The Complexity of the Distributed Constraint Satisfaction Problem

Silvia Butti ✉ 🏠 ⓘ
Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

Victor Dalmau ✉ 🏠 ⓘ
Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

Abstract

We study the complexity of the Distributed Constraint Satisfaction Problem (DCSP) on a synchronous, anonymous network from a theoretical standpoint. In this setting, variables and constraints are controlled by agents which communicate with each other by sending messages through fixed communication channels. Our results endorse the well-known fact from classical CSPs that the complexity of fixed-template computational problems depends on the template's invariance under certain operations. Specifically, we show that DCSP(Γ) is polynomial-time tractable if and only if Γ is invariant under symmetric polymorphisms of all arities. Otherwise, there are no algorithms that solve DCSP(Γ) in finite time. We also show that the same condition holds for the search variant of DCSP.

Collaterally, our results unveil a feature of the processes' neighbourhood in a distributed network, its iterated degree, which plays a major role in the analysis. We explore this notion establishing a tight connection with the basic linear programming relaxation of a CSP.

1 Introduction

The Constraint Satisfaction Problem (CSP) consists of a collection of variables and a collection of constraints where each constraint specifies the valid combinations of values that can be taken simultaneously by the variables in its scope. The goal is to decide if there exists an assignment of the elements of a domain to the variables which satisfies all constraints. The CSP is a very rich mathematical framework that is widely used both as a fruitful paradigm for theoretical research, and as a powerful tool for applications in AI, such as scheduling and planning [21, 17].

38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021).
Editors: Markus Bläser and Benjamin Monmege; Article No. 20; pp. 20:1–20:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

While, in its full generality, the finite-domain CSP is known to be NP-complete, applying specific restrictions on the instances can yield tractable subclasses of the problem. One of the most studied approaches consists in requiring that, in each constraint, the set of allowed combinations for its values be drawn from a prescribed set Γ, usually called the constraint language or the template. Thanks to the proof of the CSP dichotomy conjecture obtained separately in [9] and [28], which culminated a decades-long research program, it is possible to determine the complexity (P or NP-complete) of each family of CSPs, CSP(Γ), which is obtained by fixing Γ. This proof confirmed that the complexity of the constraint satisfaction problem is deeply tied to certain algebraic properties of the constraint language. Specifically, it depends on whether or not the constraint language is invariant under certain operations known as its polymorphisms. The polymorphisms of a constraint language enforce a symmetry on the space of solutions of a CSP instance that can possibly be exploited by an algorithm. This connection with algebra is also present in our work.

We study the computational complexity of the distributed counterpart of CSP, which is known as DCSP. This was introduced by Yokoo et al. [25] as a formal framework for the study of cooperative distributed problem solving. In particular, we consider a deterministic, synchronous, anonymous network of agents controlling variables and constraints, and we study the complexity of message passing algorithms on this network. A number of practical applications can be encoded in the DCSP model, for instance resource allocation tasks in wireless networks, routing, networking, and mobile technologies (see for instance [10, 6]).

We notice that this framework is general enough to encompass some simple Graph Neural Network architectures (see for example [20, 14]). In particular, when training a GNN to classify graphs, it is customary that the GNN network ignores the node label when updating its feature vector. This is, in fact, essential as otherwise there would be no way to apply the network trained on a given graph to another one. However, whereas in all variants of GNNs the computation is limited to a reduced number of operations over feature vectors, in the DCSP model the computation at each node is governed by an arbitrary algorithm. GNNs have a wide range of applications including molecule classification or image classification (see [5] for example). Recently, GNNs have been deployed to solve CSPs [22].

While there are a variety of well-performing distributed algorithms for constraint satisfaction and optimisation (see for instance [27, 19, 11]), the theoretical aspects of distributed complexity are to date not well understood. In this paper we initiate the study of the complexity of DCSP parametrized by the constraint language, obtaining a complete characterization of its tractable classes. More specifically, building on the connection between the CSP and algebra, we show that for any finite constraint language Γ, the decision problem for DCSP(Γ) is tractable whenever Γ is invariant under symmetric polymorphisms of all arities, where an operation is symmetric if its result does not depend on the order of its arguments. Otherwise, there are no message passing algorithms that solve DCSP(Γ). Collaterally, we show that the same holds for the search problem for DCSP.

Our work begins with the identification of a feature of the nodes in a distributed network, its iterated degree, which plays a major role in how messages are transmitted in the network. The iterated degree is an extension of the similar concept introduced in the study of the isomorphism problem which turns out to have a variety of alternative characterizations in terms of fractional isomorphisms, the Weisfeiler-Leman test, and definability with counting logics (see [14]). It turns out that, due to the network anonymity, in every distributed algorithm all equivalent agents (with respect to iterated degree) must necessarily behave identically at each round. A similar phenomenon has been observed independently in the context of GNNs in [20, 23] leading to further study in [3].

We use this fact to show that, under the absence of symmetric polymorphisms of any arity in Γ, it is always possible to construct two instances of DCSP(Γ), one satisfiable and the other unsatisfiable, that cannot be distinguished by any message passing algorithm in an anonymous network.

On the other hand, invariance under symmetric polymorphisms is connected with the basic linear programming relaxation of a CSP instance. More precisely, if Γ has symmetric polymorphisms of all arities then one can decide the satisfiability of every instance of CSP(Γ) by checking whether its basic linear programming relaxation is feasible (see for instance [4]). Whereas it is not clear how to directly use this fact to obtain a distributed algorithm for DCSP(Γ), it can be applied to establish a structure theorem that unveils a simple yet surprising structure in the solution space of every satisfiable instance in DCSP(Γ): it must contain a solution that assigns the same value to all variables that have the same iterated degree. The proof of the structure theorem uses the weighted majority algorithm, a weight update method that is widely used in optimisation and machine learning applications (see [2]). The structure theorem is key in the proof of the positive results as it allows to run an adapted variant of the *jpq*-consistency algorithm [16] that overcomes the absence of unique identifiers for the variables, by using instead their iterated degree.

This paper is organised as follows. In Section 2 we introduce some definitions and technical concepts about the DCSP model. In Section 3 we present the basic LP relaxation for CSPs and we show its connection to the symmetry on the solution space, culminating in the statement of the structure theorem. Section 4 is dedicated to the proof of the dichotomy theorem for the complexity of DCSP, with the hardness results in Section 4.1, the details of the distributed algorithm for tractable languages in Section 4.2, and its extension to the search problem in Section 4.3. Finally, in the Conclusion we discuss some directions into which our work could be extended.

2   Preliminaries

Constraint Satisfaction Problems

An instance I of the finite-domain *Constraint Satisfaction Problem* (CSP) is a triple (X, D, C) where X is a set of variables, D is a finite set called the domain, and C is a set of constraints where a constraint $c \in C$ is a pair $(\mathbf{s}, R)$ where $R \subseteq D^k$ for k a positive integer, R is a relation over D of arity k, and $\mathbf{s}$ is a tuple of k variables, known as the *scope* of c. We use $arity(\cdot)$ to denote the arity of a relation, tuple, or constraint and we write $x \in c$ for any variable x in the scope of c. An *assignment* $\nu : X \to D$ is said to be *satisfying* if for all constraints $c = (\mathbf{s}, R) \in C$ we have $\nu(\mathbf{s}) \in R$, where ν is applied to $\mathbf{s}$ coordinate-wise. Usually we denote the number of variables by n and the number of constraints by m.

Let Γ be a set of relations over some finite domain D, and let CSP(Γ) denote the set of CSP instances with all constraint relations lying in Γ. In this context, Γ is known as the *constraint language*. Throughout this paper, we will assume that Γ is always finite. Then, the *decision problem* for CSP(Γ) is the problem of deciding whether a satisfying assignment exists for an instance $I \in$ CSP(Γ). The *search problem* for CSP(Γ) is the problem of deciding whether a satisfying assignment exists and, if it does, to find one such assignment.

The Distributed Model

We consider the DCSP model of [25] with some small modifications. The basic idea is to assign the task of solving a constraint satisfaction problem to a multi-agent system. In the original model, which assumes that all constraints are binary [26, 27], the assumption is that

each variable is controlled by an agent, and two agents can communicate with one another if and only if they share a constraint. Here we deviate slightly from the original model to allow for non-binary constraints and we assume that both variables and constraints are controlled by distributed agents in the network. An instance of the *Distributed Constraint Satisfaction Problem* (DCSP) is a tuple (A, X, D, C, α), where X, D, and C are as in the classical CSP, A is a finite set of agents, and $\alpha : X \cup C \to A$ is a surjective function which assigns the control of each variable $x \in X$ and each constraint $c \in C$ to an agent $\alpha(x)$, $\alpha(c)$ respectively. For the purpose of this paper, we assume that there are exactly $n + m$ agents, and therefore each agent controls exactly one variable or one constraint. This can be done without loss of generality since any agent controlling multiple nodes can simulate multiple agents, each controlling a node. Under this assumption, there is a one-to-one correspondence between instances of CSP and DCSP, and thus we shall switch freely between them.

Distributed Networks and Message Passing

We now present some fundamental concepts relating to the message-passing paradigm for distributed networks. For a general introduction to distributed algorithms, we refer the reader to [12]. A *distributed system* consists of a finite set of nodes or processes, which are connected through communication channels to form a network. Any process in the network can perform events of three kinds: *send*, *receive* and *internal*. Send and receive events are self-explanatory, as they denote the sending or receiving of a message over a communication channel. Any kind of local computation performed at the process level, as well as state changes and decisions, are classified as internal events.

We assume a fully *synchronous* communication model, meaning that the send event at a process a and the corresponding receive event at a process a' can be considered *de facto* as a unique event, with no time delay. As a whole, a synchronous system proceeds in rounds, where at each round a process can perform some internal computation and then send messages to and receive messages from its neighbours. A round needs to terminate at every process before the next round begins. Note that while for simplicity we assume a synchronous network, all our algorithms can be adapted to asynchronous systems by applying a simple synchronizer. Nonetheless, we point out that our negative results rely on the network operating in synchronous rounds.

We make the fundamental assumption that the network is *anonymous*, meaning that variables, constraints and agents do not have IDs. For practical purposes, we still refer to variables and constraints with names (such as x_i, c_i), however these cannot be communicated through the channels. The assumption of anonymity can have various practical justifications: the processes may actually lack the hardware to have an ID, or they may be unable to reveal their ID due to security or privacy concerns. For instance, the basic architecture of GNNs is anonymous. This is a very desirable property as it allows to deploy GNNs in different networks than those in which they were trained.

We assume that all the processes run locally the same deterministic algorithm, therefore IDs cannot be created and deadlocks cannot be broken by for instance flipping a random coin. Hence, the lack of IDs makes the processes essentially indistinguishable from one another - except, as we will see later, for the structure of their neighbourhood in the network.

Leader election is a procedure by which the processes in a network select a single process to be the leader in a distributed way. If a leader is elected, then she can assign unique identifiers to every process. Moreover, all the information about the instance can be gathered to the leader, who can then solve the CSP locally. It is a well-known result that there does not exist a terminating deterministic algorithm for electing a leader in an anonymous ring [1].

Therefore, the assumptions of anonymity and determinism ensure that the DCSP model is intrinsically different from the (centralised) CSP framework, and open up the way for establishing novel, non-trivial complexity results. We remark that while considerable effort has been put into characterizing under what conditions an anonymous network is able to elect a leader [7, 24] or compute relations [8], our work focuses on characterizing the complexity of the DCSP as parametrised by the constraint language. Therefore, all of our algorithms work regardless of the topology of the network, and hence regardless of whether or not a leader can be elected.

The encoding of a DCSP instance into the message passing framework is straightforward. The processes correspond to the agents of the network, and there is a labelled communication channel between a variable agent $\alpha(x)$ and a constraint agent $\alpha(c)$ if and only if $x \in c$. More formally, the *Factor Graph* [11] G_I of an instance $I = (X, D, C)$ of CSP is the undirected bipartite graph with vertex set $X \cup C$ and edge set $\{\{x, c\} \mid x \in c\}$. Each edge in G_I that is incident to a variable x and a constraint c where $c = (\mathbf{s}, R)$ has a label $\ell_{x,c} = (S, R)$ for $S = \{i \mid \mathbf{s}[i] = x\}$, where for a tuple $\mathbf{t}$, $\mathbf{t}[i]$ denotes the i^{th} entry of $\mathbf{t}$.[1] Then, the message passing network corresponds to the factor graph where every node (variable or constraint) is replaced by their associated agent and every edge by a communication channel of the same label. Note that between any two nodes there is at most one channel. If privacy is a concern, we point out that labeling channels does not reveal any more information about the processes than what is strictly necessary for the problem instance to be well defined. Unless explicitly stated we only consider instances whose factor graph consists of a unique connected component. It is easy to prove (see the full version for details) that in the case that all relations are binary, the original model where only variables are controlled by agents is equivalent to our model.

At the start of an algorithm, a process only has access to very limited information. All processes know the total number n of variables in the CSP instance, the total number m of constraints, the labels of the communication channels that they are incident to in the network, and naturally whether they are controlling a variable or a constraint. During a run of the algorithm a process can acquire further knowledge from the messages that it receives from its neighbours. We assume that at any time each process is in one of a set of states, a subset of which are terminating states. When it enters a terminating state, a process performs no more send or internal events, and all receive events are disregarded. The local algorithm is then a deterministic function which determines the process' next state, and the messages it will send to its neighbours. The output of such function only depends on the process' current knowledge, on its state, and on the global time. We allow processes to send different messages through different channels. However, since processes can only distinguish the channels based on their labels, identical messages must be sent through channels with identical labels. Note that the power of the model would not decrease if only one message was allowed to be passed through all the channels, since a process can simulate sending a separate message through each channel by tagging each message with the label of the desired channel and concatenating them in a unique string. This, however, comes at the cost of increased message size. Moreover, if a process needs to broadcast multiple messages, these can be concatenated into one. We say that an algorithm terminates when all processes are in a terminating state.

---

[1] For mathematical clarity, we label edges with the relation itself. However, in algorithmic applications, every relation can be substituted with a corresponding symbol.
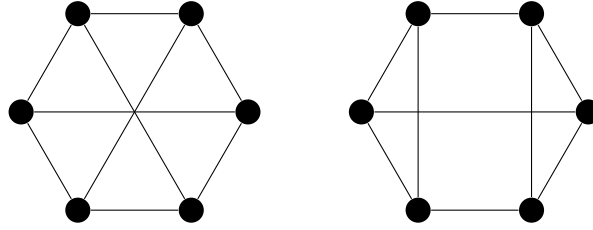
■ **Figure 1** Both graphs depicted above are 3-regular and hence they have the same iterated degree sequence. However, they are clearly not isomorphic, since the left graph is bipartite while the right one is not.

We say that a distributed algorithm solves an instance I of DCSP if the algorithm terminates and the terminating state of every process correctly states that I is satisfiable if it is, and that it is not satisfiable otherwise. Moreover, we consider the search version of DCSP, denoted DCSP-Search. In the search version, if the input instance I is satisfiable, the terminating state of every variable process $\alpha(x)$ must additionally specify a value $\nu(x) \in D$ such that $\nu : X \to D$ is a satisfying assignment. For every constraint language Γ, we denote by DCSP(Γ) and DCSP-Search(Γ) the restrictions of DCSP and DCSP-Search, respectively, to instances containing only constraint relations from Γ.

In terms of algorithmic complexity, there are a number of measures that can be of interest. Time complexity, which is our primary concern, corresponds to the total amount of time required for the algorithm to terminate, including the time needed for internal events. This is closely related to the number of rounds of the algorithm, which is another measure that we are concerned with. Message complexity and bit complexity measure the total number of messages and bits exchanged respectively. These can be bounded easily from the maximum size of a message.

Iterated Degree and Degree Sequence

We present a number of concepts from graph theory that carry over to CSPs. Their adaptation to DCSPs is straightforward in all cases. In an undirected graph G, the degree of a vertex v is the number of edges incident at v. The zeroth iterated degree of v is equal to its degree. For $k \geq 1$, the k^{th} iterated degree of v is the multiset of $(k-1)^{th}$ degrees of v's neighbours in G. The k^{th} iterated degree sequence of a graph is the multiset of k^{th} iterated degrees of its vertices.

▶ **Example 2.1.** In the context of graph theory the *colour refinement algorithm*, which calculates the iterated degree sequence of a graph, is often used as a simple heuristic for the graph isomorphism problem. If two graphs are isomorphic then they must have the same iterated degree sequence, but the opposite is not true (see for example Figure 1). ⌟

We extend the notion of iterated degree to CSPs as follows. Consider the labelled factor graph G_I of an instance I described in the previous paragraph. In what follows it will be convenient to allow instances I with a disconnected factor graph G_I. Let v be a node of G_I and denote its neighbourhood in the factor graph by $N(v)$. The (zeroth) *degree*, denoted $\delta_0(v)$, of a node in the factor graph is simply a symbol that distinguishes variables from constraints: we set $\delta_0(x) =$ "$\bullet$" for all $x \in X$ and $\delta_0(c) =$ "$\blacktriangle$" for all $c \in C$. The k^{th} *iterated*

degree[2] $(k \geq 1)$ of a node v is defined as $\delta_k(v) = \{(\ell_{v,w}, \delta_{k-1}(w)) \mid w \in N(v)\}$. We write $v \sim_\delta^k v'$ if $\delta_k(v) = \delta_k(v')$, and simply $v \sim_\delta v'$ if $v \sim_\delta^k v'$ for all $k \geq 0$. In this case, we say that v and v' are *iterated degree equivalent*. It can be shown (see the full version) that as k increases, the partition induced by $\sim_\delta^k$ gets more refined, and indeed it reaches a fixed point for some $k \leq 2n$ where $n = |X|$. The notion of iterated degree is strikingly relevant in our work as it captures what it means for two processes in a network to be indistinguishable. This implies that no distributed algorithm can differentiate between two iterated degree equivalent nodes, as we illustrate in the following result.

▶ **Proposition 2.2.** *Let $I = (A, X, D, C, \alpha)$ be an instance of* $\mathrm{DCSP}(\Gamma)$ *whose factor graph is not necessarily connected and consider two variables $v, v' \in G_I$. Then, $v \sim_\delta v'$ if and only if any terminating decision algorithm over I outputs the same decision at $\alpha(v)$ and $\alpha(v')$. Furthermore, if $v, v' \in X$ and I is satisfiable, then any terminating search algorithm outputs the same values $\nu(v) = \nu(v')$ at $\alpha(v)$ and $\alpha(v')$.*

The following is a direct consequence of Proposition 2.2. We say that two instances I and I' have the same *iterated degree sequence* if there exists a bijection γ between the nodes of G_I and the nodes of $G_{I'}$ such that for every $k \geq 0$ and every node v of G_I, the k^{th} degree of v in I is equal to the k^{th} degree of $\gamma(v)$ in I'. We note that in this case, if we construct the (disconnected) instance $I \cup I'$ containing all the variables and constraints in I and I', then $v \sim_\delta \gamma(v)$ for every node $v \in G_I$. Hence the result below follows.

▶ **Corollary 2.3.** *Let $I, I' \in \mathrm{DCSP}(\Gamma)$ have the same iterated degree sequence. Then with both inputs any terminating decision algorithm will report the same decision.*

Polymorphisms

Let R be a k-ary relation over a finite domain D. An ℓ-ary *polymorphism* of R is an operation $f : D^\ell \to D$ such that the coordinate-wise application of f to any set of ℓ tuples from R gives a tuple in R. More precisely, for any $\mathbf{t}_1, \ldots, \mathbf{t}_\ell \in R$, we have that $(f(\mathbf{t}_1[1], \ldots, \mathbf{t}_\ell[1]), \ldots, f(\mathbf{t}_1[k], \ldots, \mathbf{t}_\ell[k])) \in R$. We say that a function f is a polymorphism of a constraint language Γ if f is a polymorphism of all relations $R \in \Gamma$. Equivalently, we say that Γ is *invariant* under f. The set of polymorphisms of a constraint language Γ will be denoted by $\mathsf{Pol}(\Gamma)$. There is a particular construction of a CSP instance that is closely related to the clone of polymorphisms of the corresponding constraint language. Let Γ be a constraint language over a finite domain D. For any positive integer r, the *indicator problem of order r* for Γ is the instance $I = (X, D, C) \in \mathrm{CSP}(\Gamma)$ where $X = D^r$ and C contains for every relation $R \in \Gamma$ and for every $\mathbf{t}_1, \ldots, \mathbf{t}_r \in R$, the constraint $(\mathbf{s}, R)$ where $\mathbf{s}[i] = (\mathbf{t}_1[i], \ldots, \mathbf{t}_r[i])$ for every $i \in \{1, \ldots, arity(R)\}$. It follows easily that for every $\nu : D^r \to D$, ν satisfies I if and only if ν is a polymorphism of Γ.

An ℓ-ary operation f is said to be *symmetric* if for all $x_1, \ldots, x_\ell$ and for all permutations σ of $\{1, \ldots, \ell\}$ we have that $f(x_1, \ldots, x_\ell) = f(x_{\sigma(1)}, \ldots, x_{\sigma(\ell)})$.

▶ **Example 2.4.** Consider the Boolean relation $R = \{(0, 1), (1, 0)\}$. It is easy to see that the ternary minority operation f given by $f(x, y, z) = x \oplus y \oplus z$ is a polymorphism of R. On the other hand, one can show that R does not have symmetric polymorphisms of arity 2. In particular, let $\mathbf{t}_1 = (0, 1)$ and $\mathbf{t}_2 = (1, 0)$. Since a symmetric binary operation f needs to satisfy $f(0, 1) = f(1, 0)$, the coordinate-wise application of f to $\mathbf{t}_1, \mathbf{t}_2$ would yield a reflexive tuple, which cannot possibly belong to R. ⌟

---

[2] We remark that the notions of degree and iterated degree are well-defined concepts in graph theory. We borrow this terminology to refer to the analogous concepts in CSPs.

Our work unveils a novel structure in the space of solutions of a CSP instance that is deeply connected to the symmetry of its polymorphisms. In particular, $\mathsf{Pol}(\Gamma)$ containing symmetric polymorphisms of all arities is equivalent to the existence of a satisfying assignment to every satisfiable instance of $\mathrm{CSP}(\Gamma)$ that preserves the partition induced by $\sim_\delta$. This is the main result of the next section.

3   Basic Linear Programming relaxation

For any CSP instance $I = (X, D, C)$ there is a LP relaxation (usually called *basic LP relaxation*, see for example [18]) denoted $\mathrm{BLP}(I)$, which is defined as follows. It has a variable $v(x, d)$ for each $x \in X$ and $d \in D$, and a variable $v(c, \mathbf{t})$ for each $c \in C$ and $\mathbf{t} \in R$ where R is the constraint relation of c. All variables must take values in the range $[0, 1]$. The value of $v(x, d)$ is interpreted as the probability that v is assigned to d. Similarly, the value of $v(c, \mathbf{t})$ is interpreted as the probability that the scope of c is assigned component-wise to the tuple $\mathbf{t}$. In this paper we only deal with a feasibility problem (that is, there is no objective function). The variables are restricted by the following equations:

$$\sum_{d \in D} v(x, d) = 1 \quad \text{for all } x \in X \tag{1}$$

$$\sum_{\substack{\mathbf{t} \in R_c \\ \mathbf{t}[i] = d}} v(c, \mathbf{t}) - v(\mathbf{s}_c[i], d) = 0 \quad \text{for all } c \in C, \text{ all } i \in \{1, \dots, arity(c)\}, \text{ and all } d \in D \tag{2}$$

where we denote the relation and scope of a constraint c by R_c and $\mathbf{s}_c$ respectively. We say that BLP decides $\mathrm{CSP}(\Gamma)$ if for every instance $I \in \mathrm{CSP}(\Gamma)$, I is satisfiable whenever $\mathrm{BLP}(I)$ is feasible. We will use the following well-known result.

▶ **Theorem 3.1** (see [18]). *If Γ has symmetric polymorphisms of all arities, then BLP decides* $\mathrm{CSP}(\Gamma)$*. Moreover, if $I \in \mathrm{CSP}(\Gamma)$ is satisfiable then it has a solution ν such that for all x, x' with $v(x, d) = v(x', d)$ for all $d \in D$, we have $\nu(x) = \nu(x')$.*

The following theorem reveals a useful structure inside the solutions of the BLP.

▶ **Theorem 3.2.** *Let $I = (X, D, C)$ be an instance of $\mathrm{CSP}(\Gamma)$ such that $\mathrm{BLP}(I)$ is feasible. Then, $\mathrm{BLP}(I)$ has a feasible solution such that for every $x, x' \in X$ with $x \sim_\delta x'$ and every $d \in D$, $v(x, d) = v(x', d)$.*

Proof (Sketch). We start by rewriting the program in the form

$$\exists ? \mathbf{v} \in [0, 1]^V \quad B\mathbf{v} \geq \mathbf{b} \tag{3}$$

by replacing every equality $a = b$ by the inequalities $a \geq b$ and $-a \geq -b$.

Let us use W and V to denote the rows and columns of B respectively. The main idea of the proof is to apply the Multiplicative Weight Update (MWU) algorithm, a well-known technique that is widely used in optimisation and machine learning. MWU was discovered independently by researchers of different communities; for a survey of its different variants we refer the reader to [2]. The version that is relevant to our paper is described in Algorithm 1. Assuming that a feasible solution to (3) does exist, the algorithm only requires the existence of an oracle which, given a probability W-vector $\mathbf{p}$ (i.e, a non-negative vector $\mathbf{p}$ such that the sum of all its entries is 1), outputs a vector $\mathbf{v}$ which is a solution to the weaker problem

$$\exists ? \mathbf{v} \in [0, 1]^V \quad \mathbf{p}^T B \mathbf{v} \geq \mathbf{p}^T \mathbf{b} \tag{4}$$

if one exists, or correctly states that no such vectors exist otherwise.

■ **Algorithm 1** Multiplicative Weight Update.

Initialisation: Fix $\eta \leq \frac{1}{2}$, let $\rho = \max_{[0,1]^V} \max_{w \in W} |B_w \mathbf{v} - \mathbf{b}[w]|$, and let $\mathbf{w}^{(1)}$ be a W-vector, whose entries, called weights, are initially set to 1.

for $t = 1, \ldots, T$ **do**

> Compute the probability vector $\mathbf{p}^{(t)} = \frac{1}{\Phi(t)} \mathbf{w}^{(t)}$, where $\Phi(t) = \sum_{j=1}^{|W|} \mathbf{w}^{(t)}[j]$
>
> Let $\mathbf{v}^{(t)}$ be a solution satisfying $(\mathbf{p}^{(t)})^T B \mathbf{v}^{(t)} \geq (\mathbf{p}^{(t)})^T \mathbf{b}$ given by oracle $\mathbf{O}$
>
> Compute the losses $\boldsymbol{\ell}^{(t)} = \frac{1}{\rho}(B\mathbf{v}^{(t)} - \mathbf{b})$
>
> Compute the new weights $\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)}(1 - \eta \boldsymbol{\ell}^{(t)})$

end

return $\mathbf{v} := \frac{1}{T} \sum_{t=1}^{T} \mathbf{v}^{(t)}$

Under some technical conditions that provide an upper bound on the number of rounds T necessary to achieve a given approximation (see full version for details) it follows that when $T \to \infty$ MWU converges to a solution of BLP(I). Now consider an oracle $\mathbf{O}$ that, given a W-vector $\mathbf{p}$, returns the V-vector $\mathbf{v}$ where for every $v \in V$, $\mathbf{v}[v] = 1$ if $\mathbf{p}^T B[v]$ is positive and $\mathbf{v}[v] = 0$ otherwise. Since $\mathbf{v}$ maximizes $\mathbf{p}^T B\mathbf{v}$ under the restriction $\mathbf{v} \in [0,1]^V$ it follows that $\mathbf{v}$ satisfies (4).

We note that $\sim_\delta$ induces an equivalence relation on the variables of BLP(I) (namely, $v(x,d)$ is equivalent to $v(x',d')$ whenever $x \sim_\delta x'$ and $d = d'$) which can be extended to an equivalence relation $\sim_V$ on the set V of columns in B. Similarly, $\sim_\delta$ induces an equivalence relation $\sim_W$ on the rows W of B in a natural way. Then our goal is to show that the positions of $\sim_V$-equivalent entries in the output $\mathbf{v} := \frac{1}{T} \sum_{t=1}^{T} \mathbf{v}^{(t)}$ are identical. This is done by showing by induction the more general fact that at each iteration t of the algorithm, the positions of all $\sim_V$-equivalent entries in $\mathbf{v}^{(t)}$ are identical, and that for each of the W-vectors ($\mathbf{w}^{(t)}$, $\mathbf{p}^{(t)}$, and $\boldsymbol{\ell}^{(t)}$) the positions of all $\sim_W$-equivalent entries are identical as well. ◄

We finalize the section by presenting the theorem on the structure of the solution space of CSP instances.

▶ **Theorem 3.3.** *Let Γ be a finite constraint language. The following are equivalent:*

1. *Γ has symmetric polymorphisms of all arities.*
2. *For all satisfiable instances $I = (X, D, C) \in \text{CSP}(\Gamma)$ there exists a satisfying assignment $\nu : X \to D$ such that for all pairs of variables $x, x' \in X$, if $x \sim_\delta x'$ then $\nu(x) = \nu(x')$.*

Proof. (1) $\Rightarrow$ (2). Let I be a satisfiable instance of CSP(Γ), where Γ has symmetric polymorphisms of all arities. Consider the solution of BLP(I) given by Theorem 3.2 and note that it satisfies $v(x,d) = v(x',d)$ for all $x \sim_\delta x'$ and all $d \in D$. Then, by Theorem 3.1, I has a solution ν which satisfies $\nu(x) = \nu(x')$ for all $x \sim_\delta x'$.

(2) $\Rightarrow$ (1). Let Γ satisfy (2) and let $r \geq 1$. We shall prove that Γ has a symmetric polymorphism of arity r. Let $I = (X, D, C)$ be the indicator problem of order r. Recall that every solution to I corresponds to an r-ary polymorphism of Γ, and hence the indicator problem is always satisfiable since for instance the projection to the first coordinate is a polymorphism of Γ. Let ν be a solution of the indicator problem which satisfies condition (2). It is easy to show by induction that for every tuple $(t_1, \ldots, t_r) \in D^r$, every permutation σ of $\{1, \ldots, r\}$ and every $k \geq 0$, $(t_1, \ldots, t_r) \sim_\sigma^k (t_{\sigma(1)}, \ldots, t_{\sigma(r)})$ which implies that $\nu(t_1, \ldots, t_r) = \nu(t_{\sigma(1)}, \ldots, t_{\sigma(r)})$. We conclude that ν is symmetric as required. ◄

4   The Complexity of DCSP

The primary goal of this section is to prove the main theorem of this paper, namely, the dichotomy theorem for tractability of DCSP(Γ), which we now state.

▶ **Theorem 4.1.** DCSP(Γ) *is solvable in polynomial time if and only if* Pol(Γ) *contains symmetric polymorphisms of all arities. Otherwise,* DCSP(Γ) *cannot be solved in finite time.*

We show hardness of constraint languages that do not have symmetric polymorphisms of all arities in Section 4.1 and tractability of the remaining languages in Section 4.2. In addition, in Section 4.3 we extend the decision algorithm so that, additionally, it also provides a solution to the search problem. Hence we have:

▶ **Theorem 4.2.** DCSP-Search(Γ) *is solvable in polynomial time if and only if* Pol(Γ) *contains symmetric polymorphisms of all arities. Otherwise,* DCSP-Search(Γ) *cannot be solved in finite time.*

4.1   Intractable Languages

In this section we focus on intractable languages, that is, the hardness part of Theorem 4.1.

▶ **Theorem 4.3.** *Let* Γ *be a constraint language on a finite domain* D*. If* Pol(Γ) *does not contain symmetric operations of all arities, then there is no algorithm that solves* DCSP(Γ) *in finite time.*

Schematically, the proof goes as follows. Assume that Γ does not have symmetric polymorphisms of some arity r. Consider the relation U defined by the set of solutions of the indicator problem of order r. It can be shown that if DCSP(Γ) is solvable in polynomial (or finite) time then so is DCSP($\{U\}$). Then, we show that there always exist two instances of DCSP($\{U\}$), one which is satisfiable and the other one which is not, that have the same iterated degree sequence. Therefore, any algorithm will return the same output on both instances, meaning that one of these outputs is wrong. Before embarking on the proof we state the following useful combinatorial lemma.

▶ **Lemma 4.4.** *Let* $0 < k < d$ *be positive integers. If* n *is a large enough multiple of* k*, then there exists a collection* $\mathbb{S}$ *of* n^k k-element subsets of $\{0, 1, \ldots, kn-1\}$ *satisfying the following properties:*
(a) $\mathbb{S}$ *contains every* k-element subset of $\{0, \ldots, d-1\}$
(b) *Every element of* $\{0, 1, \ldots, kn - 1\}$ *appears in the same number of sets of* $\mathbb{S}$*.*

Proof of Theorem 4.3. Assume that Pol(Γ) does not contain symmetric polymorphisms of arity r. Fix any arbitrary order $\mathbf{t}_1, \ldots, \mathbf{t}_{|D|^r}$ on the tuples of D^r and consider the relation U defined as

$$\{(f(\mathbf{t}_1), \ldots, f(\mathbf{t}_{|D|^r})) \mid f \text{ is a polymorphism of } \Gamma \text{ of arity } r\}$$

This is, U is the set of solutions of the indicator problem of order r. It follows easily (see full version) that if DCSP($\{U\}$) is not solvable in finite time then neither is DCSP(Γ). In particular, this follows from an adaptation of standard complexity reductions, given that U is pp-definable from Γ without using equality.

Partition D^r into equivalence classes where two tuples $\mathbf{t}, \mathbf{t}' \in D^r$ are related, denoted $\mathbf{t} \equiv \mathbf{t}'$, if there exists some permutation σ on $\{1, \ldots, r\}$ such that $\mathbf{t}'[i] = \mathbf{t}[\sigma(i)]$ for every $i \in \{1, \ldots, r\}$. We shall use $D^r_{\equiv}$ to refer to the collection of classes and $[\mathbf{t}]_{\equiv}$ to refer to the

class of tuple $\mathbf{t}$. For every $\mathbf{t} \in D^r$, define $k_{[\mathbf{t}]_\equiv}$ to be the number of tuples in $[\mathbf{t}]_\equiv$. Then we can choose an integer n large enough such that for every $\mathbf{t} \in D^r$, n is a multiple of $k_{[\mathbf{t}]_\equiv}$, and n satisfies Lemma 4.4 for $k = k_{[\mathbf{t}]_\equiv}$ and $d = k_{[\mathbf{t}]_\equiv} \cdot |D|$.

We are now ready to construct two instances I_1 and I_2 of DCSP($\{U\}$), which are indistinguishable with respect to their iterated degree sequence, but differ with regards to satisfiability. The two instances have the same set of variables, defined to be $\bigcup_{[\mathbf{t}]_\equiv \in D^r_\equiv} V_{[\mathbf{t}]_\equiv}$ where $V_{[\mathbf{t}]_\equiv} = \{v^i_{[\mathbf{t}]_\equiv} \mid 0 \leq i < k_{[\mathbf{t}]_\equiv} n\}$ is a set of $k_{[\mathbf{t}]_\equiv} n$ distinct variables.

We start by constructing the constraints of the unsatisfiable instance I_1, which we will do in two stages. First, for every class $[\mathbf{t}]_\equiv$, let $\mathbb{S}_{[\mathbf{t}]_\equiv}$ be the collection of $n^{k_{[\mathbf{t}]_\equiv}}$ sets of cardinality $k_{[\mathbf{t}]_\equiv}$ given by Lemma 4.4, as before with $d = k_{[\mathbf{t}]_\equiv} \cdot |D|$ and $k = k_{[\mathbf{t}]_\equiv}$. Note that each set in $\mathbb{S}_{[\mathbf{t}]_\equiv}$ defines naturally a subset of $V_{[\mathbf{t}]_\equiv}$ so we shall abuse notation and assume that $\mathbb{S}_{[\mathbf{t}]_\equiv}$ is a collection of subsets of $V_{[\mathbf{t}]_\equiv}$.

To simplify notation it will be convenient to use $\mathbb{S}$ as a shorthand for the indexed family $\{\mathbb{S}_{[\mathbf{t}]_\equiv} \mid [\mathbf{t}]_\equiv \in D^r_\equiv\}$. Now let S be $\{S_{[\mathbf{t}]_\equiv} \mid [\mathbf{t}]_\equiv \in D^r_\equiv\}$ satisfying $S_{[\mathbf{t}]_\equiv} \in \mathbb{S}_{[\mathbf{t}]_\equiv}$ for every $[\mathbf{t}]_\equiv \in D^r_\equiv$. We associate to S the constraint $(\mathbf{s}, U)$ where the scope $\mathbf{s}$ is constructed as follows. Before defining $\mathbf{s}$ we need some preparation. Recall that every coordinate of U, and hence of $\mathbf{s}$, is associated to a tuple $\mathbf{t} \in D^r$, so we can talk of the class $[\mathbf{t}]_\equiv$ to which each coordinate belongs. In particular, there are $k_{[\mathbf{t}]_\equiv}$ coordinates in $\mathbf{s}$ of class $[\mathbf{t}]_\equiv$. Hence, by fixing some arbitrary ordering we can use $\mathbf{s}^i_{[\mathbf{t}]_\equiv}$, $i = 1, \ldots, k_{[\mathbf{t}]_\equiv}$ to refer to the coordinates in $\mathbf{s}$ of class $[\mathbf{t}]_\equiv$. Then, informally, $S_{[\mathbf{t}]_\equiv}$ describes which variables from $v^0_{[\mathbf{t}]_\equiv}, \ldots, v^{k_{[\mathbf{t}]_\equiv} n - 1}_{[\mathbf{t}]_\equiv}$ to use in order to fill coordinates $\mathbf{s}^i_{[\mathbf{t}]_\equiv}$, $i = 1, \ldots, k_{[\mathbf{t}]_\equiv}$. Formally, for every $[\mathbf{t}]_\equiv \in D^r_\equiv$ and each $i = 1, \ldots, k_{[\mathbf{t}]_\equiv}$, $\mathbf{s}^i_{[\mathbf{t}]_\equiv}$ is assigned to the i^{th} element in $S_{[\mathbf{t}]_\equiv}$ in increasing order.

We add such a constraint for each of the $\Pi_{[\mathbf{t}]_\equiv \in D^r_\equiv} n^{k_{[\mathbf{t}]_\equiv}} = n^{(|D|^r)}$ possible choices for S. Therefore, after the first stage we have exactly $n^{(|D|^r)}$ constraints.

In the second stage we add more constraints which will yield the particular symmetry of I_1. Note that every permutation σ on $\{1, \ldots, r\}$ induces a permutation σ' on the coordinates of U in a natural way. Specifically, if coordinate i of U is associated to tuple $\mathbf{t}_i$, then $\sigma'(i) = j$ where $\mathbf{t}_j = (\mathbf{t}_i[\sigma(1)], \ldots, \mathbf{t}_i[\sigma(r)])$. Then, in the second stage, for each permutation σ on $\{1, \ldots, r\}$ and for every constraint $(\mathbf{s}, U)$ added in the first stage we add the constraint $(\mathbf{s}', U)$ where for every $1 \leq i \leq |D|^r$, $\mathbf{s}'[i] = \mathbf{s}[\sigma'(i)]$. Therefore, after the second stage we have a total of $m = r! \cdot n^{(|D|^r)}$ constraints as needed.

We now turn to I_2. The constraints are constructed in a similar way, but instead of using the family $\mathbb{S}$ in the first stage, we use a different family $\mathbb{S}'$. In particular, for each class $[\mathbf{t}]_\equiv$, $\mathbb{S}'_{[\mathbf{t}]_\equiv}$ is obtained by partitioning $V_{[\mathbf{t}]_\equiv}$ in $k_{[\mathbf{t}]_\equiv}$ blocks of consecutive elements, so that each block has exactly n elements. Then, $\mathbb{S}'_{[\mathbf{t}]_\equiv}$ contains the $n^{k_{[\mathbf{t}]_\equiv}}$ sets that can be obtained by selecting one element from each block. The second stage is done exactly as in I_1.

▷ Claim 4.5. I_1 and I_2 have the same iterated degree sequence.

Proof. Let $[\mathbf{t}]_\equiv \in D^r_\equiv$. First, we observe that in both instances after the first stage, every variable of $V_{[\mathbf{t}]_\equiv}$ appears in the same number of constraints. More specifically, every variable in $V_{[\mathbf{t}]_\equiv}$ appears in an n-fraction of the constraints added in stage 1. In the case of instance I_1 this is due to the fact that $\mathbb{S}_{[\mathbf{t}]_\equiv}$ satisfies condition (b) in Lemma 4.4 and in instance I_2 this follows from the fact that $\mathbb{S}'_{[\mathbf{t}]_\equiv}$ contains all possible sets obtained by choosing an element within each one of the blocks of size n. After the second stage (in both I_1 and I_2 since the second stage is common) every variable in $V_{[\mathbf{t}]_\equiv}$ still participates in an n-fraction of the total number of constraints. In addition, it follows easily that the positions of the scope in which a variable in $V_{[\mathbf{t}]_\equiv}$ participates distribute evenly among the $k_{[\mathbf{t}]_\equiv}$ positions associated to $\mathbf{t}$. That is, in both instances, we have that for every $[\mathbf{t}]_\equiv \in D^r_\equiv$, every variable $x \in V_{[\mathbf{t}]_\equiv}$, and

every position i associated to $[\mathbf{t}]_{\equiv}$ there are exactly $\frac{m}{nk_{[\mathbf{t}]_{\equiv}}}$ constraints in which x appears at position i of the scope, where $m = r! \cdot n^{|D|^r}$. Using this fact it is very easy to prove that I_1 and I_2 have the same iterated degree sequence. Formally, one could show by induction on k that for every $[\mathbf{t}]_{\equiv} \in D_{\underline{\equiv}}^r$ and $x_1, x_2 \in V_{[\mathbf{t}]_{\equiv}}$, $\delta_k^{I_1}(x_1) = \delta_k^{I_2}(x_2)$ and that for any two constraints c_1, c_2 in I_1 and I_2 respectively $\delta_k^{I_1}(c_1) = \delta_k^{I_2}(c_2)$. Here we are using $\delta_k^{I_1}(\cdot)$ and $\delta_k^{I_2}(\cdot)$ to denote the k^{th} degree of a node in the factor graphs of I_1 and I_2 respectively.   ◁

▷ **Claim 4.6.**   Instance I_1 is unsatisfiable while instance I_2 is satisfiable.

Proof. We start by showing that I_1 is not satisfiable. Assume by contradiction that I_1 has a satisfying assignment ν. For each class $[\mathbf{t}]_{\equiv}$, consider the values given by ν to the first d variables $v_0, \ldots, v_{d-1}$ in $V_{[\mathbf{t}]_{\equiv}}$. Since $d = k_{[\mathbf{t}]_{\equiv}} \cdot |D|$, it follows by the pigeon-hole principle that at least $k_{[\mathbf{t}]_{\equiv}}$ of these variables are assigned by ν to the same value of D. Let $S_{[\mathbf{t}]_{\equiv}}$ be a subset of $V_{[\mathbf{t}]_{\equiv}}$ containing $k_{[\mathbf{t}]_{\equiv}}$ of these variables (we know that this subset belongs to $\mathbb{S}_{[\mathbf{t}]_{\equiv}}$ by condition (a) of Lemma 4.4). Now consider the constraint $(\mathbf{s}, U)$ in I_1 associated to $S := \{S_{[\mathbf{t}]_{\equiv}} \mid [\mathbf{t}]_{\equiv} \in D_{\underline{\equiv}}^r\}$, which belongs to I_1. If ν is a solution to I_1, then the restriction of ν to $\mathbf{s}$ corresponds to an r-ary polymorphism of Γ. But ν assigns the same value to any two related tuples $\mathbf{t} \equiv \mathbf{t}'$, which implies that ν is symmetric, a contradiction.

We now turn our focus to I_2. Let f be any r-ary polymorphism of Γ (for example the i^{th} ($1 \leq i \leq r$) projection operation defined as $f(x_1, \ldots, x_r) = x_i$). We shall construct a solution ν of I_2 in the following way. Recall that in the definition of I_2 we have partitioned the tuples of $V_{[\mathbf{t}]_{\equiv}}$ in $k_{[\mathbf{t}]_{\equiv}}$ consecutive blocks. In the first stage, all the elements in each block are placed in the same coordinate of U. So, if $\mathbf{t}_1, \ldots, \mathbf{t}_{|D|^r}$ are the tuples associated to coordinates $1, \ldots, |D|^r$ and hence block $1, \ldots, |D|^r$ respectively, then we only need that all variables in the i^{th} block are assigned to $f(\mathbf{t}_i)$ to satisfy all constraints added in the first stage. This assignment also satisfies the constraints added in the second stage, because if f is an r-ary polymorphism of Γ, then for every permutation σ on $\{1, \ldots, r\}$, the operation $g(x_1, \ldots, x_r)$ defined as $f(x_{\sigma(1)}, \ldots, x_{\sigma(r)})$ is also a polymorphism of Γ.   ◁

To sum up, we constructed two instances I_1 and I_2, the latter of which is satisfiable while the former is not, which have the same iterated degree sequence. It follows from Corollary 2.3 that any distributed algorithm will give the same output on both instances, meaning that no algorithm can solve DCSP($\{U\}$). As anticipated at the beginning of the proof then it follows that there are also no algorithms that solve DCSP(Γ).   ◀

4.2   Tractable Languages

In this section we turn our attention to the tractable case. In particular we shall show the following:

▶ **Theorem 4.7.** *Let Γ be a constraint language that is invariant under symmetric polymorphisms of all arities. Then there is an algorithm* Alg *that solves* DCSP(Γ). *The total running time, number of rounds, and maximum message size of* Alg *are, respectively, $\mathcal{O}(n^3 m \log n)$, $\mathcal{O}(n^2)$, and $\mathcal{O}(m \log n)$ where n and m are the number of variables and constraints, respectively, of the input instance.*

Note that this implies the "if" part of Theorem 4.1. Alg is composed of two phases. In the first phase, a distributed version of the colour refinement algorithm allows every process to calculate its iterated degree. Then, thanks to Theorem 3.3 we can use the degree of a variable as its ID for the second phase, implying that a distributed adapted version of the *jpq*-consistency algorithm [16] where messages are tagged with a process' iterated degree solves the decision problem for Γ.

Distributed Colour Refinement

Let $I = (A, X, D, C, \alpha)$ be an instance of DCSP(Γ) and let $n = |X|$ and $m = |C|$. There is a very natural way to calculate an agent's iterated degree in a distributed way, both for variables and for constraints. This is a mere adaptation of the 1-dimensional Weisfeiler-Leman algorithm, also known as colour refinement, an algorithm that partitions the vertices of a graph into classes by iteratively distinguishing them on the basis of their degree (see for example [15, 14]). The algorithm proceeds in rounds. At round $k = 0$, each agent $\alpha(v)$ for $v \in X \cup C$ computes $\delta_0(v)$ and broadcasts it to all its neighbours. At round $k > 0$, each agent $\alpha(v)$ knows the $(k-1)^{th}$ degrees of its neighbours which it had received in the previous round, uses them to compute $\delta_k(v)$, and broadcasts it to its neighbours. If $k = 2n$ then for every $x, x' \in X$ satisfying $x \sim_\delta^k x'$ we have that $x \sim_\delta x'$, which implies that we can essentially regard the k^{th} iterated degree as the unique common ID for all variables that are iterated degree equivalent. Then in $2n$ rounds each agent $\alpha(v)$ can compute $\delta_\infty(v)$, where we use δ_∞ as a shorthand for δ_{2n}. As we described it, the distributed colour refinement algorithm is not particularly efficient in terms of message complexity. Although this is not necessary to achieve polynomial time, we can reduce the space required to encode $\delta_\infty(v)$.

▶ **Lemma 4.8.** *Let $s_{\max}$ denote the size of the encoding of $\delta_\infty(v)$. A modified version of the distributed colour refinement algorithm that runs over $\mathcal{O}(n^2)$ rounds achieves $s_{\max} = \mathcal{O}(\log n)$. The time at each round and the maximum size of a message are both bounded above by $\mathcal{O}(ms_{max})$.*

As we will see, the price of an increase in the number of rounds (from n to n^2) is compensated by the effect of $s_{\max}$ on both time complexity and the size of the messages.

The Distributed Consistency Algorithm

It is well known that if a constraint language Γ has symmetric operations of all arities then it satisfies the so-called *bounded width* property (see [4]). We avoid introducing the definition of bounded width as it is not needed in our results and refer the reader to [4] for reference. Then, it has been shown in [16] that if Γ has bounded width and $I \in$ CSP(Γ) satisfies a combinatorial condition called *jpq-consistency*, then I has a solution. Instead of stating literally the result in [16] we shall state a weaker version that uses a different notion of consistency, more suitable to the model of distributed computation introduced in the paper.

A *set system* S is a subset of $X \times D$. We shall use S_x to denote the set $\{d \in D \mid (x, d) \in S\}$. A walk of length ℓ (in instance I) is any sequence $x_0 c_0 \ldots c_{\ell-1} x_\ell$ where $x_0, \ldots, x_\ell$ are variables, $c_0, \ldots, c_{\ell-1}$ are constraints, and $x_i, x_{i+1} \in c_i$ for every $0 \leq i < \ell$. Note that walks are precisely the walks in the factor graph G_I (in the standard graph-theoretic sense) starting and finishing in X.

Let S be a set system, p be a walk, and $B \subseteq S_x$ where x is the starting node of p. The *propagation* of B via p under S, denoted $B +_S p$, is the subset of D defined inductively on the length ℓ of p as follows. If $\ell = 0$ then $B +_S p = B$. Otherwise, $p = p' c_{\ell-1} x_\ell$ where p' is a path of length $\ell - 1$ ending at $x_{\ell-1}$. Let $c_{\ell-1} = (\mathbf{s}, R)$. Then we define $B +_S p$ to contain all $e \in D$ such that there exists $d \in B +_S p'$ and $\mathbf{t} \in R$ such that for every $1 \leq i \leq arity(R)$, $\mathbf{t}[i]$ satisfies the following conditions:
1. $\mathbf{t}[i] \in S_{\mathbf{s}[i]}$,
2. if $\mathbf{s}[i] = x_{\ell-1}$ then $\mathbf{t}[i] = d$, and
3. if $\mathbf{s}[i] = x_\ell$ then $\mathbf{t}[i] = e$.

We are now ready to state the result from [16] that we shall use.

▶ **Theorem 4.9** (follows from [16]). *Let I be an instance of $\mathrm{CSP}(\Gamma)$ where Γ has bounded width and let S be a set system such that $S_x \neq \emptyset$ for every $x \in X$ and such that for every walk p starting and finishing at the same node x and for every $d \in S_x$, d belongs to $\{d\} +_S p$. Then I is satisfiable.*

Our goal is to design a distributed algorithm that either correctly determines that an instance I is unsatisfiable, or produces a set system S verifying the conditions of Theorem 4.9. This is not possible in general due to the fact that agents are anonymous and hence a hypothetical algorithm that would generate a walk in a distributed way would be unable to determine if the initial and end nodes are the same. However, thanks to the structure established by Theorem 3.3, this difficulty can be overcome when Γ has symmetric polymorphisms of all arities because, essentially, the iterated degree of a node can act as its unique identifier. To make this intuition precise we will need to introduce a few more definitions.

We say that a pair $(x, d) \in S$ is *S-supported* if for every walk p starting at x and finishing at a node y with $x \sim_\delta y$, we have that $\{d\} +_S p$ contains d.

▶ Remark 4.10. We note that if $(x, d) \in S$ is not S-supported and $p = x_0 c_0 \ldots x_\ell$ is a walk of minimal length among all walks witnessing that (x, d) is not S-supported then $\ell \leq n2^{|D|}$. Indeed if we let $B_i = \{d\} + x_0 c_0 \ldots x_i$, $i = 0, \ldots, \ell$ then we have that $(x_i, B_i) \neq (x_j, B_j)$ for every $0 \leq i < j \leq \ell$, since otherwise the shorter walk $x_0 c_0, \ldots, x_i, c_j, \ldots, x_\ell$ would contradict the minimality of p. Since there are n choices for each x_i and $2^{|D|}$ choices for B_i, the bound follows.   ⌟

We say that a set system S is *safe* if for every solution $\nu \in I$ we have

$$\nu(x) = \nu(y) \text{ for all } x, y \in X \text{ with } x \sim_\delta y \implies \nu(x) \in S_x \text{ for all } x \in X.$$

Then, we have

▶ **Lemma 4.11.** *Let S be a safe set system and let $(x, d) \in S$ be a pair that is not S-supported. Then $S \setminus \{(x, d)\}$ is safe.*

Our distributed consistency algorithm (that is, the second phase of Alg) works as follows. Every variable agent $\alpha(x)$ maintains a set $S_x \subseteq D$ in such a way that the set system S is guaranteed to be safe at all times. As a result of an iterative process S is modified. We shall use S^i to denote the content of S at the i^{th} iteration, where an iteration is, in turn, a loop of $T = 2n(2^{|D|} + 1) = \mathcal{O}(n)$ consecutive rounds. The rationale behind this exact value will be made clear later. Initially, S_x^0 is set to D for every $x \in X$. At iteration i for $i \geq 1$, S^i is obtained by removing all the elements in S^{i-1} that are not S^{i-1}-supported. Then, in at most $n|D| = \mathcal{O}(n)$ iterations we shall obtain a fixed point S^∞.

The key observation is that when Γ has symmetric polymorphisms of all arities, the satisfiability of I can be determined from S^∞. Indeed, if $S_x^\infty = \emptyset$ for some $x \in X$ then we can conclude from the fact that S^∞ is safe and Theorem 3.3 that I has no solution. Otherwise, S^∞ satisfies the conditions of Theorem 4.9 and, hence, I is satisfiable.

It remains to see how to compute S^{i+1} from S^i. In an initial preparation step for every iteration, every variable agent $\alpha(x)$ sends S_x^i to all its neighbours. To compute S^{i+1} the algorithm proceeds in rounds. All the messages sent are sets containing triplets of the form (δ_∞, d, B) where $d \in D$, $B \subseteq D$, and δ_∞ is the iterated degree of some variable $x \in X$. It follows from the fact that there are at most n possibilities for the degree of a variable that the size of each message is $\mathcal{O}(ns_{\max})$.

The agents controlling variables and constraints alternate. That is, variables perform internal and send events at even rounds and receive messages at odd rounds, while constraints perform internal and send events at odd rounds and receive messages at even rounds. More

specifically, in round $j = 0$ of iteration i, every variable agent $\alpha(x)$ sends to its neighbours the message M containing all triplets of the form $(\delta_\infty(x), d, \{d\})$ with $d \in S_x^i$. At round $2j$ for $j > 0$, $\alpha(x)$ computes $M = M_1 \cup \cdots \cup M_r$ where $M_1, \ldots, M_r$ are the messages it received at the end of round $2j - 1$. Subsequently, for every triplet $(\delta_\infty, d, B) \in M$ with $\delta_\infty = \delta_\infty(x)$ and $d \notin B$, $\alpha(x)$ marks d as "not S^i-supported". Finally, it sends message M to all its neighbours. This computation can be done in time $\mathcal{O}(rns_{\max}) = \mathcal{O}(mns_{\max})$ provided that each message is stored as an ordered array.

In round $2j + 1$, every constraint agent $\alpha(c)$ computes from the messages M_x (received from each neighbour $\alpha(x)$ in the previous round) the set M_x', which contains for every variable $y \in c$ and every (δ_∞, d, B) in M_y, the triplet $(\delta_\infty, d, B +_{S^i} p)$ where $p = y, c, x$. Finally, it sends to each neighbour $\alpha(x)$ the corresponding message M_x'. Note that while $\alpha(c)$ doesn't know the address of $\alpha(x)$ specifically, knowing the label of the channel that connects them is sufficient to calculate M_x' correctly and send the message accordingly. Moreover, for given y and x, $\alpha(c)$ can compute $B +_{S^i} p$ in $\mathcal{O}(1)$ time as $\alpha(c)$ knows both S_y^i and S_x^i. Hence, since the arity of the relations is fixed (as Γ is fixed) the total running time at iteration $2j + 1$ of a constraint agent $\alpha(c)$ is $\mathcal{O}(ns_{\max})$.

Now it is immediate to show by induction that for every $j \geq 0$, every $x \in X$ and $c \in C$ with $x \in c$ the message sent by $\alpha(x)$ to $\alpha(c)$ at the end of round $2j$ is precisely

$$\{(\delta_\infty(y), d, \{d\} + p) \mid y \in X, d \in S_y^i, p \text{ is a walk of length } j \text{ of the form } p = y, \ldots, x\}$$

and the message sent by $\alpha(c)$ to $\alpha(x)$ at the end of round $2j + 1$ is precisely

$$\{(\delta_\infty(y), d, \{d\} + p) \mid y \in X, d \in S_y^i, p \text{ is a walk of length } j + 1 \text{ of the form } p = y, \ldots, c, x\}.$$

By Remark 4.10 only $2n2^{|D|} = T - 2n$ iterations are needed to identify all elements in S^i that are not S^i-supported. Hence, after exactly $T - 2n$ rounds every variable agent $\alpha(x)$ computes S_x^{i+1} by removing all the elements in S^i that are marked as "not S^i-supported". If $S_x^{i+1} = \emptyset$, then $\alpha(x)$ initiates a wave, which is propagated by all its neighbours, broadcasting that an inconsistency was detected. In this case, in at most $2n$ additional rounds all agents can correctly declare that I is unsatisfiable. Otherwise, a new iteration begins.

To sum up, the distributed consistency algorithm consists of $\mathcal{O}(n)$ iterations consisting, each, of $\mathcal{O}(n)$ rounds where the total running time for internal events at a given round is $\mathcal{O}(mns_{\max})$ and the maximum size of each message transmitted is $\mathcal{O}(ns_{\max})$. Together with the bounds given by Lemma 4.8 for the distributed colour refinement phase, this completes the proof of Theorem 4.7.

4.3 The Search Algorithm

We conclude by presenting the proof of Theorem 4.2. The hardness part follows immediately from Theorem 4.1 as the search problem is as difficult as the decision problem. For the positive result we shall present an adaptation of the algorithm solving the decision version. Let I be an instance of DCSP-Search(Γ) where Γ contains symmetric polymorphisms of all arities. In what follows we shall use intensively the fact that $\mathsf{Pol}(\Gamma)$ is closed under composition. Let $J \subseteq D$ be minimal with the property that $f(D) = J$ for some unary polymorphism f in $\mathsf{Pol}(\Gamma)$. It is fairly standard to show that for every $r \geq 0$ there is a r-ary symmetric operation g such that $g(x, \ldots, x) = x$ for every $x \in J$. Indeed, let f satisfy $f(D) = J$ and let g be any r-ary symmetric polymorphism in $\mathsf{Pol}(\Gamma)$. Then the unary operation h defined by $h(x) = f \circ g(x, \ldots, x)$ is a unary polymorphism of Γ. By the choice of f we have $h(D) \subseteq J$. We note that $h(J) = J$ since otherwise h^2 would contradict the

minimality of f. Consequently, h^{-1} belongs to $\mathsf{Pol}(\Gamma)$ and, hence, the r-ary operation defined as $h^{-1} \circ f \circ g$ satisfies the claim. This implies that if we enlarge the constraint language by adding all singletons $\{d\}$, $d \in J$, the resulting constraint language, which we shall denote by Γ', still has symmetric polymorphisms of all arities. For convenience we also include D in Γ'.

The algorithm has two phases. In the first phase it runs the decision algorithm to determine whether the instance is satisfiable. As a byproduct, every variable agent $\alpha(x)$ has computed its iterated degree $\delta_\infty(x)$ and knows as well its rank in a prescribed ordering of all variable degrees $\delta_\infty^1, \ldots, \delta_\infty^r$, $r \leq n$. This (partial) order will be used to coordinate between the agents. An i-agent, $1 \leq i \leq r$ is any agent $\alpha(x)$ with $\delta_\infty(x) = \delta_\infty^i$. We also assume a fixed ordering on the elements in D. If the instance is unsatisfiable nothing else remains to be done so from now on we shall assume that the instance is satisfiable.

In the second phase the algorithm searches for a solution. Every variable agent $\alpha(x)$ maintains a set $F_x \subseteq D$ with the property that there is a solution ν that *falls within F*, i.e, such that $\nu(x) \in F_x$ for every $x \in X$. Initially every agent $\alpha(x)$ sets $F_x = D$ so it is only necessary to make sure that this condition is preserved during the execution of the algorithm. The second phase contains two nested loops. The outer loop has r iterations and the inner loop consists of at most $|D|$ iterations so that we shall use iteration (i, d) to indicate the run of the algorithm at the $i = 1, \ldots, r$ iteration of the outer loop and at the iteration d of the inner loop.

At the beginning of iteration (i, d) every variable agent $\alpha(x)$ defines $S_x \subseteq D$ to be $S_x = \{d\}$ whenever $\alpha(x)$ is an i-agent and $S_x = F_x$ elsewhere. Then it runs the distributed consistency algorithm starting at S obtaining a fixed point S^∞. We note that since all initial sets S_x belong to Γ' and Γ' contains symmetric polymorphisms of all arities then the obtained fixed point S^∞ correctly determines whether there exists a solution ν that falls within S. Then every i-agent $\alpha(x)$ checks whether $S_x^\infty = \emptyset$. In case of positive answer nothing else is done and round (i, d) finishes. Otherwise, $\alpha(x)$ sets F_x to $\{d\}$ and starts a wave to indicate to all processes that the i^{th} iteration of the outer loop is finished and that the next iteration of the outer loop can start. When the r iterations of the outer loop have been completed the set system F contains only singletons. The assignment that sets every variable $x \in X$ to the only element in F_x is necessarily a solution. This concludes the proof of Theorem 4.2.

5 Conclusion

We analysed the complexity of the distributed constraint satisfaction problem on a synchronous, anonymous network parametrised by the constraint language. We showed that, depending on the polymorphisms of Γ, DCSP(Γ) is either solvable in polynomial time, or not solvable altogether. A number of natural open questions arise in this context. For instance, it is not clear whether asynchronous networks are strictly more powerful than their synchronous counterpart. Moreover, it would be interesting to explore the role of allowing agents to make random choices - provided this is not used to create and share unique IDs.

In the spirit of [13], one could consider characterizing the structural restrictions on tractable distributed CSPs, or in other words, determining which *classes of networks* are tractable in the DCSP framework, regardless of the constraint language. The starting point for this analysis could be the work on fibrations by Boldi et al. (see for example [7, 8]). In particular, we propose the question of establishing a connection between the universal fibration of a graph and its iterated degree sequence.

References

1. Dana Angluin. Local and global properties in networks of processors (extended abstract). In Raymond E. Miller, Seymour Ginsburg, Walter A. Burkhard, and Richard J. Lipton, editors, *Proceedings of the 12th Annual ACM Symposium on Theory of Computing, April 28-30, 1980, Los Angeles, California, USA*, pages 82–93. ACM, 1980. `doi:10.1145/800141.804655`.

2. Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory Comput.*, 8(1):121–164, 2012. `doi:10.4086/toc.2012.v008a006`.

3. Pablo Barceló, Egor V. Kostylev, Mikaël Monet, Jorge Pérez, Juan L. Reutter, and Juan Pablo Silva. The logical expressiveness of graph neural networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL: `https://openreview.net/forum?id=r1lZ7AEKvB`.

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

5. Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinícius Flores Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Çaglar Gülçehre, H. Francis Song, Andrew J. Ballard, Justin Gilmer, George E. Dahl, Ashish Vaswani, Kelsey R. Allen, Charles Nash, Victoria Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet Kohli, Matthew Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases, deep learning, and graph networks. *CoRR*, abs/1806.01261, 2018. `arXiv:1806.01261`.

6. Ramon Bejar, Bhaskar Krishnamachari, Carla Gomes, and Bart Selman. Distributed constraint satisfaction in a wireless sensor tracking system. In *Workshop on Distributed Constraint Reasoning, International Joint Conference on Artificial Intelligence*, volume 4, 2001.

7. Paolo Boldi, Shella Shammah, Sebastiano Vigna, Bruno Codenotti, Peter Gemmell, and Janos Simon. Symmetry breaking in anonymous networks: Characterizations. In *Fourth Israel Symposium on Theory of Computing and Systems, ISTCS 1996, Jerusalem, Israel, June 10-12, 1996, Proceedings*, pages 16–26. IEEE Computer Society, 1996.

8. Paolo Boldi and Sebastiano Vigna. An effective characterization of computability in anonymous networks. In Jennifer L. Welch, editor, *Distributed Computing, 15th International Conference, DISC 2001, Lisbon, Portugal, October 3-5, 2001, Proceedings*, volume 2180 of *Lecture Notes in Computer Science*, pages 33–47. Springer, 2001. `doi:10.1007/3-540-45414-4_3`.

9. Andrei A. Bulatov. A dichotomy theorem for nonuniform csps. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 319–330. IEEE Computer Society, 2017. `doi:10.1109/FOCS.2017.37`.

10. Ken R. Duffy, Charles Bordenave, and Douglas J. Leith. Decentralized constraint satisfaction. *IEEE/ACM Trans. Netw.*, 21(4):1298–1308, 2013. `doi:10.1109/TNET.2012.2222923`.

11. Ferdinando Fioretto, Enrico Pontelli, and William Yeoh. Distributed constraint optimization problems and applications: A survey. *J. Artif. Intell. Res.*, 61:623–698, 2018. `doi:10.1613/jair.5565`.

12. Wan Fokkink. *Distributed algorithms: an intuitive approach*. MIT Press, 2013.

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

14. Martin Grohe. word2vec, node2vec, graph2vec, x2vec: Towards a theory of vector embeddings of structured data. In Dan Suciu, Yufei Tao, and Zhewei Wei, editors, *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2020, Portland, OR, USA, June 14-19, 2020*, pages 1–16. ACM, 2020. `doi:10.1145/3375395.3387641`.

15   Martin Grohe, Kristian Kersting, Martin Mladenov, and Pascal Schweitzer. Color refinement and its applications. *Van den Broeck, G.; Kersting, K.; Natarajan, S*, 2017.

16   Marcin Kozik. Solving CSPs using weak local consistency. *SIAM Journal on Computing, to appear*, 2020. URL: `https://marcinkozik.staff.tcs.uj.edu.pl/Solving.CSPs.using.weak.local.consistency.pdf`.

17   Andrei A. Krokhin and Stanislav Živný, editors.   *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. URL: `http://www.dagstuhl.de/dagpub/978-3-95977-003-3`.

18   Gábor Kun, Ryan O'Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 csps, and robust satisfaction. In Shafi Goldwasser, editor, *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 484–495. ACM, 2012. `doi:10.1145/2090236.2090274`.

19   Amnon Meisels.   *Distributed Search by Constrained Agents - Algorithms, Performance, Communication.* Advanced Information and Knowledge Processing. Springer, 2008. `doi:10.1007/978-1-84800-040-7`.

20   Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 4602–4609. AAAI Press, 2019. `doi:10.1609/aaai.v33i01.33014602`.

21   Francesca Rossi, Peter van Beek, and Toby Walsh, editors.   *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*.   Elsevier, 2006.   URL: `http://www.sciencedirect.com/science/bookseries/15746526/2`.

22   Jan Toenshoff, Martin Ritzert, Hinrikus Wolf, and Martin Grohe. RUN-CSP: unsupervised learning of message passing networks for binary constraint satisfaction problems. *CoRR*, abs/1909.08387, 2019. `arXiv:1909.08387`.

23   Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL: `https://openreview.net/forum?id=ryGs6iA5Km`.

24   Masafumi Yamashita and Tiko Kameda. Computing on an anonymous network. In Danny Dolev, editor, *Proceedings of the Seventh Annual ACM Symposium on Principles of Distributed Computing, Toronto, Ontario, Canada, August 15-17, 1988*, pages 117–130. ACM, 1988. `doi:10.1145/62546.62568`.

25   Makoto Yokoo, Edmund H. Durfee, Toru Ishida, and Kazuhiro Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Proceedings of the 12th International Conference on Distributed Computing Systems, Yokohama, Japan, June 9-12, 1992*, pages 614–621. IEEE Computer Society, 1992. `doi:10.1109/ICDCS.1992.235101`.

26   Makoto Yokoo, Edmund H. Durfee, Toru Ishida, and Kazuhiro Kuwabara. The distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Trans. Knowl. Data Eng.*, 10(5):673–685, 1998. `doi:10.1109/69.729707`.

27   Makoto Yokoo and Katsutoshi Hirayama. Algorithms for distributed constraint satisfaction: A review. *Auton. Agents Multi Agent Syst.*, 3(2):185–207, 2000. `doi:10.1023/A:1010078712316`.

28   Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 331–342. IEEE Computer Society, 2017. `doi:10.1109/FOCS.2017.38`.