only considers graphs), a flexibility that turns out to be useful to see different aspects of homomorphism spaces $\text{Hom}(\mathbb{G}, \mathbb{H})$ that otherwise would be meaningless. We note here that the case where \mathbb{G} is restricted has also been studied. In particular, there is an important line of research that studies the homomorphism existence problem when \mathbb{G} belongs to a prescribed set of relational structures and \mathbb{H} is an arbitrary relational structure (e.g., see [31]). The case where both \mathbb{G} and \mathbb{H} are restricted has also been investigated (e.g., see [20]).

There is a vast quantity of literature concerning graph homomorphisms and their properties through the lens of statistical physics [7, 23, 12]. In this context, it is very common to encode a *spin system* as a pair of relational structures \mathbb{G} and \mathbb{H} , where \mathbb{G} contains a set of variables/particles and \mathbb{H} contains the set of values/spins that each particle could take, imposing *hard constraints* on them, i.e., disregarding configurations of values that do not satisfy all the given constraints. In practical terms, all this reduces to the study of the maps from \mathbb{G} to \mathbb{H} that are homomorphisms, individually and as a set. In particular, many important parameters of a spin system such as free energy and entropy can be learned from studying such a set of homomorphisms.

In [13], it was proved that many of the properties of graph homomorphisms used in the above areas are equivalent to a single structural property of graphs, namely, *dismantlability*. In this paper we follow a similar approach and study properties of CSPs over general

relational structures that we put into basically three categories: (1) dismantlability, (2) connectedness, and (3) mixing. Furthermore, as a consequence of our results, we established a connection with a fourth notion not initially contemplated in [13]: (4) *finite duality*.

Dismantlability. A graph is said to be dismantlable if it can be reduced to a single vertex by removing vertices whose neighborhood is contained in the neighborhood of some other vertex. Such transformations are called *folds*, and they can be viewed as *retractions* of a very particular kind. Ultimately, this kind of action allow us to reduce \mathbb{H} and $\text{Hom}(\mathbb{G}, \mathbb{H})$ by replacing the appearance of certain spins with others that have equal or more freedom of allocation. Dismantlable graphs were introduced in [50], based on ideas already present in [36] in the context of lattices, and have been intensively studied in combinatorics. Dismantlability can be generalized in a natural way to relational structures. Indeed, some variants of this notion have been used in the study of CSPs. In particular, dismantlability has been applied in [16] to the problem of enumerating all solutions of $\text{Hom}(\mathbb{G}, \mathbb{H})$ with polynomial delay. Also, it has played a major role in the study of CSPs definable in first-order logic [21, 42].

Connectedness. When \mathbb{G} is finite, it is often useful to convert $\text{Hom}(\mathbb{G}, \mathbb{H})$ into an auxiliary graph and explore the connectivity properties of it. The set of edges of $\text{Hom}(\mathbb{G}, \mathbb{H})$ can be defined in a variety of ways, usually the most suitable to the problem at hand. For example, it is common to say that two elements from $\text{Hom}(\mathbb{G}, \mathbb{H})$ are close (and therefore adjacent in the graph) if the Hamming distance between them is smaller than a certain threshold. The particular case when this threshold is 1 has been intensively studied, motivated initially by the fact that the connectedness of the solution space for SAT problems over random instances is linked to the performance of standard satisfiability algorithms, such as WalkSAT or DPLL [1, 40]. This has given rise to a general framework called *reconfiguration* [35] that goes way beyond homomorphisms (see also [49] for a recent survey). Work in this area encompasses both structural questions (under which conditions is $\text{Hom}(\mathbb{G}, \mathbb{H})$ connected?) and algorithmic ones (what is the complexity of, deciding, given \mathbb{G} and \mathbb{H} as input, whether $\text{Hom}(\mathbb{G}, \mathbb{H})$ is connected? Its diameter? The shortest path between two given members of $\text{Hom}(\mathbb{G}, \mathbb{H})$? Etc.). In the context of spin systems, the connectedness of $\text{Hom}(\mathbb{G}, \mathbb{H})$ is related to processes that consists on periodically updating the spin of a single or a small set of particles (e.g., irreducibility of *Glauber dynamics*). We also consider an alternative way to define adjacency in $\text{Hom}(\mathbb{G}, \mathbb{H})$ via *links* as in [42]. This notion of adjacency is linked to the so-called *finite duality property*, which is another of the main themes of our work.

Mixing. Mixing properties have been intensively studied in statistical physics and related areas (see [2, 4, 8, 10, 11, 19, 51]), and are usually applied when the set of particles in \mathbb{G} is very large or infinite. In this case, it can be very useful to be able to “glue” together partial homomorphisms, provided their domains are far from each other. There are several properties that formalize this phenomenon and it is common to establish hierarchies among them. More concretely, given a metric in \mathbb{G} , it is natural to ask whether there exists some uniform gap such that for any two subsets V and W of particles sufficiently far apart (in terms of the gap), and for any pair of homomorphisms $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$, we can find a third one, γ , such that restriction of γ to V and W coincides with the restrictions of ϕ and ψ on V and W , respectively. On the contrary, whenever the information content of a given set (at least partially) determines the information content of another set (i.e., the possible values that the variables on it can take), no matter how far it is, such a phenomenon has been called *long range action* in previous work [14].

Similar phenomena are used in the related area of approximate computing of partition functions, where many algorithms are based on decay of correlations between values of remote elements of \mathbb{G} , which allows for approximation of partition functions based only on local neighborhoods of variables [24, 54].

Finite duality and logic characterizations. Homomorphism duality often helps to design a solution algorithm for a CSP or establish useful properties. A relational structure \mathbb{H} is said to have *homomorphism duality* if there is a set \mathcal{O} of relational structures—called *obstructions*—such that a relational structure \mathbb{G} has a homomorphism to \mathbb{H} if and only if no relational structure from \mathcal{O} is homomorphic to \mathbb{G} . Sometimes the set of obstructions is very simple, say, any bipartite graph has homomorphic duality, where \mathcal{O} is the set of all odd cycles. If \mathcal{O} can be chosen finite, we say that \mathbb{H} has *finite duality*.

Homomorphism duality is closely related to another property of CSPs. Let \mathcal{L} be a logic language such as first order, second order, etc. The problem of deciding homomorphisms to a relational structure \mathbb{H} is said to be *expressible in \mathcal{L}* if there is a formula Φ in the language \mathcal{L} such that \mathbb{G} has a homomorphism to \mathbb{H} if and only if Φ is true on \mathbb{G} . It is known, for instance, that \mathbb{H} has a set of obstructions consisting of relational structures of bounded treewidth if and only if the corresponding homomorphism problem is expressible in Datalog [26], or that \mathbb{H} has finite duality if and only if the corresponding problem is expressible in first order logic [3]. For a survey on dualities for CSP, see [17].

Our results. In this paper we extend the approach taken in [13] to the case of homomorphisms between relational structures, that includes, besides graphs, very natural special cases, like *digraphs*, *hypergraphs*, and *shifts of finite type* (see Example 2.2). In particular, we show (Theorem 3.2) that, for a relational structure \mathbb{H} , the three following conditions are equivalent: (A) \mathbb{H}^2 dismantles to a substructure of its diagonal, that is, the substructure of \mathbb{H}^2 induced by the set $\{(a, a) \mid a \in \mathbb{H}\}$; (B) for any \mathbb{G} , the homomorphism graph $\text{Hom}(\mathbb{G}, \mathbb{H})$ is connected; and (C) for any \mathbb{G} , the space $\text{Hom}(\mathbb{G}, \mathbb{H})$ satisfies certain mixing properties. Furthermore, there are several contexts where it is natural to work by forcing a certain subset of variables to take each of them a particular value and work with the remaining ones. This idea inspired a refined version of Theorem 3.2, namely Theorem 3.6, which can be regarded as the study of *boundary long range actions*, i.e., long range action phenomena where some boundary configuration is fixed. In particular, even in the case of graphs, this leads to new results that characterizes conditions on \mathbb{H} that guarantee connectedness and mixing properties when some variables are fixed to take particular values.

As a byproduct of our results, we obtain two applications. On the one hand, we establish a link with *strong spatial mixing* (e.g., see [24]) and *topological strong spatial mixing* (introduced in [9]). These two last properties have played an important role in the development of deterministic approximate counting algorithms. In this paper we address the following question: What fixed targets \mathbb{H} are suitable for both of these properties to hold for any \mathbb{G} ? On the other hand, we establish a connection with finite duality, which allows us to reprove the main theorem in [42]. We hope that our work opens the possibility of developing new counting techniques based on this approach in a very general setting.

Organization. The paper is organized as follows: In Section 2, we introduce most of the objects and terminology relevant to this work. In particular, we define relational structures, the natural maps and operations on them, and some useful constructions. In addition, we introduce the three relevant properties around homomorphisms: dismantlability, connectedness, and mixing. Next, in Section 3, we state the two main theorems of this paper,

Theorem 3.2 and its refinement, Theorem 3.6, and we illustrate how these two theorems relate to the work of Brightwell and Winkler and generalizes it. Next, in Section 4, we prove all the equivalences which constitute Theorem 3.6. In Section 5, we define Gibbs measures on sets of homomorphisms and explore the consequences of our results in relation to spatial mixing properties of such measures. Finally, in Section 6, we establish a meaningful connection between our results, topological strong spatial mixing, and finite duality of relational structures.

2. PRELIMINARIES

Let H be a countable (finite or denumerable) set and k a positive integer. The set of k -tuples over H is denoted by H^k . A (k -ary) **relation** R over H is a subset $R \subseteq H^k$. The elements of a relation R will be denoted in boldface, e.g., \mathbf{a} , \mathbf{b} , etc., and a_i will denote the i th entry of \mathbf{a} for $1 \leq i \leq k$.

Given another countable set G and a map $\phi : G \rightarrow H$, for a k -tuple \mathbf{a} over G we shall use $\phi(\mathbf{a})$ to denote the k -tuple over H obtained after applying ϕ to \mathbf{a} componentwise. If $V \subseteq G$, we will denote by $\phi|_V$ the restriction of ϕ to V . Furthermore, if ψ is another map with domain H , we shall use $\psi \circ \phi$ to denote the composition of ψ with ϕ , i.e., the map $x \mapsto \psi(\phi(x))$.

A **signature** τ is a collection of relation symbols R , each of them with an associated arity. For a given signature τ , a **relational structure** (with signature τ)—or simply, a **τ -structure**— \mathbb{H} consists of a countable set H called the **universe** of \mathbb{H} and a relation $R(\mathbb{H}) \subseteq H^k$ for each $R \in \tau$. Thus relation symbols serve as markers that allow one to associate relations in different structures with the same signature (e.g., see [34]).

We shall use the same capital letter to denote the universe of a τ -structure, e.g., H is the universe of \mathbb{H} . We will usually consider τ to be a fixed signature, and \mathbb{G} and \mathbb{H} to be τ -structures with universes G and H , respectively.

The **degree** of an element in a relational structure is defined as the number of tuples in which it occurs. A relational structure is **locally finite** if every element in its universe has finite degree and is said to be **finite** if its universe is finite.

Remark 2.1. A digraph \mathbb{G} (with self-loops allowed) is a very particular case of a relational structure, where the signature τ consists of a unique relation symbol E of arity 2. Moreover, graphs correspond to the digraph case where the edge-set $E(\mathbb{G})$ is a symmetric relation.

A map $\phi : G \rightarrow H$ is said to be a **homomorphism** from \mathbb{G} to \mathbb{H} if, for every relation symbol $R \in \tau$,

$$\mathbf{a} \in R(\mathbb{G}) \implies \phi(\mathbf{a}) \in R(\mathbb{H}).$$

We will denote by $\text{Hom}(\mathbb{G}, \mathbb{H})$ the set of all homomorphisms from \mathbb{G} to \mathbb{H} .

Example 2.2. A particular example of CSPs that cannot be represented in the setting of Brightwell and Winkler (that is, as homomorphisms of graphs) is the case of *shifts of finite type*, a fundamental object in dynamical systems and probability [44, 45, 46].

Given a positive integer d , consider the signature $\tau = \{R_1, \dots, R_d\}$, where R_i is a 2-ary relation for all $1 \leq i \leq d$. We consider two τ -structures \mathbb{G} and \mathbb{H} . Here, \mathbb{G} will be an infinite relational structure with universe $G = \mathbb{Z}^d$ and relations $R_i(\mathbb{G})$, $1 \leq i \leq d$, representing the usual d -dimensional hypercubic lattice and the adjacency of pairs of elements in it. On the other hand, \mathbb{H} will be a finite relational structure with universe H and $R_i(\mathbb{H})$, $1 \leq i \leq d$, will represent pairs of “colors” from H that are allowed to be adjacent in the canonical i th direction of the lattice. Then, $X = \text{Hom}(\mathbb{G}, \mathbb{H})$ is known as a **d -dimensional nearest-neighbor shift of finite type**, a set of colorings of \mathbb{Z}^d with not necessarily *isotropic* adjacency rules

(i.e., we do not need to have the same restrictions in every direction), and any such object can be represented in this way (see also [53] for a similar approach). It is well-known that every tiling of \mathbb{Z}^d can be represented as a nearest-neighbor shift of finite type and vice versa [52].

More generally, it can be checked that for any shift of finite type X defined on a countable group Γ (see [18] for an introduction to the subject), i.e., not necessarily nearest-neighbor nor restricted to $\Gamma = \mathbb{Z}^d$, there exists a signature τ and τ -structures \mathbb{G} and \mathbb{H} such that $X = \text{Hom}(\mathbb{G}, \mathbb{H})$.

Example 2.3. Hypergraph homomorphisms can be also naturally encoded using relational structures. A **hypergraph** is a pair (V, E) where V is a set of elements called *nodes* and E is a collection of subsets of V called *hyperedges*. If $H = (V, E)$, $H' = (V', E')$ are hypergraphs, then a mapping $f : V \rightarrow V'$ is a homomorphism from H to H' if for every hyperedge $e \in E$, $\bigcup_{v \in e} \{f(v)\}$ is a hyperedge of E' . Now, we can encode H and H' as relational structures over a signature τ containing a unique relation symbol R of arity m , where m is the maximum size of any hyperedge in $E \cup E'$. In particular, H is encoded by the relational structure \mathbb{H} with universe V and

$$R(\mathbb{H}) = \{(v_1, \dots, v_m) \mid \{v_1, \dots, v_m\} \in E\}.$$

Analogously, H' can be encoded in a similar way. It is not difficult to see that every mapping $f : V \rightarrow V'$ is a homomorphism from H to H' if and only if it is a homomorphism of their associated relational structures.

A relational structure \mathbb{J} is a **substructure** of \mathbb{H} if $J \subseteq H$ and, for every relation symbol $R \in \tau$, we have that $R(\mathbb{J}) \subseteq R(\mathbb{H})$. Furthermore, if for every k -ary $R \in \tau$, we have that $R(\mathbb{J}) = R(\mathbb{H}) \cap J^k$, then we say that \mathbb{J} is the substructure of \mathbb{H} **induced by J** . If $J \subseteq H$ and $\phi : H \rightarrow J$ is a homomorphism acting as the identity on J , then ϕ is said to be a **retraction**.

The **product** of \mathbb{H}_1 and \mathbb{H}_2 , denoted $\mathbb{H}_1 \times \mathbb{H}_2$, is the τ -structure with universe $H_1 \times H_2$ where, for every k -ary relation symbol $R \in \tau$, we have that $R(\mathbb{H}_1 \times \mathbb{H}_2)$ consists of all tuples $((a_1, b_1), \dots, (a_k, b_k))$ with $(a_1, \dots, a_k) \in R(\mathbb{H}_1)$ and $(b_1, \dots, b_k) \in R(\mathbb{H}_2)$. We shall denote by \mathbb{H}^2 the product $\mathbb{H} \times \mathbb{H}$. The **projections** $\pi_1, \pi_2 : H^2 \rightarrow H$ are the maps $(a, b) \mapsto a$ and $(a, b) \mapsto b$, respectively, for $(a, b) \in H^2$. An element (a, b) of H^2 is **diagonal** if $a = b$. The **diagonal set** of H^2 , denoted $\Delta(H^2)$, is the set of its diagonal elements. Similarly, the **diagonal structure** of \mathbb{H}^2 , denoted $\Delta(\mathbb{H}^2)$, is the substructure of \mathbb{H}^2 induced by $\Delta(H^2)$. A substructure \mathbb{K} of \mathbb{H}^2 is **symmetric** whenever $(a, b) \in K$ if and only if $(b, a) \in K$. Notice that \mathbb{H}^2 is always symmetric.

In this paper, we will study properties of \mathbb{H} and how they relate to other properties of $\text{Hom}(\mathbb{G}, \mathbb{H})$ for arbitrary \mathbb{G} . We mainly consider three families of properties, namely, *dismantling* of \mathbb{H} , *connectedness* of some particular graphs with vertex set $\text{Hom}(\mathbb{G}, \mathbb{H})$, and *mixing* properties of $\text{Hom}(\mathbb{G}, \mathbb{H})$. By the end of the article, we also show how all these properties are related to a fourth one, namely, *finite duality*.

2.1. Dismantling. Let \mathbb{H} be a τ -structure and let a, b be elements in its universe H . We say that b **dominates** a (in \mathbb{H}) if for every k -ary $R \in \tau$, any $i \in \{1, \dots, k\}$, and any $(a_1, \dots, a_k) \in R(\mathbb{H})$ with $a_i = a$, we also have that

$$(a_1, \dots, a_{i-1}, b, a_{i+1}, \dots, a_k) \in R(\mathbb{H}).$$

Additionally, if $a \neq b$, then we say that a is **dominated** (in \mathbb{H}).

A sequence of τ -structures $\mathbb{J}_0, \dots, \mathbb{J}_\ell$ is a **dismantling sequence** if for every $0 \leq j < \ell$ there exist $a_j, b_j \in J_j$ such that b_j dominates a_j in \mathbb{J}_j , and \mathbb{J}_{j+1} is the substructure of \mathbb{J}_j induced by $J_j \setminus \{a_j\}$. In this case, we say that \mathbb{J}_0 **dismantles to** \mathbb{J}_ℓ . We can alternatively

denote a dismantling sequence by giving the initial τ -structure \mathbb{J}_0 and the sequence of elements $a_0, \dots, a_{\ell-1}$. We say that \mathbb{H} is **dismantlable** if it dismantles to a τ -structure such that its universe is a singleton.

Note that for every $0 \leq j < \ell$ there is a natural retraction r_j from \mathbb{J}_j to \mathbb{J}_{j+1} , where r_j maps a_j to b_j and acts as the identity elsewhere. We call such retractions a **fold**. By successive composition, one can define a retraction (namely, $r_{j'-1} \circ \dots \circ r_j$) from \mathbb{J}_j to $\mathbb{J}_{j'}$ for every $j \leq j'$.

Example 2.4. In the context of d -dimensional nearest-neighbor SFTs from Example 2.2, suppose that $d = 2$ and consider $\tau = \{R_1, R_2\}$ with R_i a 2-ary relation symbol for $i = 1, 2$. Then, we can consider τ -structure \mathbb{G} and \mathbb{H} such that

- the universe of \mathbb{G} is \mathbb{Z}^2 ,

$$R_1(\mathbb{G}) = \{(g, g + (1, 0)) : g \in \mathbb{Z}^2\}, \text{ and}$$

$$R_2(\mathbb{G}) = \{(g, g + (0, 1)) : g \in \mathbb{Z}^2\};$$

- the universe of \mathbb{H} is $H = \{a, b, c\}$,

$$R_1(\mathbb{H}) = \{(a, a), (a, b), (b, a), (b, b), (b, c), (c, b)\}, \text{ and}$$

$$R_2(\mathbb{H}) = \{(a, a), (a, b), (b, a), (b, b), (b, c), (c, a)\}.$$

We can notice that c folds to b , and b folds to a . We will see that the fact that \mathbb{H} can be dismantled to a single element will imply that $\text{Hom}(\mathbb{G}, \mathbb{H})$ has special properties. We also notice that any $\text{Hom}(\mathbb{G}, \mathbb{H})$ can be seen as a *Wang tiling* of \mathbb{Z}^2 and vice versa [52].

It is well known that if \mathbb{H} dismantles to some substructure \mathbb{K} , then this dismantling can be found in a greedy manner. Formally, we have the following lemma.

Lemma 2.5 ([42, Lemma 5.1]). *If \mathbb{H} dismantles to \mathbb{K} and $a \in H \setminus K$ is dominated in \mathbb{H} , then the substructure of \mathbb{H} induced by $H \setminus \{a\}$ dismantles to \mathbb{K} .*

Let $J \subseteq H$. We say that \mathbb{H} is **J -non-foldable** if every dominated element in \mathbb{H} belongs to J . When $J = \emptyset$, i.e., if no fold is possible at all, J -non-foldable relational structures are a generalization from the graph case of *stiff graphs*, used in [13]. There, it is proved that if we start from a given graph, one will always reach the same stiff graph up to isomorphism when dismantling greedily [13, Theorem 4.4]. It can be checked that in the context of relational structures the following is true: for every $J \subseteq H$, dismantling greedily elements in $H \setminus J$ will always lead to isomorphic relational structures (moreover, using an isomorphism that preserves J).

2.2. Walks in relational structures. We define a **walk**¹ w in a τ -structure \mathbb{H} to be a sequence

$$a_0, i_1, (R_1, \mathbf{a}_1), j_1, a_1, \dots, a_{n-1}, i_n, (R_n, \mathbf{a}_n), j_n, a_n$$

for some $n \geq 0$, such that, for all $1 \leq \ell \leq n$,

- $R_\ell \in \tau$, $\mathbf{a}_\ell \in R_\ell(\mathbb{H})$, $i_\ell \neq j_\ell$, and
- $a_{\ell-1} = \mathbf{a}_\ell[i_\ell]$ and $a_\ell = \mathbf{a}_\ell[j_\ell]$.

In this case, we will say that w **joins** a_0 (the **starting point**) and a_n (the **ending point**), and that the **length** of the walk w is n . Notice that if a walk w joins a_0 and a_n , then there is another walk w' that joins a_n and a_0 obtained by just reversing the order of indices. The **distance** $\text{dist}(a, b)$ between two elements $a, b \in H$ is defined to be the smallest length

¹See for example [41]. Our notion of walk can be alternatively seen as a walk in the incidence multigraph [42] of \mathbb{H} .

among all the walks w that join a and b . The distance $\text{dist}(V, W)$ between sets $V, W \subseteq H$ is defined to be the minimum distance between an element from V and an element from W . If V is a singleton $\{a\}$, we will just write $\text{dist}(a, W)$ instead of $\text{dist}(\{a\}, W)$.

Note that the definition of walk above coincides with the standard definition of walk when \mathbb{H} is a graph. However, in the case of graphs it will be convenient to describe the walk merely as the list a_0, \dots, a_n of its nodes, as usual.

A τ -structure \mathbb{H} is **connected** if there is a walk joining any pair of elements of its universe H and a **connected component** is any induced substructure that is connected and maximal in the sense of inclusion. A walk w is a **cycle** if $n > 0$, the starting and ending nodes are the only repeated nodes, and for all $1 \leq \ell < \ell' \leq n$, we have that $(R_\ell, \mathbf{a}_\ell) \neq (R_{\ell'}, \mathbf{a}_{\ell'})$. A τ -structure \mathbb{T} is a **τ -forest** if it has no cycles.

If, additionally, it is connected then it is a **τ -tree**. Usually, τ -trees are defined using the notion of incidence multigraph (e.g., see [42]). It is easy to verify that the definition given here is equivalent.

Remark 2.6. *We note here that some subtleties appear when encoding undirected graphs using relational structures (where tuples are ordered). It is usual (for example as in the proof of Lemma 3.3) to encode the edge-set of an undirected graph as a binary symmetric relation since many properties (such as the existence of homomorphisms) are preserved under this encoding. However, some of the connectedness notions (e.g., cycle, tree, etc.) introduced in this section do not correspond to the standard graph-theoretical notions when graphs are encoded as relational structures in this way. Instead, if one encodes the edge set of a graph with a binary relation containing one tuple per edge (ordering the nodes in any arbitrary way) then the concepts introduced in this section correspond, over graphs, to the usual graph-theoretic meaning of the term.*

2.3. Forest of walks. Given a τ -structure \mathbb{H} , we proceed to define a new τ -structure $\mathbb{T}_{\mathbb{H}}$. The universe $T_{\mathbb{H}}$ of $\mathbb{T}_{\mathbb{H}}$ consists of all the walks w in \mathbb{H} . For a k -ary $R \in \tau$, we define $R(\mathbb{T}_{\mathbb{H}})$ as follows: for all $\mathbf{a} = (a_1, \dots, a_k) \in R(\mathbb{H})$, for all $1 \leq i \leq k$, and for all walks w ending in a_i , we include in $R(\mathbb{T}_{\mathbb{H}})$ the tuple $(w_1, \dots, w_{i-1}, w, w_{i+1}, \dots, w_k)$, where w_j , $j \neq i$, is the walk obtained from w by extending it with i , (R, \mathbf{a}) , j , a_j .

We note that $\mathbb{T}_{\mathbb{H}}$ does not contain cycles and has exactly $|H|$ connected components, i.e., $|H|$ τ -trees in correspondence with the possible starting element in H for a walk in $T_{\mathbb{H}}$. It is easy to check that for every substructure \mathbb{I} of \mathbb{H} , the τ -structure $\mathbb{T}_{\mathbb{I}}$ is a substructure of $\mathbb{T}_{\mathbb{H}}$.

Remark 2.7. If \mathbb{H} is connected and we consider a slight modification of this previous definition, where the walks are asked to be non-backtracking [41] (i.e., for every $1 \leq \ell < n$, we have that either $i_\ell \neq j_{\ell+1}$, or $j_\ell \neq i_{\ell+1}$, or $(R_\ell, \mathbf{a}_\ell) \neq (R_{\ell+1}, \mathbf{a}_{\ell+1})$) then we obtain that each connected component of the resulting τ -structure corresponds to the *universal covering tree* of \mathbb{H} [39, 41]. In particular, they are all the same up to isomorphism.

Note that, by construction, the map $\rho_{\mathbb{H}} : T_{\mathbb{H}} \rightarrow H$ that sends every walk w in $T_{\mathbb{H}}$ to its ending point, that from now on we refer as the **label map**, defines a homomorphism from $\mathbb{T}_{\mathbb{H}}$ to \mathbb{H} . Furthermore,

Lemma 2.8. *Assume that \mathbb{H} is J -non-foldable for some $J \subseteq H$ and let U be a cofinite subset of $T_{\mathbb{H}}$ containing $\rho_{\mathbb{H}}^{-1}(J)$. Then, every homomorphism in $\text{Hom}(\mathbb{T}_{\mathbb{H}}, \mathbb{H})$ that agrees with $\rho_{\mathbb{H}}$ in U is identical to $\rho_{\mathbb{H}}$.*

Proof. Given $n \geq 0$, let W_n be the set of walks of length at least n in \mathbb{H} (notice that $W_n \subseteq W_{n-1}$ and $W_0 = T_{\mathbb{H}}$). We shall show that any $\rho' \in \text{Hom}(\mathbb{T}_{\mathbb{H}}, \mathbb{H})$ that agrees with $\rho_{\mathbb{H}}$ in

$W_n \cup \rho_{\mathbb{H}}^{-1}(J)$ for arbitrary n also agrees with $\rho_{\mathbb{H}}$ in W_{n-1} . Let w be any walk of length $n-1$ and let a be its ending point. We first show that $\rho'(w)$ dominates a in \mathbb{H} . Indeed, let $R \in \tau$ and let $\mathbf{a} = (a_1, \dots, a_k) \in R(\mathbb{H})$, where a appears, say, in the i th coordinate. By construction, $R(\mathbb{T}_{\mathbb{H}})$ contains the tuple $\mathbf{w} = (w_1, \dots, w_{i-1}, w, w_{i+1}, \dots, w_k)$, where for every $j \neq i$, w_j is obtained by concatenating $i, (R, \mathbf{a}), j, a_j$ at the end of w . Since w_j has length n for every $j \neq i$, it follows by assumption that $\rho'(w_j) = a_j$. That is, $\rho'(\mathbf{w})$ (which must be a tuple in $R(\mathbb{H})$) is obtained by replacing, in \mathbf{a} , a_i by $\rho'(w)$.

Hence, we have shown that $\rho'(w)$ dominates a in \mathbb{H} . Since \mathbb{H} is J -non-foldable it follows that either $\rho'(w) = a$ (and, hence, $\rho'(w) = \rho_{\mathbb{H}}(w)$) or $a \in J$ (and, hence, $\rho'(w) = \rho_{\mathbb{H}}(w)$ since $w \in \rho_{\mathbb{H}}^{-1}(J)$). To conclude the proof it is only necessary to observe that, since U is a cofinite set containing $\rho_{\mathbb{H}}^{-1}(J)$, it follows that any homomorphism that agrees with $\rho_{\mathbb{H}}$ in U , agrees as well in $W_n \cup \rho_{\mathbb{H}}^{-1}(J)$ for sufficiently large n . \square

2.4. Graphs of homomorphisms. Let \mathbb{G} and \mathbb{H} be τ -structures and suppose that \mathbb{H} is finite. We define two different kinds of graphs with vertex set $\text{Hom}(\mathbb{G}, \mathbb{H})$.

The first notion has been heavily studied, from an algorithmic perspective, in the context of the so-called *CSP reconfiguration problem* (see [32] and the references therein) and also from a structural point of view in the special case when \mathbb{G} and \mathbb{H} are graphs [13, 55]. We define $C(\mathbb{G}, \mathbb{H})$ as the (reflexive) graph with vertex set $\text{Hom}(\mathbb{G}, \mathbb{H})$ such that for every $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$, ϕ and ψ are adjacent if and only if ϕ and ψ differ in at most one value, i.e., there exists at most one $x \in G$ such that $\phi(x) \neq \psi(x)$. More generally, for any $n \geq 1$ we can define $C_n(\mathbb{G}, \mathbb{H})$ on $\text{Hom}(\mathbb{G}, \mathbb{H})$ by declaring ϕ and ψ adjacent if they differ in at most n values (in particular, $C(\mathbb{G}, \mathbb{H}) = C_1(\mathbb{G}, \mathbb{H})$).

A second notion of graph of homomorphisms appears in [42] and uses the notion of *links*. The **1-link** \mathbb{L} (with signature τ) is the τ -structure with universe $\{0, 1\}$, where $R(\mathbb{L}) = \{0, 1\}^k$ for every k -ary $R \in \tau$. Define a (di)graph $L(\mathbb{G}, \mathbb{H})$ with vertex set $\text{Hom}(\mathbb{G}, \mathbb{H})$ as follows: set $\phi \rightarrow \psi$ —i.e., a directed edge starting from ϕ and ending in ψ —if for any k -ary $R \in \tau$ and any $(x_1, \dots, x_k) \in R(\mathbb{G})$, we have that $(\gamma_1(x_1), \dots, \gamma_k(x_k)) \in R(\mathbb{H})$ whenever $\gamma_1, \dots, \gamma_k \in \{\phi, \psi\}$. Alternatively, one can say that ϕ and ψ are joined by a directed edge if there exists a homomorphism from \mathbb{L} to $\text{Hom}(\mathbb{G}, \mathbb{H})$ (see [42, Section 5.2]), mapping 0 to ϕ and 1 to ψ . Notice that the symmetry in the definition of 1-link implies that $L(\mathbb{G}, \mathbb{H})$ is, in fact, an undirected graph.

We say that a graph of homomorphisms (i.e., $C(\mathbb{G}, \mathbb{H})$, $C_n(\mathbb{G}, \mathbb{H})$, $L(\mathbb{G}, \mathbb{H})$) is *connected* if for every pair of maps $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$ agreeing on all but finitely many elements there exists a walk that joins them. Notice that if \mathbb{G} is finite, this coincides with the usual notion of connectedness in graph theory.

Clearly, $C_n(\mathbb{G}, \mathbb{H})$ is a subgraph of $C_{n+1}(\mathbb{G}, \mathbb{H})$. In contrast, $C_n(\mathbb{G}, \mathbb{H})$ and $L(\mathbb{G}, \mathbb{H})$ are not included in one another in general. However, we will establish (see Lemma 4.2) a meaningful relationship between both of them, by characterizing the connectivity properties of one in terms of the other.

Note that there is a one-to-one correspondence between the elements in $\text{Hom}(\mathbb{L} \times \mathbb{G}, \mathbb{H})$ and the edges of $L(\mathbb{G}, \mathbb{H})$. More generally, for $\ell \geq 1$ we define the ℓ -**link** \mathbb{L}_ℓ (with signature τ) as the τ -structure with universe $\{0, 1, \dots, \ell\}$, where $R(\mathbb{L}_\ell) = \cup_{i=0}^{\ell-1} \{i, i+1\}^k$, for every k -ary $R \in \tau$. In other words, the ℓ -link is a sequence of 1-links with their endpoints identified. Then the following result is immediate:

Lemma 2.9. *For every map $\phi : \{0, 1, \dots, \ell\} \times G \rightarrow H$ and every $1 \leq i \leq \ell$, let $\phi(i) : G \rightarrow H$ be the map defined by $\phi(i)(x) \mapsto \phi(i, x)$ for $x \in G$. Then, $\phi \in \text{Hom}(\mathbb{L}_\ell \times \mathbb{G}, \mathbb{H})$ if and only if $\phi(0), \dots, \phi(\ell)$ is a walk in $L(\mathbb{G}, \mathbb{H})$.*

2.5. Mixing properties. Given τ -structures \mathbb{G} and \mathbb{H} , it is useful to study properties in $\text{Hom}(\mathbb{G}, \mathbb{H})$ that allow us to *glue together* partially defined homomorphisms. This kind of properties are usually referred in the literature as *mixing properties*.

We say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is (V, W) -**mixing** for $V, W \subseteq G$, if for every $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$, there exists a map $\gamma \in \text{Hom}(\mathbb{G}, \mathbb{H})$ that agrees with ϕ on V and agrees with ψ on W . Given $g \geq 0$, we say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is **strongly irreducible with gap g** if it is (V, W) -mixing for every V, W such that $\text{dist}(V, W) \geq g$. We say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is **strongly irreducible** if it is strongly irreducible with gap g for some g .

A strengthening of strong irreducibility is the following property, introduced in [9]. Given $g \geq 0$, we say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is **topologically strong spatial mixing (TSSM) with gap g** if for every $V, W, S \subseteq G$ such that $\text{dist}(V, W) \geq g$ and for all $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$ that agree on S , there exists $\gamma \in \text{Hom}(\mathbb{G}, \mathbb{H})$ that agrees with ϕ on $V \cup S$ and agrees with ψ on $S \cup W$. We say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is **topologically strong spatial mixing** if it is TSSM with gap g for some g .

Clearly, $\text{Hom}(\mathbb{G}, \mathbb{H})$ is TSSM only if $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly irreducible but not vice versa (see [9, 10] for some counterexamples).

An antithesis of having good mixing properties is the existence of homomorphisms which are *frozen*. We say that $\phi \in \text{Hom}(\mathbb{G}, \mathbb{H})$ is a **frozen homomorphism** if for any cofinite set $U \subseteq G$, the only homomorphism $\psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$ such that $\psi|_U = \phi|_U$ is $\psi = \phi$ itself. Notice that Lemma 2.8, when $J = \emptyset$, says that $\rho_{\mathbb{H}}$ is a frozen homomorphism in $\text{Hom}(\mathbb{T}_{\mathbb{H}}, \mathbb{H})$.

3. MAIN THEOREMS

In this section we present the two main theorems of our work, which characterize in several ways a special class of relational structures. Both theorems consist of a generalization of some of the equivalences characterizing *dismantlable graphs* that appear in [13, Theorem 4.1]—which were developed only for the case of graphs—in two directions. First, Theorem 3.2 (or the *simple theorem*) extends [13] from graphs to arbitrary relational structures. Second, Theorem 3.6 (or the *refined theorem*), shows how the equivalences in Theorem 3.2 can be rephrased in terms of stronger properties with respect to special subsets of the universe of the given relational structure.

3.1. The case of graphs. The following theorem is a rephrasing of the equivalences that appear in [13, Theorem 4.1] which are relevant to us. We will use this as a prototypical example of the kind of results that we are aiming for, where we split the properties in 3 main categories: (A) dismantlability, (B) connectedness, and (C) mixing.

We suppose that graphs are allowed to have loops. Notice that a loopless graph cannot be dismantlable.

Theorem 3.1 ([13, Theorem 4.1]). *Let \mathbb{H} be a finite graph. The following are equivalent:*

- (A) \mathbb{H} is dismantlable;
- (B) $C(\mathbb{G}, \mathbb{H})$ is connected for every finite graph \mathbb{G} ; and
- (C) there exists $g \geq 0$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly irreducible with gap g for every graph \mathbb{G} .

3.2. First theorem: A parallel with the graph case. The following theorem shows that different dismantling, connectedness, and mixing notions are equivalent. It can be seen as a generalization of Theorem 3.1 to relational structures.

Theorem 3.2. *Let \mathbb{H} be a finite τ -structure with universe H . Then the following are equivalent:*

- (A1s) \mathbb{H} dismantles to a substructure \mathbb{I} such that \mathbb{I}^2 dismantles to its diagonal;
- (A2s) \mathbb{H}^2 dismantles to a substructure of its diagonal;
- (B1s) $C(\mathbb{G}, \mathbb{H})$ is connected for every locally finite τ -structure \mathbb{G} ;
- (B2s) there exists some $n \geq 1$ such that $C_n(\mathbb{G}, \mathbb{H})$ is connected for every finite τ -structure \mathbb{G} ;
- (B3s) $C(\mathbb{L} \times \mathbb{H}^2, \mathbb{H})$ is connected;
- (B4s) $L(\mathbb{G}, \mathbb{H})$ is connected for every finite τ -structure \mathbb{G} ;
- (B5s) the projections π_1 and π_2 are connected in $L(\mathbb{H}^2, \mathbb{H})$;
- (C1s) there exists $g \geq 0$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly irreducible with gap g for every τ -structure \mathbb{G} ; and
- (C2s) there exists $g \geq 0$ such that $\text{Hom}(T_{\mathbb{H}^2}, \mathbb{H})$ is $(\{x\}, W)$ -mixing for all $x \in T_{\mathbb{H}^2}$ and $W \subseteq T_{\mathbb{H}^2}$ with $\text{dist}(x, W) \geq g$.

As a consequence of our results we have the following lemma.

Lemma 3.3. *A graph \mathbb{H} is dismantlable if and only if \mathbb{H}^2 dismantles to a subset of its diagonal.*

Proof. In Theorem 3.2, we prove that, for a finite τ -structure \mathbb{H} , we have that \mathbb{H}^2 dismantles to a substructure of its diagonal if and only if there exists $g \geq 0$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly irreducible with gap g for all τ -structures \mathbb{G} . In particular, this applies if $\tau = \{E\}$, the usual binary relation of adjacency in graphs. Therefore, by Theorem 3.1, these two properties are also equivalent to \mathbb{H} being dismantlable, and we conclude. \square

In other words, thanks to Lemma 3.3, at least in the realm of graphs, we can freely replace “dismantlable” by “the square dismantles to a substructure of its diagonal”.

Still, it is not necessary to invoke Theorem 3.2 and Theorem 3.1 in order to prove Lemma 3.3, as this can be directly proved even in a more general setting.

A relation R of arity k is *symmetric*² if for every tuple $(t_1, \dots, t_k) \in R$, and every permutation $\sigma : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$, $(t_{\sigma(1)}, \dots, t_{\sigma(k)}) \in R$.

Lemma 3.4. *Assume that $\tau = \{R\}$ contains a unique relation symbol and let \mathbb{H} be a relational structure such that $R(\mathbb{H})$ is symmetric. Then \mathbb{H} is dismantlable if and only if \mathbb{H}^2 dismantles to a subset of its diagonal.*

Proof. It is immediate that if a is dominated in \mathbb{H} , then \mathbb{H}^2 dismantles to \mathbb{K}^2 , where \mathbb{K} is the substructure of \mathbb{H} induced by $H \setminus \{a\}$, which implies the “only if” part of the statement. For the converse, assume that \mathbb{H}^2 dismantles to a subset of its diagonal and let \mathbb{K} be a \emptyset -non-foldable structure obtained from \mathbb{H} by a sequence of folds until no further folding is possible. Then \mathbb{K}^2 dismantles to its diagonal as well (this fact is easy to see but to avoid repeating arguments given elsewhere we refer here to Remark 3.8). Assume, towards a contradiction, that the universe of \mathbb{K} is not a singleton. In consequence, the universe of \mathbb{K}^2 is not a singleton either and there must exist some element (a, b) dominated in \mathbb{K}^2 . Let (a', b') be the element in \mathbb{K}^2 dominating (a, b) and assume that $a' \neq a$ (the case $b' \neq b$ is analogous). We shall prove that a' dominates a in \mathbb{K} , contradicting the fact that \mathbb{K} is \emptyset -non-foldable. Indeed, let (a_1, \dots, a_k) be any tuple in $R(\mathbb{K})$ and let $i \in \{1, \dots, k\}$ with $a_i = a$. Note that b is not an isolated element (meaning it appears in some tuple of $R(\mathbb{K})$) since this would contradict the fact that \mathbb{K} is \emptyset -non-foldable. Note that $R(\mathbb{K})$ is symmetric as well,

²Not to be confused with the notion of symmetry introduced in Section 2.

implying that there exists a tuple $(b_1, \dots, b_k) \in R(\mathbb{K})$ with $b_i = b$. Since (a, b) is dominated by (a', b') in \mathbb{K}^2 , it follows that $R(\mathbb{K}^2)$ contains the tuple obtained after replacing (a_i, b_i) by (a', b') in $((a_1, b_1), \dots, (a_k, b_k))$. This implies that the tuple obtained replacing a_i by a' in (a_1, \dots, a_k) belongs to $R(\mathbb{K})$. \square

Question 3.5. While it is always true that if \mathbb{H} dismantles to a single element, then \mathbb{H}^2 dismantles to a substructure of the diagonal it is important to notice that the equivalence between “dismantlable” and “the square dismantles to a substructure of its diagonal” is not true for general relational structures. For example, given $\tau = \{R\}$ for R a binary relation symbol, we can take \mathbb{H} such that $H = \{0, 1\}$ and $R(\mathbb{H}) = \{(0, 1)\}$. Then, \mathbb{H} is not dismantlable, but \mathbb{H}^2 dismantles to its diagonal structures. It is an interesting open question to determine for which relational structures, besides those consisting of a unique symmetric relation, both notions coincide.

In Section 4 we shall prove Theorem 3.2. Indeed, we shall prove a refinement of it (Theorem 3.6). This refined theorem, which we believe is interesting on its own, is motivated by the fact that, sometimes, it is natural—particularly in the context of statistical physics—to work by forcing each particle of some subset to each take a particular spin and work with the remaining ones. For example, this is a common scenario when the particles in the boundary of a given set in a lattice are fixed to take particular spins and we want to study the distribution of spins in the interior of the set, conditioned on such boundary configuration. These ideas inspired the refined version, which can be regarded as the study of *boundary long range actions*, i.e., long range action phenomena where some boundary configuration is fixed, very similar to the concept of *boundary phase transition* in relation to phase transitions [47]. In order to state this stronger version, we need the following definitions.

3.3. Some refined definitions. Let \mathbb{H} and \mathbb{G} be τ -structures where \mathbb{G} is possibly infinite. Let ϕ_1, \dots, ϕ_t be a sequence of homomorphisms in $\text{Hom}(\mathbb{G}, \mathbb{H})$ and let $J \subseteq H$. We say that ϕ_1, \dots, ϕ_t is **J -preserving** if for every $x \in G$ such that $\phi_1(x) = \phi_t(x) = a \in J$, we have that $\phi_i(x) = a$ for every $1 \leq i \leq t$.

A **J -walk** is a J -preserving walk. Furthermore, we say that a graph of homomorphisms is **J -connected** if for every pair of maps $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$ agreeing on all but finitely many elements there exists a J -walk that joins them. Notice that if $J = \emptyset$, the definition of J -connectedness coincides with the definition of connectedness introduced in Section 2.4.

Given $J \subseteq H$, we say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is **(V, W) -mixing with respect to J** if $\text{Hom}(\mathbb{G}, \mathbb{H})$ is (V, W) -mixing and the map γ can be chosen so that $\gamma(x) = \phi(x) = \psi(x)$ for all $x \in (\phi, \psi)^{-1}(\Delta(J^2))$, where $(\phi, \psi)^{-1}$ denotes the inverse of the map $(\phi, \psi) : \mathbb{G} \rightarrow \mathbb{H}^2$ given by $x \mapsto (\phi(x), \psi(x))$; i.e., the map γ coincides with ϕ in V , with ψ in W , and with both of them for x such that $\phi(x) = \psi(x) \in J$. We say that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is **strongly J -irreducible** with gap g if it is (V, W) -mixing with respect to J for all V, W such that $\text{dist}(V, W) \geq g$.

It is easy to check that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly H -irreducible with gap g if and only if $\text{Hom}(\mathbb{G}, \mathbb{H})$ is TSSM with gap g . Indeed, the equivalence comes from the fact that the set S in the definition of TSSM is always a subset of $(\phi, \psi)^{-1}(\Delta(H^2))$ and the strongly H -irreducible property is equivalent to the TSSM case where $S = (\phi, \psi)^{-1}(\Delta(H^2))$.

3.4. Second theorem: A refinement.

Theorem 3.6. *Let \mathbb{H} be a finite τ -structure with universe H and let $J \subseteq H$. Then the following are equivalent:*

- (A1) \mathbb{H} dismantles to a substructure \mathbb{I} whose universe contains J and such that \mathbb{I}^2 dismantles to its diagonal;
- (A2) \mathbb{H}^2 dismantles to a substructure \mathbb{K} where its universe K satisfies $\Delta(J^2) \subseteq K \subseteq \Delta(H^2)$;
- (B1) $C(\mathbb{G}, \mathbb{H})$ is J -connected for every locally finite τ -structure \mathbb{G} ;
- (B2) there exists some $n \geq 1$ such that $C_n(\mathbb{G}, \mathbb{H})$ is J -connected for every finite τ -structure \mathbb{G} ;
- (B3) $C(\mathbb{L} \times \mathbb{H}^2, \mathbb{H})$ is J -connected;
- (B4) $L(\mathbb{G}, \mathbb{H})$ is J -connected for every finite τ -structure \mathbb{G} ;
- (B5) the projections π_1 and π_2 are J -connected in $L(\mathbb{H}^2, \mathbb{H})$;
- (C1) there exists $g \geq 0$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly J -irreducible with gap g for every τ -structure \mathbb{G} ; and
- (C2) there exists $g \geq 0$ such that $\text{Hom}(\mathbb{T}_{\mathbb{H}^2}, \mathbb{H})$ is $(\{x\}, W)$ -mixing with respect to J for all $x \in T_{\mathbb{H}^2}$ and $W \subseteq T_{\mathbb{H}^2}^2$ with $\text{dist}(x, W) \geq g$.

Remark 3.7. To our knowledge, if $J \neq \emptyset$, Theorem 3.6 has not been known before even in the graph case. In addition, notice that from Theorem 3.6 it follows, by taking $J = H$, that \mathbb{H}^2 dismantles to its full diagonal if and only if there exists $g \geq 0$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is TSSM with gap g for every τ -structure \mathbb{G} .

As a byproduct of our results (in particular of Lemma 4.3 in Section 4) we derive the following effective procedure to decide whether the statements of Theorem 3.6 are satisfied.

Remark 3.8. Let \mathbb{H} be a finite τ -structure and let $J \subseteq H$. If there is a dismantling sequence (for \mathbb{H} and J) as in Theorem 3.6(A1), then there is one that can be obtained in the following greedy manner:

- First step. Starting from \mathbb{H} , iteratively fold elements in $H \setminus J$ in any arbitrary way until no further dismantling is possible, obtaining a J -non-foldable τ -structure \mathbb{I} .
- Second step. Starting from \mathbb{I}^2 , iteratively fold symmetric pairs (a, b) and (b, a) in an arbitrary way until no further dismantling is possible.

4. PROOFS

The following implications are immediate:

$$(B1) \Rightarrow (B2), (B1) \Rightarrow (B3), (B4) \Rightarrow (B5), \text{ and } (C1) \Rightarrow (C2).$$

The rest of the section contains several lemmas from which the remaining implications follow according the following table:

Lemma 4.2	$(B2) \Rightarrow (B4)$ $(B3) \Rightarrow (B5)$	Lemma 4.5	$(A1) \Rightarrow (A2)$
		Lemma 4.7	$(A1) \Rightarrow (B1)$
Lemma 4.3	$(B5) \Rightarrow (A1)$	Lemma 4.8	$(A1) \Rightarrow (C1)$
Lemma 4.4	$(A2) \Rightarrow (A1)$	Lemma 4.9	$(C2) \Rightarrow (A1)$

Lemma 4.1. Let \mathbb{I} be a J -non-foldable τ -structure and \mathbb{K} be a symmetric τ -structure obtained by folding only non-diagonal elements of \mathbb{I}^2 (note that $J^2 \subseteq K$). Assume that \mathbb{K} is minimal, i.e., \mathbb{K} has no proper substructure with the same property. Then, \mathbb{K} is $\Delta(J^2)$ -non-foldable.

Proof. Let (a, b) be any dominated element in \mathbb{K} . We first shall prove that (a, b) is diagonal by contradiction. Let $(c, d) \neq (a, b)$ be an element dominating (a, b) in \mathbb{K} . Since \mathbb{K} is symmetric, it follows that both (b, a) and (d, c) are present in K as well. If $(c, d) \neq (b, a)$,

then we can fold both (a, b) and (b, a) in \mathbb{K} , contradicting the minimality of \mathbb{K} . Otherwise, it follows (as we shall prove straight away) that (b, b) dominates both (a, b) and (b, a) , contradicting again the minimality of \mathbb{K} .

We only need to prove that (b, b) dominates (a, b) , as the other case is analogous. Let $R \in \tau$, let $((a_1, b_1), \dots, (a_k, b_k))$ be any tuple in $R(\mathbb{K})$, and let j such that $(a_j, b_j) = (a, b)$. Since (b, a) dominates (a, b) ,

$$((a_1, b_1), \dots, (a_{j-1}, b_{j-1}), (b, a), (a_{j+1}, b_{j+1}), \dots, (a_k, b_k)) \in R(\mathbb{K}).$$

It follows that $(a_1, \dots, a_{j-1}, b, a_{j+1}, \dots, a_k) \in R(\mathbb{I})$. Since (b_1, \dots, b_k) is also a tuple in $R(\mathbb{I})$ and $b_j = b$, it follows that

$$((a_1, b_1), \dots, (a_{j-1}, b_{j-1}), (b, b), (a_{j+1}, b_{j+1}), \dots, (a_k, b_k)) \in R(\mathbb{K}).$$

To complete the proof, we shall show that a is dominated in \mathbb{I} which, by the assumptions on \mathbb{I} , implies that $a \in J$. Let (c, d) be any element dominating (a, a) in \mathbb{K} . Clearly, c or d is different from a , so assume, w.l.o.g., that $a \neq c$. We shall show that c dominates a in \mathbb{I} . Indeed, let $R \in \tau$, let (a_1, \dots, a_k) be any tuple in $R(\mathbb{I})$, and let j such that $a_j = a$. Then, $((a_1, a_1), \dots, (a_k, a_k)) \in R(\mathbb{K})$ and thus, $((a_1, a_1), \dots, (a_{j-1}, a_{j-1}), (c, d), (a_{j+1}, a_{j+1}), \dots, (a_k, a_k)) \in R(\mathbb{K})$ implies that $(a_1, \dots, a_{j-1}, c, a_{j+1}, \dots, a_k) \in R(\mathbb{I})$, concluding the proof. \square

Lemma 4.2 ([(B2) \Rightarrow (B4), (B3) \Rightarrow (B5)]). *Given a finite τ -structure \mathbb{G} , if $C_n(\mathbb{L}_n \times \mathbb{G}, \mathbb{H})$ is J -connected for some $n \geq 1$, then $L(\mathbb{G}, \mathbb{H})$ is J -connected.*

Proof. Let $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$. We shall show that there is a J -walk in $L(\mathbb{G}, \mathbb{H})$ from ϕ to ψ . By Lemma 2.9, the maps $h, h' : \{0, \dots, n\} \times G \rightarrow H$, where $h(i) = \phi$ and $h'(i) = \psi$ for every $i \in \{0, \dots, n\}$, belong trivially to $\text{Hom}(\mathbb{L}_n \times \mathbb{G}, \mathbb{H})$. Since we are assuming that $C_n(\mathbb{L}_n \times \mathbb{G}, \mathbb{H})$ is J -connected, it follows that there exists a J -walk $h = h^0, \dots, h^s = h'$ in $C_n(\mathbb{L}_n \times \mathbb{G}, \mathbb{H})$ joining them, for some $s \geq 0$. Proceeding by induction, we will construct a walk in $L(\mathbb{G}, \mathbb{H})$ connecting $h^0(0)$ with $h^j(0)$ for every $j \in \{0, \dots, s\}$. The base case, $j = 0$, is trivial. Now assume that the statement holds for some $j < s$. Since h^j and h^{j+1} differ in at most n values, this implies that $(h^j(0), \dots, h^j(n))$ and $(h^{j+1}(0), \dots, h^{j+1}(n))$ must have an entry in common. Hence, let $i^* \in \{0, \dots, n\}$ be such that $h^j(i^*) = h^{j+1}(i^*)$. Since, by Lemma 2.9, there are walks from $h^j(0)$ to $h^j(i^*)$ and from $h^{j+1}(i^*)$ to $h^{j+1}(0)$, we are done.

Hence, we have shown that there is a walk in $L(\mathbb{G}, \mathbb{H})$ connecting $h^0(0) = \phi$ and $h^s(0) = \psi$. It remains to show that the walk we have just constructed is J -preserving. Let $h^j(i)$, for $j \in \{0, \dots, s\}$ and $i \in \{0, \dots, n\}$, be any element in the walk, and let $x \in G$ such that $\phi(x) = \psi(x) = a \in J$. Since $h^0(i)(x) = \phi(x) = a$, $h^s(i)(x) = \psi(x) = a$, and the walk h^0, \dots, h^s is J -preserving, it follows that $h^j(i) = a$ as well. \square

Lemma 4.3 ([(B5) \Rightarrow (A1)]). *Suppose that the projections π_1 and π_2 are J -connected in $L(\mathbb{H}^2, \mathbb{H})$. Let \mathbb{I} be any J -non-foldable structure obtained by dismantling \mathbb{H} and let \mathbb{K} be the symmetric $\Delta(J^2)$ -non-foldable structure given by Lemma 4.1. Then \mathbb{K} contains only diagonal elements.*

Proof. Let \mathbb{I} and \mathbb{K} be as in the statement. It suffices to show that \mathbb{K} contains only diagonal elements. Let r be the natural retraction from \mathbb{H}^2 to \mathbb{K} obtained from a successive composition of folds. Let $\pi_1 = h_1, \dots, h_t = \pi_2$ be a J -walk in $L(\mathbb{H}^2, \mathbb{H})$ connecting π_1 and π_2 and consider the family of maps $\phi_1, \dots, \phi_{2t-1} : K \rightarrow K$, where

$$\phi_i = \begin{cases} r(h_i, \pi_2) & \text{if } i \leq t, \\ r(\pi_2, h_{2t-i}) & \text{if } i > t. \end{cases}$$

It follows directly from the construction that every ϕ_i is an endomorphism of \mathbb{K} (i.e., a homomorphism from \mathbb{K} to \mathbb{K}). If \mathbb{K} contains some non-diagonal element (c, d) , then it must also contain (d, c) . Hence, $\phi_1 \neq \phi_{2t-1}$, since $\phi_1(c, d) = r(c, d) = (c, d)$ and $\phi_{2t-1}(c, d) = r(d, c) = (d, c)$. It follows that there exists some $i \leq t$ such that $\phi_1 \neq \phi_i$ or there exists some $i \geq t$ such that $\phi_i \neq \phi_{2t-1}$. We shall consider only the first case since the proof for the other case is symmetric. Let i be the minimum such that $\phi_1 \neq \phi_i$. Also, let $(a, b) \in K$ with the property that $\phi_i(a, b) \neq \phi_1(a, b) = \phi_{i-1}(a, b)$. We shall prove that $\phi_i(a, b)$ dominates (a, b) in \mathbb{K} . Indeed, let $R \in \tau$ of arity k , let $((a_1, b_1), \dots, (a_k, b_k))$ be any tuple in $R(\mathbb{K})$, and let $j \in \{1, \dots, k\}$ be such that $(a_j, b_j) = (a, b)$. Consider now the k -tuple (c_1, \dots, c_k) , where $c_\ell = h_i(a_\ell, b_\ell) = h_i(a, b)$ for $j = \ell$ and $c_\ell = h_{i-1}(a_\ell, b_\ell)$ for $\ell \neq j$. Since h_{i-1} and h_i are adjacent in $L(\mathbb{H}^2, \mathbb{H})$, it follows that $(c_1, \dots, c_k) \in R(\mathbb{H})$. Hence, the tuple (d_1, \dots, d_k) with $d_\ell = r(c_\ell, b_\ell)$ belongs to $R(\mathbb{K})$. Notice that, by definition, $d_j = r(h_i(a, b), b) = \phi_i(a, b)$ and $d_\ell = r(h_{i-1}(a_\ell, b_\ell), b_\ell) = \phi_{i-1}(a_\ell, b_\ell) = (a_\ell, b_\ell)$ for $\ell \neq j$.

We have just shown that $\phi_i(a, b)$ dominates (a, b) in \mathbb{K} . We know that \mathbb{K} is $\Delta(J^2)$ -non-foldable which does not lead yet to contradiction as it could be the case that $a = b$ and $a \in J$. In this case, note that $h_1(a, b) = h_t(a, b) = a$. Since $a \in J$ and h_1, \dots, h_t is J -preserving, it follows that $h_j(a, b) = a$ for every $1 \leq j \leq t$. This contradicts the fact that $\phi_i(a, b) \neq \phi_1(a, b)$. \square

Note that Remark 3.8 is a direct consequence of Lemma 4.3.

Lemma 4.4 ([(A2) \Rightarrow (A1)]). *If \mathbb{H} dismantles to a substructure \mathbb{I} whose universe contains J and such that \mathbb{I}^2 dismantles to its diagonal, then \mathbb{H}^2 dismantles to a substructure \mathbb{K} where its universe K satisfies $\Delta(J^2) \subseteq K \subseteq \Delta(H^2)$.*

Proof. Assume that \mathbb{H}^2 dismantles to a substructure \mathbb{K} where its universe K satisfies $\Delta(J^2) \subseteq K \subseteq \Delta(H^2)$. We can assume that $K \neq \Delta(H^2)$ since otherwise there is nothing to prove. Let $\mathbb{H}^2 = \mathbb{J}_0, \dots, \mathbb{J}_\ell = \mathbb{K}$ the dismantling sequence and \mathbb{J}_i the last structure in the sequence whose domain contains $\Delta(H^2)$. Consequently, \mathbb{J}_i contains a diagonal element (a, a) that is dominated by some other element (b, c) . Assume that $b \neq a$ (the other case is analogous). We claim that b dominates a in \mathbb{H} . Indeed, let $R \in \tau$, let (a_1, \dots, a_k) be any tuple in $R(\mathbb{H})$ and let j such that $a_j = a$. Note that $R(\mathbb{J}_i)$ contains tuple $((a_1, a_1), \dots, (a_k, a_k))$ and, consequently, it also contains $((a_1, a_1), \dots, (a_{j-1}, a_{j-1}), (b, c), (a_{j+1}, a_{j+1}), \dots, (a_k, a_k))$. It follows that $R(\mathbb{H})$ contains $(a_1, \dots, a_{j-1}, b, a_{j+1}, \dots, a_k)$ and we are done. Then, every element of the form (x, a) or (a, x) is dominated by (x, b) or (b, x) in \mathbb{H}^2 , respectively, and hence, \mathbb{H}^2 dismantles to \mathbb{I}^2 where \mathbb{I} is the structure obtained by dismantling a in \mathbb{H} . It follows by Lemma 2.5 that \mathbb{I}^2 dismantles to \mathbb{K} . Iterating this argument we obtain statement (A1). \square

Lemma 4.5 ([(A1) \Rightarrow (A2)]). *Assume that \mathbb{H} dismantles to a substructure \mathbb{I} whose universe contains J and such that \mathbb{I}^2 dismantles to its diagonal. Then, there is a dismantling sequence $\mathbb{H}^2 = \mathbb{J}_0, \dots, \mathbb{J}_\ell = \Delta(\mathbb{I}^2)$ such that, for every $0 < i \leq \ell$, if u_i is the element folded in \mathbb{J}_{i-1} to obtain \mathbb{J}_i and u_i is a diagonal element, then u_i is dominated by a (different) diagonal element in \mathbb{J}_{i-1} .*

Proof. Let a_1, \dots, a_r be a sequence of elements to be folded to obtain \mathbb{I} from \mathbb{H} . We construct a sequence of elements to be folded to obtain $\Delta(\mathbb{I}^2)$ from \mathbb{H}^2 as follows: In a first stage, if a_1 was dominated by some element b_1 in \mathbb{H} , we fold all elements of the form (x, a_1) , (a_1, x) , or (a_1, a_1) to (x, b_1) , (b_1, x) , or (b_1, b_1) , respectively. In a second stage, one folds (again, in an arbitrary order) all elements of the form (x, a_2) , (a_2, x) , and (a_2, a_2) that are still left, and continues in the same manner until one obtains \mathbb{I}^2 . At this point, one

proceeds dismantling \mathbb{I}^2 to its diagonal as originally was done. It is easy to see that the sequence finally obtained satisfies the desired properties. \square

We will require the following definition. Let $\mathbb{H}^2 = \mathbb{J}_0, \dots, \mathbb{J}_\ell = \Delta(\mathbb{I}^2)$ be the sequence provided by Lemma 4.5. For every $0 < i \leq \ell$, let s_i be the fold of \mathbb{J}_{i-1} into \mathbb{J}_i defined in the natural way (i.e., s_i acts as the identity on J_i and maps u_i to any element that dominates it in \mathbb{J}_{i-1}) and define, for every $0 < i \leq \ell$,

$$(1) \quad r_i := s_i \circ \dots \circ s_1$$

and r_0 to be the identity. Note that r_i defines a retraction of \mathbb{H}^2 into \mathbb{J}_i and that, again by Lemma 4.5, we can assume that r_i maps every diagonal element into a diagonal element. Considering this, we have the following additional lemma, which was inspired by [13, Lemma 5.2].

Lemma 4.6. *Under the assumptions of Lemma 4.5, let \mathbb{G} be a (possibly infinite) τ -structure, $\phi \in \text{Hom}(\mathbb{G}, \mathbb{H}^2)$, $X \subseteq G$, and let $\omega : \mathbb{G} \rightarrow \mathbb{H}^2$ be the map defined as*

$$\omega(x) = \begin{cases} (r_{\ell - \text{dist}(x, X)} \circ \phi)(x), & \text{if } \text{dist}(x, X) \leq \ell, \\ \phi(x), & \text{otherwise,} \end{cases}$$

where r_i is as in Formula (1) and $\text{dist}(x, X)$ is the distance from x to X in \mathbb{G} . Then, ω is a homomorphism from \mathbb{G} to \mathbb{H}^2 . Furthermore, if X is finite and \mathbb{G} is locally finite, then ϕ and ω are connected in $C(\mathbb{G}, \mathbb{H}^2)$ by a $\Delta(J^2)$ -preserving walk.

Proof. Given $Y \subseteq G$, for every $i = 0, \dots, \ell$, let $\omega_{i,Y}$ be the map defined as

$$\omega_{i,Y}(x) = \begin{cases} (r_{i - \text{dist}(x, Y)} \circ \phi)(x), & \text{if } \text{dist}(x, Y) \leq i, \\ \phi(x), & \text{otherwise.} \end{cases}$$

Note that $\omega = \omega_{\ell, X}$. By induction, we shall prove that $\omega_{i,X}$ is a homomorphism from \mathbb{G} to \mathbb{H}^2 and, if X is finite and \mathbb{G} is locally finite, then ϕ and $\omega_{i,X}$ are connected in $C(\mathbb{G}, \mathbb{H}^2)$ by a J_i -preserving walk.

The base case, $i = 0$, is trivial. For the inductive case ($i - 1 \Rightarrow i$), we can assume that the map $\omega' = \omega_{i-1, X'}$ is a homomorphism in $\text{Hom}(\mathbb{G}, \mathbb{H}^2)$, where X' is the set of all elements of G at distance at most 1 from X .

Claim. For every $Z \subseteq X$, the map Ψ_Z that acts as $s_i \circ \omega'$ in Z and as ω' elsewhere, is a homomorphism from \mathbb{G} to \mathbb{H}^2 .

Indeed, let R be any relation symbol in τ and let $\mathbf{x} = (x_1, \dots, x_k)$ be any tuple in $R(\mathbb{G})$. By inductive hypothesis, $\omega'(\mathbf{x})$ belongs to $R(\mathbb{H}^2)$. If $\{x_1, \dots, x_k\} \cap Z = \emptyset$, then $\Psi_Z(\mathbf{x}) = \omega'(\mathbf{x})$ and nothing needs to be done. So, assume that \mathbf{x} contains some element from Z . Since $Z \subseteq X$, it follows that $\{x_1, \dots, x_k\} \subseteq X'$. Consequently, $\omega'(\mathbf{x}) = (r_{i-1} \circ \phi)(\mathbf{x})$, which is a tuple of $R(\mathbb{J}_{i-1})$. Let u_i be the element that is folded in \mathbb{J}_{i-1} to obtain \mathbb{J}_i and note that $\Psi_Z(\mathbf{x})$ is obtained by replacing, in $\omega'(\mathbf{x})$, some (possibly zero) occurrences of u_i by $s_i(u_i)$. Since $s_i(u_i)$ dominates u_i in \mathbb{J}_{i-1} , it follows that $\Psi_Z(\mathbf{x})$ belongs to $R(\mathbb{J}_{i-1})$ (and hence to $R(\mathbb{H}^2)$), finishing the proof of the claim.

It follows that $\omega = \Psi_X \in \text{Hom}(\mathbb{G}, \mathbb{H}^2)$. Furthermore, assume that $X = \{x_1, \dots, x_n\}$ is finite and \mathbb{G} is locally finite. Then

$$\Psi_\emptyset, \Psi_{\{x_1\}}, \Psi_{\{x_1, x_2\}}, \dots, \Psi_{\{x_1, x_2, \dots, x_n\}}$$

is a walk in $C(\mathbb{G}, \mathbb{H}^2)$ connecting $\omega' = \Psi_\emptyset$ and $\omega = \Psi_X$.

Since \mathbb{G} is locally finite, it follows that X' is finite. By inductive hypothesis, there is a J_{i-1} -preserving walk joining ϕ and ω' in $C(\mathbb{G}, \mathbb{H}^2)$. Hence, by concatenating the two walks, it follows that ϕ and ω are also connected in $C(\mathbb{G}, \mathbb{H}^2)$.

It remains to see that the walk thus constructed is J_i -preserving. Let $x \in G$ such that $\phi(x) = \omega(x) = (a, b) \in J_i$. Since $(a, b) \in J_i$ it follows from the definition of ω' that $\omega'(x) = \phi(x) = (a, b)$. Hence it is only necessary to observe that the walk joining ϕ and ω' is J_i -preserving (by inductive hypothesis and $J_i \subseteq J_{i-1}$) and that the walk joining ω' and ω is also J_i -preserving (directly from its construction). \square

Lemma 4.7 ([(A1) \Rightarrow (B1)]). *If \mathbb{H} dismantles to a substructure \mathbb{I} whose universe contains J and such that \mathbb{I}^2 dismantles to its diagonal, then $C(\mathbb{G}, \mathbb{H})$ is J -connected for every locally finite τ -structure \mathbb{G} .*

Proof. Assume that statement (A1) holds. Let r_0, \dots, r_ℓ be the maps provided by Lemma 4.5 (see Equation (1)), let \mathbb{G} be any locally finite τ -structure, let $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$, and let D be the set of elements in which ϕ and ψ disagree, which we can assume is finite. Then, the map $\Phi : G \rightarrow \mathbb{H}^2$ with $x \mapsto \Phi(x) = (\phi(x), \psi(x))$ for $x \in G$, defines a homomorphism from \mathbb{G} to \mathbb{H}^2 . Since D is finite, it follows from Lemma 4.6 that there is $\Delta(J^2)$ -preserving walk Ψ_1, \dots, Ψ_t in $C(\mathbb{G}, \mathbb{H}^2)$ that joins Φ and the map $\omega : G \rightarrow \mathbb{H}^2$ sending $x \in G$ to $(r_{\ell-j} \circ \Phi)(x)$, where j is the minimum between ℓ and the distance from x to D . It is not difficult to see that $\omega(x)$ is a diagonal element for every $x \in G$. Indeed, if $x \notin D$, it follows from the fact (seen just right after Lemma 4.5) that every retraction $r_{\ell-j}$ maps diagonal elements into diagonal elements and, if $x \in D$, it follows from the fact that $r_\ell(\mathbb{H}^2) = \Delta(\mathbb{I}^2)$. Hence, $\pi_1 \circ \Psi_t = \pi_2 \circ \Psi_t$. Consequently, $\pi_1 \circ \Psi_1, \pi_1 \circ \Psi_2, \dots, \pi_1 \circ \Psi_t = \pi_2 \circ \Psi_t, \pi_2 \circ \Psi_{t-1}, \dots, \pi_2 \circ \Psi_1$ defines a walk that joins $\pi_1 \circ \Psi_1 = \phi$ and $\pi_2 \circ \Psi_1 = \psi$. It remains to see that the walk thus defined is J -preserving. To do so, we just have to notice that if x is an element in \mathbb{G} such that $\phi(x) = \psi(x) \in J$, then we have by construction that $\Psi_1(x) = (x, x)$ and $\Psi_t(x) = (x, x)$, and then use the fact that Ψ_1, \dots, Ψ_t is $\Delta(J^2)$ -preserving. \square

Lemma 4.8 ([(A1) \Rightarrow (C1)]). *If \mathbb{H} dismantles to a substructure \mathbb{I} whose universe contains J and such that \mathbb{I}^2 dismantles to its diagonal, then there exists $g \geq 0$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is strongly J -irreducible with gap g for every τ -structure \mathbb{G} .*

Proof. Assume that statement (A1) holds. Let r_0, \dots, r_ℓ be the maps provided by Lemma 4.5 (see Equation (1)) and define g to be 2ℓ . Let \mathbb{G} be a τ -structure, $V, W \subseteq G$ with $\text{dist}(V, W) \geq g$, and $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$. Then, the map $x \mapsto (\phi(x), \psi(x))$ is a homomorphism from \mathbb{G} to \mathbb{H}^2 . Let ω be the homomorphism from \mathbb{G} to \mathbb{H}^2 provided by Lemma 4.6 with map $x \mapsto (\phi(x), \psi(x))$ and where $X \subseteq G$ is the set of all elements at distance at least ℓ from $V \cup W$. It follows by construction that $\omega(x) = (\phi(x), \psi(x))$ for every $x \in V \cup W$ and that $\omega(X) \subseteq \Delta(\mathbb{I}^2)$.

Define $h(x)$ to be $\pi_1(\omega(x))$ if x is at distance less than ℓ to V and $\pi_2(\omega(x))$ otherwise. We shall prove that h defines a homomorphism from \mathbb{G} to \mathbb{H} . Let R be any relation symbol in τ of arity k , and let $\mathbf{x} = (x_1, \dots, x_k)$ be any tuple in $R(\mathbb{G})$. The fact that h is a homomorphism follows directly from the following claim: $h(\mathbf{x}) = \pi_1(\omega(\mathbf{x}))$ if the minimum distance j from any element in the tuple to V is at most $\ell - 1$ and $h(\mathbf{x}) = \pi_2(\omega(\mathbf{x}))$ otherwise. To prove the claim, notice that, since all the elements in \mathbf{x} are at distance j or $j + 1$ from V , we only need to consider the case when $j = \ell - 1$. Let x_i be any element in \mathbf{x} . If the distance of x_i to V is $\ell - 1$, then $h(x_i) = \pi_1(\omega(x_i))$ by definition. Assume, otherwise, that the distance of x_i to V is ℓ . Since the distance from V to W is at least $g = 2\ell$, then x_i is at distance at least ℓ from W , and hence $x_i \in X$. Hence, $\omega(x_i)$ is a diagonal element, and hence $\pi_2(\omega(x_i)) = \pi_1(\omega(x_i))$.

Furthermore, since for every $x \in V \cup W$, $\omega(x) = (\phi(x), \psi(x))$, it follows that h agrees with ϕ on V and ψ on W .

Now, let $x \in G$ be any element such that $\phi(x) = \psi(x) \in J$. Note that, by construction, ω is $\Delta(I^2)$ -preserving. Since $J \subseteq I$, it follows that $\omega(x) = (\phi(x), \psi(x))$. Since $h(x)$ is either the first or second projection of $\omega(x)$, it follows that $h(x) = \phi(x) = \psi(x)$. \square

Recall that the definition of the label map $\rho_{\mathbb{H}} : T_{\mathbb{H}} \rightarrow H$ that sends every walk w in $T_{\mathbb{H}}$ to its ending point (see Section 2.3).

Lemma 4.9 [(C2) \Rightarrow (A1)]. *If there exists $g \geq 0$ such that $\text{Hom}(T_{\mathbb{H}^2}, \mathbb{H})$ is $(\{x\}, W)$ -mixing with respect to J with gap g for all $x \in T_{\mathbb{H}^2}$ and $W \subseteq T_{\mathbb{H}^2}^2$, then \mathbb{H} dismantles to a substructure \mathbb{I} whose universe contains J and such that \mathbb{I}^2 dismantles to its diagonal.*

Proof. For the sake of contradiction, let's suppose that statement (C2) holds, but statement (A1) does not. Let \mathbb{I} be any J -non-foldable relational structure obtained by dismantling \mathbb{H} , let \mathbb{K} be the symmetric $\Delta(J^2)$ -non-foldable relational structure given by Lemma 4.1, let r be a retraction of \mathbb{H}^2 onto \mathbb{K} defined in the natural way, and let g be any value given by statement (C2).

Since statement (A1) does not hold, it follows that K contains a non-diagonal element $a = (a_1, a_2)$. Let w_0 be the (unique) walk of length 0 starting at a , let $V = \{w_0\}$, and let W_g be the set containing all walks in $T_{\mathbb{H}^2}$ of length at least g . Let ϕ and ψ be the homomorphisms from $T_{\mathbb{H}^2}$ to \mathbb{H} defined as $\phi := \pi_1 \circ \rho_{\mathbb{H}^2}$ and $\psi := \pi_2 \circ \rho_{\mathbb{H}^2}$. From statement (C2), it follows that there exist homomorphisms $h_1, h_2 : T_{\mathbb{H}^2} \rightarrow \mathbb{H}$ such that h_1 agrees with ψ on V and with ϕ on W_g , h_2 agrees with ϕ on V and with ψ on W_g , and $h_1(w) = h_2(w) = \phi(w) = \psi(w)$ for all $w \in \gamma^{-1}(\Delta(J^2))$.

Since $T_{\mathbb{K}}$ is a substructure of $T_{\mathbb{H}^2}$, it follows that the mapping $\rho' : T_{\mathbb{K}} \rightarrow K$ defined as $\rho'(w) = r(h_1(w), h_2(w))$ is a homomorphism from $T_{\mathbb{K}}$ to \mathbb{K} . By construction, ρ' agrees with $\rho_{\mathbb{K}}$ in W_g and every $w \in T_{\mathbb{K}}$ with $\rho_{\mathbb{K}}(w) \in \Delta(J^2)$. Hence, by Lemma 2.8, ρ' must be identical to $\rho_{\mathbb{K}}$, implying, in particular, that $\rho_{\mathbb{K}}$ and ρ' agree in w_0 . However,

$$\rho'(w_0) = r(h_1(w_0), h_2(w_0)) = r(a_2, a_1) = (a_2, a_1),$$

where the last equality follows from the fact that \mathbb{K} is symmetric. We obtain a contradiction, since $\rho_{\mathbb{K}}(w_0) = a = (a_1, a_2)$ and a is a non-diagonal element. \square

5. GIBBS MEASURES AND APPLICATIONS

5.1. Basic definitions. Given a finite τ -structure \mathbb{H} with universe H , a **weight function** for \mathbb{H} is a map $\lambda : H \rightarrow \mathbb{R}^+$.

Let \mathbb{G} be a locally finite τ -structure. If $V \subseteq G$ is a finite set and $\phi \in \text{Hom}(\mathbb{G}, \mathbb{H})$, we define $\mathbb{P}_{V, \phi}$ to be the probability measure on $\text{Hom}(\mathbb{G}, \mathbb{H})$ given by

$$\mathbb{P}_{V, \phi}(\{\psi\}) := \begin{cases} Z_{V, \phi}(\lambda)^{-1} \prod_{x \in V} \lambda(\psi(x)) & \text{if } \psi|_V \cup \phi|_{G \setminus V} \in \text{Hom}(\mathbb{G}, \mathbb{H}), \\ 0 & \text{otherwise,} \end{cases}$$

for $\psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$, where $\psi|_V \cup \phi|_{G \setminus V}$ is the map that coincides with ψ in V and with ϕ in $G \setminus V$, and $Z_{V, \phi}(\lambda)$ is a normalization constant—the *partition function*—defined as

$$Z_{V, \phi}(\lambda) := \sum_{\substack{\psi \in \text{Hom}(\mathbb{G}, \mathbb{H}) \\ \psi|_V \cup \phi|_{G \setminus V} \in \text{Hom}(\mathbb{G}, \mathbb{H})}} \prod_{x \in V} \lambda(\psi(x)).$$

We will call the collection of probability measures $\{\mathbb{P}_{V, \phi}\}$, the **Gibbs** $(\mathbb{G}, \mathbb{H}, \lambda)$ -**specification**. The **boundary** of a set $V \subseteq G$, denoted by ∂V , is defined as the set of elements in G at

distance exactly 1 from V . Notice that $\mathbb{P}_{V,\phi}$ depends exclusively on $\phi|_{\partial V}$. Now, consider events of the form

$$A(\phi, V) = \{\psi \in \text{Hom}(\mathbb{G}, \mathbb{H}) : \psi|_V = \phi|_V\}.$$

Next, consider the σ -algebra \mathcal{F} generated by all events of the form $A(\phi, V)$ for V finite, and define $\mathcal{M}(\mathbb{G}, \mathbb{H})$ to be the set of probability measures on $(\text{Hom}(\mathbb{G}, \mathbb{H}), \mathcal{F})$.

A measure $\mu \in \mathcal{M}(\mathbb{G}, \mathbb{H})$ is a **Gibbs measure** for the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification if for any finite $V \subseteq G$ and for all $\psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$,

$$\mu(A(\psi, V) | A(\phi, G \setminus V)) = \mathbb{P}_{V,\phi}(\{\psi\}) \text{ for } \mu\text{-a.e. } \phi \in \text{Hom}(\mathbb{G}, \mathbb{H}).$$

In other words, the probability distribution of a random ψ inside a finite V conditioned on its values outside V to coincide with those of ϕ , depends only on the values of $\psi|_V$ and on the boundary, $\phi|_{\partial V}$. Furthermore, the conditional distribution is the same as for $\mathbb{P}_{V,\phi}$ (see also [13, Definition 2.1]).

If $\text{Hom}(\mathbb{G}, \mathbb{H}) \neq \emptyset$, then there always exists at least one Gibbs measure [30, Chapter 4]. A fundamental question in statistical physics is whether there exists a unique Gibbs measure or multiple for a given Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification.

5.2. Non-uniqueness and spatial mixing properties. In [13], it is shown that if \mathbb{H} is a graph and it is dismantlable (or equivalently, by Lemma 3.3, its square dismantles to a subgraph of its diagonal), then, for any graph \mathbb{G} of bounded degree (and therefore, locally finite), there exists some λ such that there is a unique Gibbs measure [13, Theorem 7.2]. Conversely, in [13] it is also proved that if \mathbb{H} is a non-dismantlable graph, then there exists \mathbb{G} (of bounded degree) such that for any λ there exists multiple Gibbs measures [13, Theorem 8.2].

Here, following a similar path, we show that when extending this question to arbitrary relational structures, the first implication does not remain true in general, but the second still holds. More exactly,

Proposition 5.1. *There exists a finite τ -structure \mathbb{H} such that \mathbb{H}^2 dismantles to a substructure of its diagonal and a τ -structure \mathbb{G} of bounded degree such that for any λ there exists multiple Gibbs measures for the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification. Moreover, \mathbb{H} can be chosen so that \mathbb{H}^2 dismantles to its full diagonal $\Delta(\mathbb{H}^2)$.*

Proposition 5.2. *Let \mathbb{H} be a finite τ -structure. If \mathbb{H}^2 does not dismantle to a substructure of $\Delta(\mathbb{H}^2)$, then there exists a τ -structure \mathbb{G} of bounded degree such that for any λ there exists multiple Gibbs measures for the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification.*

Before proving these two results, we introduce and explore some *spatial mixing* properties in this same context.

Definition 5.3. *Given $J \subseteq H$, we say that a Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification satisfies **spatial J -mixing (J-SM)** if there exists constants $C, \alpha > 0$ such that for all $\phi_1, \phi_2 \in \text{Hom}(\mathbb{G}, \mathbb{H})$, for all finite $V \subseteq G$, and for all $x \in V$ and $a \in H$,*

$$(2) \quad \left| \mathbb{P}_{V,\phi_1}(\{\psi(x) = a\}) - \mathbb{P}_{V,\phi_2}(\{\psi(x) = a\}) \right| \leq C \cdot \exp(-\alpha \cdot \text{dist}(x, D_V^J(\phi_1, \phi_2))),$$

where

$$D_V^J(\phi_1, \phi_2) = \{x \in \partial V : (\phi_1(x), \phi_2(x)) \in H^2 \setminus \Delta(J^2)\}$$

and $\{\psi(x) = a\}$ refers to the event that a random ψ takes the value a at x .

The definition of J -SM unifies and interpolates two well-known properties. If $J = \emptyset$, then $D_V^\emptyset(\phi_1, \phi_2) = \partial V$ and Equation (2) corresponds to the definition of **weak spatial mixing (WSM)**, i.e., \emptyset -SM. On the other hand, if $J = H$, then $D_V^H(\phi_1, \phi_2) = \{x \in \partial V : \phi_1(x) \neq \phi_2(x)\}$ and Equation (2) corresponds to the definition of **strong spatial mixing (SSM)**, i.e., H -SM.

In general, spatial mixing properties are forms of *correlation decay* that have been of interest because of their many applications. On the one hand, WSM is related with uniqueness of Gibbs measures and the absence of phase transitions [24]. On the other hand, SSM is a strengthening of WSM and it is related to the absence of *boundary phase transitions* [47] and has connections with the existence of FPTAS for #P-hard counting problems [5, 54], mixing time of Glauber dynamics [24], and efficient approximation algorithms for thermodynamic quantities [29, 9].

In [10], there were explored sufficient and necessary conditions for a graph \mathbb{H} to have, for any graph \mathbb{G} of bounded degree, the existence of a weight function λ such that the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification satisfies WSM and SSM. In particular, it was proved that dismantlability was equivalent to the existence of Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specifications satisfying WSM for any graph \mathbb{G} of bounded degree, and therefore uniqueness, since WSM implies it. In addition, it was observed that a direct consequence is that a necessary condition for SSM to hold is that \mathbb{H} is dismantlable, because SSM implies WSM. However, it was also shown that it is not a sufficient condition. Here, we strengthen this necessary condition and extend it to the realm of relational structures.

Proposition 5.4. *If \mathbb{H}^2 does not dismantle to a substructure of $\Delta(\mathbb{H}^2)$ whose universe contains $\Delta(J^2)$, then there exists a τ -structure \mathbb{G} of bounded degree such that the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification does not satisfy J -SM for any λ .*

Two direct corollaries of this fact are the following.

Corollary 5.5. *If \mathbb{H}^2 does not dismantle to some substructure of the diagonal $\Delta(\mathbb{H}^2)$, then there exists a τ -structure \mathbb{G} of bounded degree such that the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification does not satisfy WSM for any λ .*

Corollary 5.6. *If \mathbb{H}^2 does not dismantle to the full diagonal $\Delta(\mathbb{H}^2)$, then there exists a τ -structure \mathbb{G} of bounded degree such that the Gibbs $(\mathbb{G}, \mathbb{H}, \lambda)$ -specification does not satisfy SSM for any λ .*

5.3. Proofs of the propositions. Before proving the propositions, we will need to state and prove the following lemma.

Lemma 5.7. *Let \mathbb{H} be a finite τ -structure with universe H and let $J \subseteq H$. Suppose that \mathbb{H}^2 does not dismantle to a substructure of $\Delta(\mathbb{H}^2)$ whose universe contains $\Delta(J^2)$. Then, there exists an infinite τ -structure of bounded degree \mathbb{G} , $x_0 \in G$, and $\rho \in \text{Hom}(\mathbb{G}, \mathbb{H}^2)$ such that, for $\phi_i = \pi_i \circ \rho$ ($i = 1, 2$) and any cofinite set $U \subseteq G \setminus \{x_0\}$ with $\rho^{-1}(\Delta(J^2)) \subseteq U$, we have that*

$$(3) \quad X_1 \cap X_2 = \emptyset,$$

for $X_i := \{\psi(x_0) : \psi \in \text{Hom}(\mathbb{T}_{\mathbb{K}}, \mathbb{H}), \psi|_U = \phi_i|_U\}$ and $i = 1, 2$.

Proof. Suppose that \mathbb{H}^2 does not dismantle to a substructure \mathbb{K} with universe K satisfying $\Delta(J^2) \subseteq K \subseteq \Delta(H^2)$. Let \mathbb{I} be any J -non-foldable relational structure obtained by dismantling \mathbb{H} and let \mathbb{K} be the $\Delta(J^2)$ -non-foldable relational structure given by Lemma 4.1. It follows easily (for example, see the proof of Lemma 4.5) that \mathbb{H}^2 dismantles to \mathbb{K} and that

there is a retraction r from \mathbb{H}^2 to \mathbb{K} such that the image of every diagonal element is also a diagonal element (that is, such that $r(\Delta(H^2)) \subseteq \Delta(H^2)$).

By our assumption, \mathbb{K} contains some non-diagonal element (a_1, a_2) . Let $w_0 \in T_{\mathbb{K}}$ be the (unique) walk of length 0 starting in (a_1, a_2) , and let $U \subseteq T_{\mathbb{K}}$ be any cofinite set such that $w_0 \notin U$ and $\rho_{\mathbb{K}}^{-1}(\Delta(J^2)) \subseteq U$.

Note that $\text{Hom}(T_{\mathbb{K}}, \mathbb{H})$ contains mappings $\phi_i = \pi_i \circ \rho_{\mathbb{K}}$ for $i = 1, 2$. Also, let

$$X_i = \{\psi(w_0) : \psi \in \text{Hom}(T_{\mathbb{K}}, \mathbb{H}), \psi|_U = \phi_i|_U\}.$$

We claim that

$$(4) \quad X_1 \cap X_2 = \emptyset.$$

Let us prove it by contradiction. Assume that $\psi_1(w_0) = \psi_2(w_0)$, where $\psi_i \in \text{Hom}(T_{\mathbb{K}}, \mathbb{H})$ and $\psi_i|_U = \phi_i|_U$ for $i = 1, 2$. Note that $\Psi(x_0) = r \circ (\psi_1(x_0), \psi_2(x_0))$ defines a homomorphism from $T_{\mathbb{K}}$ to \mathbb{K} that agrees with $\rho_{\mathbb{K}}$ in U . Since \mathbb{K} is $\Delta(J^2)$ -non-foldable and $\rho_{\mathbb{K}}^{-1}(\Delta(J^2)) \subseteq U$, it follows by Lemma 2.8 that $\Psi(w_0) = (a_1, a_2)$. However, this is a contradiction since $(\psi_1(w_0), \psi_2(w_0))$ is a diagonal element and r sends diagonal elements to diagonal elements. Finally, by identifying \mathbb{G} with $T_{\mathbb{K}}$, ρ with $\rho_{\mathbb{K}}$, and a with w_0 , we conclude. \square

Proof of Proposition 5.2. Let \mathbb{G} , $x_0 \in G$, and ρ be as in Lemma 5.7 for $J = \emptyset$. Let λ be any weight function for \mathbb{H} . Construct a Gibbs measure μ_1 by taking *weak limits* using $\phi_1 = \pi_1 \circ \rho$, this is to say, we consider the sequence of measures $\{\mathbb{P}_{V_n, \phi_1}\}_n$, where V_n is an increasing sequence (in the sense of inclusion) of finite sets containing x_0 eventually exhausting \mathbb{G} . Such a sequence must have a subsequence $\{\mathbb{P}_{V_{n_k}, \pi_1 \circ \rho}\}_k$ *weakly converging* to a Gibbs measure μ_1 (i.e., $\lim_{k \rightarrow \infty} \mathbb{P}_{V_{n_k}, \pi_1 \circ \rho}(A(\phi, V)) = \mu_1(A(\phi, V))$ for all ϕ and finite V). This is a standard argument used for constructing Gibbs measures [30, Chapter 4]. Similarly, the sequence $\{\mathbb{P}_{V_{n_k}, \pi_2 \circ \rho}\}_k$ must contain a subsequence converging to a Gibbs measure μ_2 .

For every $a \in H$, we have that

$$\lim_{k \rightarrow \infty} \mathbb{P}_{V_{n_k}, \phi_1}(\{\psi(x_0) = a\}) = \mu_1(\{\psi(x_0) = a\}).$$

Hence, choose any $a \in H$ such that $\mu_1(\{\psi(x_0) = a\}) > 0$ (that it has to exist, since μ_1 is a probability measure). It follows from Equation (4), by setting $U = G \setminus V_{n_k}$, that

$$\lim_{k \rightarrow \infty} \mathbb{P}_{V_{n_k}, \phi_2}(\{\psi(x_0) = a\}) = 0.$$

Consequently, $\mu_2(\{\psi(x_0) = a\}) = 0$ and hence, $\mu_1 \neq \mu_2$. \square

Proof of Proposition 5.4. This is a direct consequence of the results proved in Proposition 5.2. Let us fix some weight function λ for \mathbb{H} . Assume \mathbb{H}^2 does not dismantle to a substructure of $\Delta(\mathbb{H}^2)$ whose universe contains $\Delta(J^2)$ and let \mathbb{G} , x_0 , and ρ as in Lemma 5.7. Define $\phi_i = \pi_i \circ \rho$ for $i = 1, 2$. Pick m large enough and let $V \subseteq G$ be the set of all elements $x \in G$ that are at distance less than m from x_0 and such that $\rho(x) \notin \Delta(J^2)$.

Since \mathbb{P}_{V, ϕ_1} is a probability measure, there exists some element $a \in H$ such that

$$\mathbb{P}_{V, \phi_1}(\{\psi(x_0) = a\}) \geq \frac{1}{|H|}.$$

It follows from Equation (4) in Proposition 5.2 (by setting $U = G \setminus V$), that

$$\mathbb{P}_{V, \phi_2}(\{\psi(x_0) = a\}) = 0.$$

Consequently, we have that $\mathbb{P}_{V,\phi_1}(\{\psi(x_0) = a\}) - \mathbb{P}_{V,\phi_2}(\{\psi(x_0) = a\}) \geq \frac{1}{|H|}$. Note that $\text{dist}(x_0, D_V^J(\phi_1, \phi_2)) \geq m$. Indeed, it follows from the definition of V that if the distance of $x \in \partial V$ to x_0 is less than m , then $\rho(x) \in \Delta(J^2)$, which implies that $(\phi_1(x), \phi_2(x)) \notin H^2 \setminus \Delta(J^2)$. It follows that for any constants $C, \alpha > 0$, the quantity $C \cdot \exp(-\alpha \cdot \text{dist}(x_0, D_V^J(\phi_1, \phi_2)))$ from the definition of J -SM can be made arbitrarily small by choosing m large enough. Therefore, J -SM cannot hold. \square

Proof of Proposition 5.1. Let $\tau = \{R_1, R_2, R_3\}$ be a signature with R_i a 2-ary relation for $i = 1, 2, 3$. Consider the τ -structure \mathbb{H} with universe $H = \{0, 1, 2\}$ and

- $R_1(\mathbb{H}) = \{(0, 0), (0, 1), (1, 0)\}$,
- $R_2(\mathbb{H}) = \{(1, 1), (1, 2), (2, 1)\}$, and
- $R_3(\mathbb{H}) = \{(2, 2), (2, 0), (0, 2)\}$.

It can be checked that \mathbb{H}^2 dismantles to $\Delta(\mathbb{H}^2)$. Indeed, it suffices to fold $(0, 1)$ and $(1, 0)$ to $(0, 0)$, $(1, 2)$ and $(2, 1)$ to $(1, 1)$, and $(0, 2)$ and $(2, 0)$ to $(2, 2)$. This τ -structure is intimately related to the so-called *hardcore model*, a system well studied in combinatorics and statistical physics [28, 54] consisting of a Gibbs $(\mathbb{G}, \mathbb{H}, \lambda_{\text{HC}})$ -specification for an arbitrary graph \mathbb{G} , the graph \mathbb{H} with universe $\{a, b\}$ (where we think that a is a *particle* and b is a *non-particle*) and edge relation $E = \{(a, b), (b, b)\}$, and a weight function λ_{HC} such that $\lambda_{\text{HC}}(b) = 1$ and $\lambda_{\text{HC}}(a) > 0$, usually called *activity*.

Now, consider the τ -structure \mathbb{G} consisting of 3 copies of the Δ -regular tree for an arbitrary $\Delta \geq 6$, so that the graph adjacency relation in the i th copy is given by R_i , for $i = 1, 2, 3$. Let $\lambda : \{0, 1, 2\} \rightarrow \mathbb{R}^+$ be an arbitrary weight function. Then, we can think that we have 3 copies of the hardcore model (with values $\{0, 1\}$, $\{1, 2\}$, and $\{2, 0\}$, respectively) on a Δ -regular tree with activities $\lambda(1)/\lambda(0)$, $\lambda(2)/\lambda(1)$, and $\lambda(0)/\lambda(2)$, respectively (i.e. the ratios between the weight of particle versus non-particle). By symmetry, w.l.o.g., suppose that $\lambda(1)/\lambda(0) \geq 1$. The *critical activity* for the hardcore model in a Δ -regular tree is given by the formula $\frac{(\Delta-1)^\Delta}{(\Delta-2)^\Delta}$ [37, 54], i.e., if the activity is below this threshold there exists a unique Gibbs measure (*subcritical regime*), and if it is below, there exist multiple ones (*supercritical regime*). Since $\frac{(\Delta-1)^\Delta}{(\Delta-2)^\Delta} < 1 \leq \lambda(1)/\lambda(0)$ for $\Delta \geq 6$, then the first copy corresponds to a hardcore model in a Δ -regular tree in the supercritical regime, inducing multiple Gibbs measure on $\text{Hom}(\mathbb{G}, \mathbb{H})$. \square

6. FINITE DUALITY REVISITED

Throughout this section all relational structures are assumed to be finite. We say that a τ -structure \mathbb{H} is a **core** if every homomorphism from \mathbb{H} to \mathbb{H} is one-to-one. An **obstruction** to \mathbb{H} is a τ -structure \mathbb{O} that admits no homomorphism to \mathbb{H} ; the obstruction \mathbb{O} is **critical** if every *proper* substructure (i.e., any substructure different from \mathbb{O} itself) admits a homomorphism to \mathbb{H} . A relational structure \mathbb{H} is said to have **finite duality** if it has only finitely many critical obstructions $\{\mathbb{O}_1, \dots, \mathbb{O}_m\}$. This implies that, for every τ -structure \mathbb{G} ,

$$\text{Hom}(\mathbb{O}_i, \mathbb{G}) \neq \emptyset \text{ for some } 1 \leq i \leq m \iff \text{Hom}(\mathbb{G}, \mathbb{H}) = \emptyset.$$

We say that a τ -structure \mathbb{H} **contains all constants** if for every $a \in H$ there exists $R_a \in \tau$ such that $R_a(\mathbb{H}) = \{a\}$. Note that every such relational structure is a core.

The main result in [42] states that a core relational structure \mathbb{H} has finite duality if and only if \mathbb{H}^2 dismantles to its diagonal. In this section we shall see how this result follows from Theorem 3.6. In addition, we shall show that, when \mathbb{H} contains all constants, having

finite duality is equivalent to having finitely many critical τ -tree obstructions, which was not previously known.

Theorem 6.1. *Let \mathbb{H} be a finite τ -structure which is a core. Then, the following are equivalent:*

- (A1c) \mathbb{H}^2 dismantles to its full diagonal;
- (D1c) \mathbb{H} has finitely many critical obstructions.

Furthermore, if \mathbb{H} contains all the constants, then the following statement is also equivalent:

- (D2c) \mathbb{H} has finitely many critical τ -tree obstructions.

Proof. As mentioned earlier, the equivalence (A1c) \Leftrightarrow (D1c) was shown in [42]. Here we provide an alternative proof.

(A1c) \Rightarrow (D1c). Assume that statement (A1c) holds. It follows that statement Theorem 3.6(C1) holds for any $J \subseteq H$ (although we note that in this proof it suffices the case $J = \emptyset$). We claim that the diameter of the critical obstructions of \mathbb{H} is bounded, where the diameter of a relational structure is the maximum distance between any pair of its elements. It follows easily (see [42, Lemma 2.4]) that our claim implies statement (D1c).

To prove the claim, assume towards a contradiction that there exists a critical obstruction \mathbb{G} containing two elements x and y at distance at least $g + 2$ (where g is the gap given by statement Theorem 3.6(C1)). Let $R \in \tau$ and let $\mathbf{x} = (x_1, \dots, x_k) \in R(\mathbb{G})$ in which x occurs. Since \mathbb{G} is critical, it follows that the substructure obtained by removing \mathbf{x} from $R(\mathbb{G})$ has a homomorphism ϕ to \mathbb{H} . Similarly, the substructure obtained by removing \mathbf{y} from some relation $S(\mathbb{G})$ a tuple $\mathbf{y} = (y_1, \dots, y_{k'})$ where y occurs, has a homomorphism ψ to \mathbb{H} .

Hence, both ϕ and ψ define homomorphisms from the substructure \mathbb{K} obtained from \mathbb{G} by removing both \mathbf{a} from $R(\mathbb{G})$ and \mathbf{b} from $S(\mathbb{G})$. Note that the distance, in \mathbb{K} , from $V = \{x_1, \dots, x_k\}$ to $W = \{y_1, \dots, y_{k'}\}$ is at least g . It follows from Theorem 3.6(C1) that there is a homomorphism γ from \mathbb{K} to \mathbb{H} that agrees with ϕ on V and with ψ on W . Consequently, γ defines a homomorphism from \mathbb{G} to \mathbb{H} , a contradiction.

(D1c) \Rightarrow (D2c). It is immediate, since critical tree obstructions are critical obstructions.

(D2c) \Rightarrow (A1c). Let us prove the contrapositive. Assume that \mathbb{H}^2 does not dismantle to the diagonal or, equivalently, that \mathbb{H} does not satisfy statement Theorem 3.6(A1) when $J = H$. It follows that \mathbb{H} does not satisfy statement Theorem 3.6(C2) either. Let g be an arbitrary gap. Since the statement Theorem 3.6(C2) fails, it follows by standard compactness arguments that there is a finite substructure \mathbb{G} of $\mathbb{T}_{\mathbb{H}^2}$ and $V, W \subseteq G \subseteq \mathbb{T}_{\mathbb{H}^2}$ with $\text{dist}(V, W) \geq g$ such that $\text{Hom}(\mathbb{G}, \mathbb{H})$ is not (V, W) -mixing with respect to H . That is, there exists mappings $\phi, \psi \in \text{Hom}(\mathbb{G}, \mathbb{H})$ such that there is no mapping in $\text{Hom}(\mathbb{G}, \mathbb{H})$ that agrees with ϕ on V , with ψ on W , and with both on every element $x \in G$ such that $\phi(x) = \psi(x)$.

Now, let \mathbb{K} be the τ -structure obtained from \mathbb{G} by coloring every element $x \in V$ according to ϕ , every element in W according to ψ , and every element x such that $\phi(x) = \psi(x) \in J$ according to either ϕ or ψ . Since there is no homomorphism from \mathbb{K} to \mathbb{H} and, consequently, there is a substructure \mathbb{I} of \mathbb{K} that is a critical obstruction of \mathbb{H} . Since \mathbb{I} is critical, then it is connected. Consequently, since \mathbb{G} is a substructure of $\mathbb{T}_{\mathbb{H}^2}$ and $\mathbb{T}_{\mathbb{H}^2}$ does not contain cycles, it follows that \mathbb{I} is a τ -tree. Clearly, $V \cap \mathbb{I}$ is nonempty since otherwise the mapping $x \mapsto \psi(x)$ would define a homomorphism from \mathbb{I} to \mathbb{H} . Similarly, $W \cap \mathbb{I} \neq \emptyset$. Since \mathbb{I} is connected and the distance in \mathbb{G} (and hence in \mathbb{I}) from V to W is at least g , it follows that $|\mathbb{I}| \geq g + 1$. Since g is arbitrary, we have completed the proof.

(D1c) \Rightarrow (A1c). Assume that \mathbb{H} satisfies statement (D1c). Let \mathbb{H}_c obtained by endowing \mathbb{H} with all constants. Formally, if τ is the signature of \mathbb{H} , then τ_c is the new signature containing a new relation symbol R_a for every $a \in H$, and \mathbb{H}_c is the τ_c -structure obtained from \mathbb{H} by setting $R_a(\mathbb{H}_c) = \{a\}$ for every $a \in H$. We shall show that \mathbb{H}_c has also finite duality. Consequently, \mathbb{H}_c satisfies statement (D2c) and, hence, \mathbb{H}_c^2 dismantles to the diagonal, implying that \mathbb{H}^2 dismantles to the diagonal as well.

The proof of this claim is fairly standard. Let \mathbb{K}_c be a minimal critical obstruction of \mathbb{H}_c and let \mathbb{G} be the τ -structure constructed in the following way. In a first stage, consider the disjoint union of \mathbb{H} and \mathbb{K} , where \mathbb{K} is the τ -structure obtained by removing all constants from \mathbb{K}_c (that is, \mathbb{K} is the τ -structure obtained from \mathbb{K}_c by removing all relations in $\tau_c \setminus \tau$). In a second stage, we glue some elements from \mathbb{H} and \mathbb{K} . In particular we glue every element a in \mathbb{H} to every element $b \in R_a(\mathbb{K}_c)$. We shall show that \mathbb{G} is not homomorphic to \mathbb{H} . Assume towards a contradiction that there is ϕ that defines a homomorphism from \mathbb{G} to \mathbb{H} . Clearly, the restriction of ϕ to H , that we shall denote $\phi|_H$, defines a homomorphism of \mathbb{H} that must be one-to-one since \mathbb{H} is a core. Since $\phi|_H$ is one-to-one it follows that $\phi|_H^{-1}$ is also a homomorphism from \mathbb{H} to \mathbb{H} and, hence, $\phi|_H^{-1} \circ \phi$ defines a homomorphism from \mathbb{G} to \mathbb{H} that acts as the identity on \mathbb{H} . It follows that $\phi|_H^{-1} \circ \phi$ defines a homomorphism from \mathbb{K}_c to \mathbb{H}_c , a contradiction.

Let \mathbb{J} be any substructure of \mathbb{G} which is a critical obstruction of \mathbb{H} . It follows easily from the criticality of \mathbb{K}_c that \mathbb{J} contains \mathbb{K} . Since \mathbb{H} has finite duality it follows that there is a bound on the size of \mathbb{J} and, hence, of \mathbb{K} . \square

Remark 6.2. Notice that for the implications (A1c) \Rightarrow (D1c) and (D1c) \Rightarrow (D2c) we didn't need the fact that \mathbb{H} is a core.

It has been shown in [27] that if a τ -structure \mathbb{H} has finite duality, then there exists some finite set $\{\mathbb{T}_1, \dots, \mathbb{T}_m\}$ of τ -trees such that for every τ -structure \mathbb{I} not homomorphic to \mathbb{H} , there exists some $1 \leq i \leq m$ such that \mathbb{T}_i is homomorphic to \mathbb{I} but not homomorphic to \mathbb{H} . We want to note that the equivalence between statements (D2c) and (D1c) does not follow from this fact. Indeed, direction (D2c) \Rightarrow (D1c) does not hold when we do not require that the τ -structure \mathbb{H} is equipped with constants as witnessed by the case when \mathbb{H} is the oriented 3-cycle. Note that, in this case, \mathbb{H} satisfies (D2c) since every τ -tree is homomorphic to \mathbb{H} and, hence, \mathbb{H} has no critical τ -tree obstructions at all. However, since any oriented cycle whose length is not a multiple of 3 is a critical obstruction of \mathbb{H} , it follows that \mathbb{H} does not satisfy (D1c).

The next corollary is direct.

Corollary 6.3. *Let \mathbb{H} be a finite τ -structure which is a core. Then, the following are equivalent:*

- (A1c) \mathbb{H}^2 dismantles to its full diagonal;
- (B1c) $C(\mathbb{G}, \mathbb{H})$ is H -connected for every locally finite τ -structure \mathbb{G} ;
- (C1c) $\text{Hom}(\mathbb{G}, \mathbb{H})$ is topologically strong spatial mixing for every τ -structure \mathbb{G} ; and
- (D1c) \mathbb{H} has finitely many critical obstructions.

Notice that Corollary 6.3 has novel consequences even in the graph case. Despite graphs that have finitely many critical obstructions are the graphs without any edge, one can always consider, given a graph \mathbb{H} , a version of it endowed with all constants. Formally, we consider the τ -structure \mathbb{H} with $\tau = \{E\} \cup \{R_a\}_{a \in H}$, universe H , $E(\mathbb{H}) = E(\mathbb{H})$, and $R_a(\mathbb{H}) = \{a\}$.

First, notice that, since $\overline{\mathbb{H}}$ contains all constants, $\overline{\mathbb{H}}$ is a core. In addition, if \mathbb{H} dismantles to its full diagonal, then $\overline{\mathbb{H}}$ also dismantles to its full diagonal. Therefore, by the previous results, $\overline{\mathbb{H}}$ has a finite set $\{\mathbb{O}_1, \dots, \mathbb{O}_m\}$ of critical obstructions. Now suppose that we are given another graph \mathbb{G} and a partial H -coloring of \mathbb{G} , this is to say, a function $\phi_U : U \rightarrow H$, where $U \subseteq G$. A natural question is whether ϕ_U can be extended to a homomorphism $\phi \in \text{Hom}(\mathbb{G}, \overline{\mathbb{H}})$. This computational problem is usually called *homomorphism extension problem* [43], and turns out to be equivalent to the *retraction problem* [25].

Consider the auxiliary τ -structure $\overline{\mathbb{G}}^{\phi_U}$ with universe G , $E(\overline{\mathbb{G}}^{\phi_U}) = E(\mathbb{G})$, and $R_a(\overline{\mathbb{G}}^{\phi_U}) = \{\phi_U^{-1}(a)\}$ for $a \in H$. The previous results tells us that such extension exists if and only if $\text{Hom}(\mathbb{O}_i, \overline{\mathbb{G}}^{\phi_U}) = \emptyset$ for all $1 \leq i \leq m$. In other words, we only need to check $\overline{\mathbb{G}}^{\phi_U}$ locally to see that this is possible. It is easy to extend this observation to the more general framework of relational structures.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their careful reading of our manuscript and their many insightful comments and suggestions.

The first author was supported by CONICYT/FONDECYT Postdoctorado 3190191 and ERC Starting Grants 678520 and 676970. The second author was supported by an NSERC Discovery grant. The third author was supported by MICCIN grants TIN2016-76573-C2-1P and PID2019-109137GB-C22, and Maria de Maeztu Units of Excellence Programme MDM-2015-0502. The fourth author was supported by an NSERC Discovery grant and FRQNT.

This work was done in part while the first three authors were visiting the Simons Institute for the Theory of Computing at University of California, Berkeley.

REFERENCES

- [1] Dimitris Achlioptas, Paul Beame, and Michael Molloy. Exponential bounds for DPLL below the satisfiability threshold. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004*, pages 139–140, 2004.
- [2] Noga Alon, Raimundo Briceño, Nishant Chandogtia, Alexander Magazinov, and Yinon Spinka. Mixing properties of colorings of the \mathbb{Z}^d lattice. *Combinatorics, Probability and Computing, to appear (arXiv:1903.11685)*, 2019.
- [3] Albert Atserias. On digraph coloring problems and treewidth duality. *Eur. J. Comb.*, 29(4):796–820, 2008.
- [4] Jung-Chao Ban and Chih-Hung Chang. Tree-shifts: Irreducibility, mixing, and the chaos of tree-shifts. *Trans. Amer. Math. Soc.*, 369(12):8389–8407, 2017.
- [5] Antar Bandyopadhyay and David Gamarnik. Counting without sampling: Asymptotics of the log-partition function for certain statistical physics models. *Random Structures & Algorithms*, 33(4):452–479, 2008.
- [6] Manuel Bodirsky. The complexity of constraint satisfaction problems (invited talk). In *32nd International Symposium on Theoretical Aspects of Computer Science, STACS 2015, March 4-7, 2015, Garching, Germany*, pages 2–9, 2015.
- [7] Christian Borgs, Jennifer Chayes, László Lovász, Vera T Sós, and Katalin Vesztegombi. Counting graph homomorphisms. In *Topics in discrete mathematics*, pages 315–371. Springer, 2006.
- [8] Mike Boyle, Ronnie Pavlov, and Michael Schraudner. Multidimensional sofic shifts without separation and their factors. *Trans. Amer. Math. Soc.*, 362(9):4617–4653, 2010.
- [9] Raimundo Briceño. The topological strong spatial mixing property and new conditions for pressure approximation. *Ergodic Theory Dynam. Systems*, 38(5):1658–1696, 2018.
- [10] Raimundo Briceño and Ronnie Pavlov. Strong spatial mixing in homomorphism spaces. *SIAM J. Discrete Math.*, 31(3):2110–2137, 2017.
- [11] Raimundo Briceño, Kevin McGoff, and Ronnie Pavlov. Factoring onto \mathbb{Z}^d subshifts with the finite extension property. *Proc. Amer. Math. Soc.*, 146(12):5129–5140, 2018.
- [12] Graham R. Brightwell and Peter Winkler. Graph homomorphisms and phase transitions. *J. Combin. Theory Ser. B*, 77(2):221–262, 1999.

- [13] Graham R. Brightwell and Peter Winkler. Gibbs measures and dismantlable graphs. *J. Combin. Theory Ser. B*, 78(1):141–166, 2000.
- [14] Graham R. Brightwell and Peter Winkler. Graph homomorphisms and long range action. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 63:29–48, 2004.
- [15] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 319–330, 2017.
- [16] Andrei A. Bulatov, Víctor Dalmau, Martin Grohe, and Dániel Marx. Enumerating homomorphisms. *J. Comput. Syst. Sci.*, 78(2):638–650, 2012.
- [17] Andrei A. Bulatov, Andrei A. Krokhnin, and Benoît Larose. Dualities for constraint satisfaction problems. In *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*, pages 93–124, 2008.
- [18] Tullio Ceccherini-Silberstein and Michel Coornaert. *Cellular automata and groups*. Springer Science & Business Media, 2010.
- [19] Tullio Ceccherini-Silberstein and Michel Coornaert. On the density of periodic configurations in strongly irreducible subshifts. *Nonlinearity*, 25(7):2119, 2012.
- [20] David A. Cohen, Martin C. Cooper, Peter G. Jeavons, and Stanislav Živny. Binary constraint satisfaction problems defined by excluded topological minors. *Inf. Comput.*, 264:12–31, 2019.
- [21] Víctor Dalmau, Andrei A. Krokhnin, and Benoît Larose. First-order definable retraction problems for posets and reflexive graph. In *19th IEEE Symposium on Logic in Computer Science (LICS 2004), 14-17 July 2004, Turku, Finland, Proceedings*, pages 232–241, 2004.
- [22] Rina Dechter and David Cohen. *Constraint processing*. Morgan Kaufmann, 2003.
- [23] Martin Dyer and Catherine Greenhill. The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 17(3-4):260–289, 2000.
- [24] Martin Dyer, Alistair Sinclair, Eric Vigoda, and Dror Weitz. Mixing in time and space for lattice spin systems: A combinatorial view. *Random Structures & Algorithms*, 24(4):461–479, 2004.
- [25] Tomás Feder and Pavol Hell. List homomorphisms to reflexive graphs. *J. Comb. Theory, Ser. B*, 72(2):236–250, 1998.
- [26] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.
- [27] Jan Foniok, Jaroslav Nešetřil, and Claude Tardif. Generalised dualities and maximal finite antichains in the homomorphism order of relational structures. *Eur. J. Comb.*, 29(4):881–899, 2008.
- [28] David Galvin and Jeff Kahn. On phase transition in the hard-core model on \mathbb{Z}^d . *Comb. Probab. Comput.*, 13(2):137–164, 2004.
- [29] David Gamarnik and Dmitriy Katz. Sequential cavity method for computing free energy and surface pressure. *Journal of Statistical Physics*, 137(2):205, 2009.
- [30] Hans-Otto Georgii. *Gibbs Measures and Phase Transitions*, volume 9 of *De Gruyter Studies in Mathematics*. Berlin, 2 edition, 2011.
- [31] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM (JACM)*, 54(1):1–24, 2007.
- [32] Tatsuhiko Hatanaka, Takehiro Ito, and Xiao Zhou. Complexity of reconfiguration problems for constraint satisfaction. *CoRR*, abs/1812.10629, 2018.
- [33] Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2004.
- [34] Wilfrid Hodges. *A shorter Model Theory*. Cambridge University Press, 1997.
- [35] Takehiro Ito, Erik D. Demaine, Nicholas J. A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theor. Comput. Sci.*, 412(12-14):1054–1065, 2011.
- [36] David Kelly and Ivan Rival. Crowns, fences, and dismantlable lattices. *Canadian Journal of Mathematics*, 26(5):1257–1271, 1974.
- [37] Frank P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 379–395, 1985.
- [38] P. Kolaitis and M. Vardi. Conjunctive-Query Containment and Constraint Satisfaction. *Journal of Computer and System Sciences*, 61:302–332, 2000.
- [39] Marcin Kozik. Weak consistency notions for all the CSPs of bounded width. In *31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16*, pages 633–641, 2016.
- [40] Florent Krzakała, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *PNAS*, 104(25):10318–10323, 2007.

- [41] Gábor Kun, Ryan O'Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *Innovations in Theoretical Computer Science, ICT '12*, pages 484–495, 2012.
- [42] Benoît Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Log. Methods Comput. Sci.*, 3(4):4:6, 22, 2007.
- [43] Benoît Larose and László Zádori. The complexity of the extendibility problem for finite posets. *SIAM J. Discret. Math.*, 17(1):114–121, 2003.
- [44] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge university press, 1995.
- [45] Douglas Lind and Klaus Schmidt. Symbolic and algebraic dynamical systems. In *Handbook of dynamical systems*, volume 1, pages 765–812. Elsevier, 2002.
- [46] Brian Marcus and Joachim Rosenthal. *Codes, systems, and graphical models*, volume 123. Springer Science & Business Media, 2012.
- [47] Fabio Martinelli, Enzo Olivieri, and Roberto H. Schonmann. For 2-D lattice spin systems weak mixing implies strong mixing. *Communications in Mathematical Physics*, 165(1):33–47, 1994.
- [48] Marc Mézard and Andrea Montanari. *Information, Physics, and Computation*. Oxford University press, 2009.
- [49] Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 11(4):52, 2018.
- [50] Richard Nowakowski and Peter Winkler. Vertex-to-vertex pursuit in a graph. *Discrete Mathematics*, 43(2-3):235–239, 1983.
- [51] Ronnie Pavlov and Michael Schraudner. Entropies realizable by block gluing \mathbb{Z}^d shifts of finite type. *Journal d'Analyse Mathématique*, 126(1):113–174, 2015.
- [52] Klaus Schmidt. Multi-dimensional symbolic dynamical systems. In *Codes, systems, and graphical models*, pages 67–82. Springer, 2001.
- [53] Michael Schraudner. A matrix formalism for conjugacies of higher-dimensional shifts of finite type. *Colloq. Math.*, 110(2):493–515, 2008.
- [54] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing, STOC '06*, pages 140–149, New York, NY, USA, 2006. ACM.
- [55] Marcin Wrochna. Homomorphism reconfiguration via homotopy. *SIAM Journal on Discrete Mathematics*, 34(1):328–350, 2020.
- [56] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 331–342, 2017.

FACULTAD DE MATEMÁTICAS, PONTIFICIA UNIVERSIDAD CATÓLICA DE CHILE, SANTIAGO, CHILE
E-mail address: raimundo.briceno@mat.uc.cl

SCHOOL OF COMPUTING SCIENCE, SIMON FRASER UNIVERSITY, BURNABY BC, CANADA
E-mail address: abulatov@sfu.ca

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGIES, UNIVERSITAT POMPEU FABRA,
BARCELONA, SPAIN
E-mail address: victor.dalmau@upf.edu

LACIM, UNIVERSITÉ DU QUÉBEC A MONTRÉAL, MONTRÉAL, CANADA
E-mail address: blarose@lacim.ca