



Constraint satisfaction problems: Convexity makes AllDifferent constraints tractable[☆]

Michael Fellows^a, Tobias Friedrich^{b,*}, Danny Hermelin^c, Nina Narodytska^d,
Frances Rosamond^a

^a Charles Darwin University, Darwin, Northern Territory, Australia

^b Friedrich-Schiller-Universität Jena, Germany

^c Max-Planck-Institut für Informatik, Saarbrücken, Germany

^d NICTA and University of New South Wales, Sydney, Australia

ARTICLE INFO

Article history:

Received 10 March 2012

Received in revised form 16 November 2012

Accepted 30 November 2012

Communicated by O. Watanabe

ABSTRACT

We examine the complexity of constraint satisfaction problems that consist of a set of AllDiff constraints. Such CSPs naturally model a wide range of real-world and combinatorial problems, like scheduling, frequency allocations, and graph coloring problems. As this problem is known to be NP-complete, we investigate under which further assumptions it becomes tractable. We observe that a crucial property seems to be the convexity of the variable domains and constraints. Our main contribution is an extensive study of the complexity of Multiple AllDiff CSPs for a set of natural parameters, like maximum domain size and maximum size of the constraint scopes. We show that, depending on the parameter, convexity can make the problem tractable even though it is provably intractable in general. Interestingly, the convexity of constraints is the key property in achieving fixed parameter tractability, while the convexity of domains does not usually make the problem easier.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Constraint satisfaction is a general framework that allows structure preserving encoding of many real world problems, including formal verification, vehicle routing and scheduling. A CONSTRAINT SATISFACTION PROBLEM (CSP) consists of a set of variables that must be assigned values from their respective domains in such a way that a set of constraints is satisfied. In general, finding a solution of a CSP is an NP-complete problem. Hence, much research has been devoted to finding restricted classes of CSPs that admit polynomial time algorithms. This area of research includes two main directions. The first direction exploits structural properties of relations between variables and constraints of a CSP instance [12,18]. The second direction investigates limitations on the types of constraints in the CSP constraint language, so-called *tractable constraint languages* [23,41].

In this work we continue this line of research and investigate CSPs under both structural restrictions and limitations on constraint types. Such combinations allow us to identify interesting classes of CSPs that have expressive constraint languages and are still tractable under some natural conditions on relations between variables and constraints in the framework of parameterized complexity.

[☆] A preliminary conference version [16] appeared in the 22nd International Joint Conferences on Artificial Intelligence (IJCAI 2011). This work was done while all five authors were visiting the University of Newcastle, Australia.

* Corresponding author. Tel.: +49 3641 946321; fax: +49 3641 946322.

E-mail address: friedrich@uni-jena.de (T. Friedrich).

We restrict the type of allowed constraints to the most common constraint, called the AllDiff constraint [29,38]. The AllDiff constraint requires a set of variables to take pairwise distinct values. We call a CSP with the language that consists of only AllDiff constraints the MULTIPLE ALLDIFF CONSTRAINT SATISFACTION PROBLEM (MAD-CSP). Finding a solution of the MAD-CSP is NP-complete due to a straightforward reduction from a graph coloring problem.

Our structural limitations of MAD-CSPs are inspired by a large body of research on global constraints in constraint programming. We consider a class of MAD-CSPs that possess the *convexity of constraints* property. This property means that the variables are linearly ordered so that all AllDiff constraints are defined over intervals of variables in this order. Such MAD-CSPs can be seen as a generalization of the SLIDE constraint where AllDiff is the relation to slide [3]. We also investigate another restriction on CSPs that comes from work on global constraints over variables with interval domains that we call *convex domains*. This means that every domain constitutes an interval of values. It is interesting to consider the convex domain restriction in the context of MAD-CSPs, as it has proved useful for constructing efficient inference algorithms for global constraints. Many global constraints, like NValue [2] or overlapping AllDiff [4], that are NP-hard to reason about become polynomially solvable if the domains of the variables are intervals of values. We further say that a MAD-CSP instance is *biconvex* if it has both convex domains and convex constraints.

To motivate the AllDiff constraint and our notion of convexity, consider the following examples of real-world and combinatorial problems that can be naturally encoded as MAD-CSP problems:

- The time table of a workshop on CSP might be modeled with one variable per speaker with the domain of the variables corresponding to the availability of the speakers [37,45]. For example, speakers A and B might be available during the first and third sessions and speaker C during all sessions. We introduce three variables S_1 , S_2 and S_3 with domains $D(S_1) = \{1, 3\}$, $D(S_2) = \{1, 3\}$, and $D(S_3) = \{1, 2, 3\}$. One (convex) AllDiff constraint containing all variables encodes that there are no two talks at the same time. If the available times of the speakers are intervals, this MAD-CSP instance is biconvex.
- Scheduling of n jobs with given start and end times on m machines can be modeled as a MAD-CSP with one variable per job and values corresponding to machines on which the job can be run. Each job can be performed by a subset of machines. Suppose we have 3 jobs with the following time intervals (ordered by their starting points and then by the end points): J_1 has to be performed in the interval $[1, 2]$ and can be assigned to one of the machines $\{1, 3\}$, J_2 in the interval $[2, 4]$ on one of the machines $\{1, 3\}$, and J_3 in the interval $[3, 5]$ on one of the machines $\{2, 3\}$. The domains of variables J_i are as follows: $D(J_1) = D(J_2) = \{1, 3\}$ and $D(J_3) = \{2, 3\}$. Convex AllDiff constraints then ensure that each machine only runs one job per time slot: AllDiff $[J_1, J_2]$ and AllDiff $[J_2, J_3]$. Our results demonstrate that this scheduling problem is FPT parameterized by the total number of machines.
- The classical graph coloring problem where two adjacent vertices should get different colors can be described as a MAD-CSP with one variable per vertex and one AllDiff constraint per edge. If there is no restriction on the set of colors used, the domains are convex. If the graph is an interval graph, also the constraints are convex (see Theorem 4.1).
- Finding frequency allocations can be modeled with one variable per cell, and domains corresponding to the available frequencies. The AllDiff constraints enforce that nearby cells get different frequencies to avoid distortions. The domains here are convex as well.
- Solving a Sudoku of size $N \times N$ can be modeled with one variable per cell and one constraint per row, column, and box. The domain of variable is either all possible numbers or the respective clue of the cell. Hence this MAD-CSP also has convex domains.
- The n -queens problem can be modeled as MAD-CSP as described in [34].

It should be noted that finding a solution of a MAD-CSP with convex domains and constraints is still an NP-complete problem due to a reduction from the list coloring of an interval graph problem (Section 4). Hence, we use the more fine-scaled framework of parameterized complexity [14,17,36] to analyze its complexity.

1.1. Our results

We examine MAD-CSPs in several dimensions [16]. First, the domains and/or constraints can be convex as observed in the examples above. Second, there are various natural parameters which might be fixed or bounded for typical problem instances. We explore several commonly used parameters of the problem to identify whether bounding these parameters leads to a fixed parameter tractable MAD-CSP: the number of variables, the number of values (universe size), the maximum size of the constraint scopes (arity), and the maximum domain size. From the theoretical point of view, our parameterizations allow us to reveal the core of the hardness of this problem. From a practical point of view, they can be used to construct efficient filtering algorithms for the Multiple AllDiff constraint. For more information about filtering algorithms for global constraints we refer the reader to [46]. Note that even though some of these parameters, like the number of variables or the number of values, may not typically be bounded in some applications, during the backtracking search of a constraint solver these parameters become bounded due to the solver's branching and inference mechanisms (see e.g. [42,44]). Hence efficient algorithms for bounded parameters are of high practical relevance.

Table 1
Overview of our results.

Parameter	n = number of variables	m = size of the universe	k = maximum size of constraint scopes	ℓ = maximum domain size
No convexity	FPT (Theorem 3.1)	para-NPC (Lemma 2.1)	para-NPC (Lemma 2.1)	para-NPC (Lemma 2.1)
Convex domains	FPT (Theorem 3.1)	para-NPC (Lemma 2.1)	para-NPC (Lemma 2.1)	para-NPC (Lemma 2.1)
Convex constraints	FPT (Theorem 3.1)	FPT (Theorem 4.5)	FPT (Theorem 6.2)	para-NPC (Lemma 5.1)
Bi-convex	FPT (Theorem 3.1)	FPT (Theorem 4.5)	FPT (Theorem 6.2)	Open ¹

Our results for the four chosen parameters, depending on the convexity of the domains and constraints, are presented in Table 1. For the first parameter (number of variables), we give an FPT algorithm independent of the convexity (Theorem 3.1). For this case we can also show that there is a polynomial-sized problem kernel (Theorem 3.3). For the second and third parameters (size of the universe and maximum size of the constraint scopes), we give an FPT algorithm if the constraints are convex (Theorems 4.5 and 6.2). For most of the remaining cases we can prove that there is no FPT algorithm unless $P = NP$ (Lemmas 2.1 and 5.1). More precisely, we show that these cases are para-NPC, that is, NP-complete for some (small) fixed parameter. There is one case left open. This is the case of bi-convex MAD-CSP instances when parameterized by maximum domain size. We conjecture that this case is in FPT, but we can prove this only for the aggregate parameter domain size and number of constraints (Theorem 5.5). See Section 5 for details.

1.2. Related work

The AllDiff constraint first appeared in the ALICE constraint programming language [29]. The complexity of a single AllDiff constraint is examined in several papers by different authors [30,31,35,37,38,45]. Recently, much research has been devoted to constructing efficient algorithms for conjunctions of AllDiff constraints [4,27,28]. Such algorithms offer strong inference power and can prevent a constraint solver from diving into exponentially large unsatisfiable search sub-trees [25,47]. Unfortunately, in most cases the conjunction of global constraints leads to a global constraint that is NP-hard. This is also the case for the AllDiff constraint: Even for the conjunction of two AllDiff constraints, a polynomial algorithm does not exist unless $P = NP$ [22,27]. This motivates a search for tractable subclasses [4,27,28]. Bessiere et al. [4] give a polynomial-time algorithm for a conjunction of two AllDiff constraints if the domains are convex. We extend this to multiple AllDiff constraints and explore different parameters and restrictions for which the problem becomes tractable.

In the context of general CSPs, much research has been done on tractable constraint satisfaction problems. One line of this research deals with tractable constraint languages (e.g. [9,10,23,41]); another exploits structural properties of CSPs [12,15,18,32,33]. Marx [32,33] investigated general CSPs from the parameterized complexity point of view, and Samer and Szeider [40] identified relevant parameters in general constraint satisfaction problems.

It should be noted that the problem that we consider is a special case of the Same-Relation constraint studied in [24]. Jefferson et al. [24] consider a binary constraint satisfaction problem with an identical binary relation over each edge of the primal constraint graph and a set of unary constraints. They showed that such binary CSPs are intractable for many structures of constraint graphs (cf. Definition 2.1), like cliques, bipartite graphs, grids, or directed acyclic graphs. In our work, we continue this research and identify additional restrictions on the Same-Relation constraint that make this constraint fixed parameter tractable.

In [1,11,26] a similar INTERVAL CONSTRAINED COLORING problem was studied. The objective of this problem is to assign a color to a set of variables such that a set of constraints is satisfied. Each constraint is a convex set of variables which specifies exactly the number of variables that belong to each color class. Despite the similarity, this problem is neither a generalization nor a specialization of MAD-CSPs.

1.3. Organization of the paper

We first formally define CSP and MAD-CSP in Section 2. In Sections 3 and 4 we examine the cases of fixed number of variables and values, respectively. Section 5 discusses fixed maximum domain size. In the last section, we study the fixed maximum size of the constraint scopes. Note that although each section covers a different parameter, the proofs of Section 6 build on results from Section 5, which in turn depend on Section 4.

¹ We know that this case is FPT for several aggregate parameters like domain size and the treewidth of the constraint graph (Lemma 5.2), domain size and the treewidth of the domain graph (Lemma 5.3), and domain size and the number of constraints (Theorem 5.5), but for the single parameter domain size it remains open. See Section 5 for details.

2. Formal model

2.1. General CSP

Formally, a CSP instance is a triple $(\mathcal{X}, \mathcal{U}, \mathcal{C})$, where $\mathcal{X} := \{x_1, \dots, x_n\}$ is a set of variables, $\mathcal{U} \subseteq \mathbb{N}$ is a set of possible values (also called the universe), and \mathcal{C} is a set of constraints. A constraint C is defined over a set of variables, called the *scope* of the constraint C and denoted $\text{scope}(C) \subseteq \mathcal{X}$. A constraint C can be represented extensionally as a table of valid (or invalid) assignments to variables in its scope, or intensionally by giving an expression or a formula involving the variables in the constraint scope [43]. We consider only constraints represented intensionally. Each constraint in \mathcal{C} is a triplet (S, R, m) , where S is an m -tuple of variables (constraint scope) and R is an arbitrary m -ary relation over \mathcal{U} (constraint relation). We will slightly abuse notation by writing $x \in C$ to indicate that variable $x \in \mathcal{X}$ is in the scope S of some constraint $C \in \mathcal{C}$. Also, we assume w.l.o.g. that every variable occurs in at least one constraint scope, and every domain element occurs in at least one constraint relation. A solution to a CSP instance is a function (assignment) τ from the set of variables \mathcal{X} to the set of values \mathcal{U} satisfying that for each constraint (S, R, m) with $S = (x_{i_1}, x_{i_2}, \dots, x_{i_m})$ the tuple $(\tau(x_{i_1}), \tau(x_{i_2}), \dots, \tau(x_{i_m}))$ is a member of R . If there is such an assignment we say that the CSP instance is *satisfiable*, and otherwise we say it is *unsatisfiable*.

2.2. MAD-CSP

We will focus on CSP instances which include only AllDiff constraints, henceforth referred to as MULTIPLE ALLDIFF CSP (MAD-CSP) instances. A MAD-CSP instance is a quadruple $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$, where \mathcal{X} and \mathcal{U} are defined as above, \mathcal{D} is a function assigning domains to variables, and \mathcal{C} is the set of AllDiff constraints. The function \mathcal{D} is defined from \mathcal{X} to $2^{\mathcal{U}}$, restricting variable $x \in \mathcal{X}$ to have values only from $\mathcal{D}(x)$. The set \mathcal{C} is a set of subsets of $2^{\mathcal{X}}$ enforcing the requirement that all variables occurring together in one constraint must be assigned different values. We assume w.l.o.g. that all MAD-CSP instances are in a “canonic” form, that is, the universe is $\mathcal{U} = \{1, 2, \dots, |\mathcal{U}|\}$; there are no values in the definitions of $\mathcal{D}(x)$ which are not used by any variable x ; there is no variable with an empty domain; there are no empty constraints; and there is no constraint which is a proper subset of another constraint. A solution to a MAD-CSP instance is a function $\tau: \mathcal{X} \rightarrow \mathcal{U}$ such that:

- (i) $\forall x \in \mathcal{X} : \tau(x) \in \mathcal{D}(x)$.
- (ii) $\forall C \in \mathcal{C}$ and $\forall x, y \in C : x \neq y \Rightarrow \tau(x) \neq \tau(y)$.

Specifically, we are interested in domains and constraints with a particular structure. We say that a domain $\mathcal{D}(x)$ of a variable x is *convex*, if $\mathcal{D}(x) = \{i, i+1, i+2, \dots, j\}$ for some $1 \leq i \leq j \leq m$. Analogously, a constraint $C \in \mathcal{C}$ is called *convex*, if $C = \{x_i, x_{i+1}, x_{i+2}, \dots, x_j\}$ for some $1 \leq i \leq j \leq n$. A MAD-CSP instance has *convex domains (constraints)* if all domains (constraints) are convex, and it is *bi-convex* if it has convex domains and constraints simultaneously.

Our work focuses on analyzing MAD-CSP instances under the framework of parameterized complexity. Readers are referred to [14,17,36] for relevant concepts and definitions. We consider the following parameters for MAD-CSP:

- Parameter $n := |\mathcal{X}|$ measuring the number of variables.
- Parameter $m := |\mathcal{U}|$ measuring the size of the universe.
- Parameter $k := \max_{C \in \mathcal{C}} |C|$ measuring the maximum size of the constraint scopes.
- Parameter $\ell := \max_{x \in \mathcal{X}} |\mathcal{D}(x)|$ measuring the maximum domain size.

It is not difficult to see that MAD-CSP instances can naturally model the GRAPH COLORING problem, where we are given a graph $G := (V, E)$, and an integer k , and the goal is to find a function (k -coloring) $f: V \rightarrow \{1, \dots, k\}$ with $f(u) \neq f(v)$ for all $\{u, v\} \in E$. To reduce GRAPH COLORING to MAD-CSP, we let $\mathcal{X} = V$, $\mathcal{U} = \{1, \dots, k\}$, $\mathcal{D}(x) = \mathcal{U}$ for all $x \in \mathcal{X}$, and $\mathcal{C} = E$. Then the MAD-CSP instance is satisfiable iff G has a k -coloring. Since GRAPH COLORING is known to be NP-complete for instances with $k \geq 3$ [19], we have the following hardness result for MAD-CSP.

Lemma 2.1. *MAD-CSP is NP-complete even when restricted to convex domain instances with $m = 3$, $k = 2$, and $\ell = 3$.*

We conclude this section by introducing the notion of a constraint graph, a concept playing an important role in many works relating to CSP [21,33,40]. After presenting a formal definition of this notion, we state an important theorem by Gottlob et al. [21] relating the fixed-parameter tractability of CSP instances to a structural parameter of the constraint graph, namely the *treewidth* parameter (see e.g. [13] for a formal definition).

Definition 2.1. The *constraint graph* of a CSP instance $(\mathcal{X}, \mathcal{U}, \mathcal{C})$ is the underlying graph of the hypergraph $(\mathcal{X}, \mathcal{C})$. That is, it is the graph with vertex set \mathcal{X} , where two variables $x, y \in \mathcal{X}$ are connected by an edge iff there is a constraint $C \in \mathcal{C}$ with $x, y \in C$.

Theorem 2.2 ([12]). *CSP can be solved in $m^{t+1} \cdot n^{O(1)}$ time and space, where t is the treewidth of the constraint graph.*

3. Fixed number of variables

We begin our discussion by considering the instances of MAD-CSP with a fixed number of variables n . We will present two parameterized algorithms for this parameter, both with asymptotically similar dependencies on the parameter, and both working even when neither the domains nor the constraints are convex. Note that general CSP is not FPT when parameterized by n , as for example k -Clique can be expressed as a CSP with k variables [40]. The first of the two algorithms is given in the theorem below.

Theorem 3.1. *MAD-CSP can be solved in $O(n! nm)$ time, independent of the underlying convexity assumption.*

Proof. We examine the following algorithm:

1. Try all permutations π of $\{1, 2, \dots, n\}$.
2. For i from 1 to n , assign $x_{\pi(i)}$ the smallest valid value from $\mathcal{D}(x_{\pi(i)})$.
3. If this is possible for all variables, return yes.
4. If there is no permutation for which a valid assignment is found, return no.

We can easily bound the worst-case runtime of this algorithm by $O(n! nm)$ which is FPT for the parameter being the number of variables n . It remains to prove correctness of the algorithm. It is clear that if the algorithm returns yes, it has found a valid assignment and the answer is correct. To show that it is also correct if it returns no, it is sufficient to prove the following:

- (1) If the given instance is satisfiable, then there is a permutation π of $\{1, 2, \dots, n\}$ such that the algorithm finds a valid assignment.
In order to prove claim (1), we first show the following claim.
- (2) If the given instance is satisfiable, then there is an assignment such that there is a variable x which is assigned to $\min \mathcal{D}(x)$.

To show claim (2), assume the contrary, that is, there is a satisfiable instance where in all valid assignments no variable x is assigned to $\min \mathcal{D}(x)$. In this case, take any valid assignment and observe that the value 1 cannot have been assigned to any variable. By the definition of our problem, (cf. Section 2), there must be a variable x with value 1 in its domain. Reassigning x to 1 gives a valid assignment in which x is assigned to $\min \mathcal{D}(x)$. This proves claim (2).

It remains to show claim (1). Given a satisfiable instance, we use claim (2) and choose the first variable in the permutation π to be a variable x which is assigned to $\min \mathcal{D}(x)$ in a valid assignment. To choose the second variable in the permutation, we reduce the given instance by the already fixed variable x . More precisely, we remove the variable x , the value $\min \mathcal{D}(x)$, and all values which are solely used by x from the instance. This new instance must also be satisfiable and claim (2) gives the second variable in the required permutation for claim (1). Continuing in the same way for the third, fourth, etc. variable, gives the desired permutation and proves claim (1). The theorem is thus proven. \square

Our second algorithm shows that the MAD-CSP problem is not only FPT when parameterized by the number of variables, it actually also has a polynomial kernel in this parameter (see e.g. [6] for a discussion on this important concept in parameterized complexity). Our kernelization algorithm is based on a simple reduction rule, which states, roughly speaking, that if a variable has a large enough domain it can be deleted. The justification of this rule is given in the following simple lemma.

Lemma 3.2. *Let $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ be a MAD-CSP instance and $x \in \mathcal{X}$ be a variable with $|\mathcal{D}(x)| \geq n$. Then $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ is satisfiable iff $(\mathcal{X}', \mathcal{U}', \mathcal{D}', \mathcal{C}')$ is satisfiable, where $\mathcal{X}' := \mathcal{X} \setminus \{x\}$, $\mathcal{U}' := \bigcup_{x \in \mathcal{X}'} \mathcal{D}(x)$, $\mathcal{D}' := \mathcal{D}|_{\mathcal{X}'}$, and $\mathcal{C}' := \{C \setminus \{x\} : C \in \mathcal{C} \text{ and } C \neq \{x\}\}$.*

Proof. If there is an assignment satisfying the instance $(\mathcal{X}', \mathcal{U}', \mathcal{D}', \mathcal{C}')$, then since $|\mathcal{X}'| = |\mathcal{X}| - 1 = n - 1 < |\mathcal{D}(x)|$, there must be a value in $\mathcal{D}(x)$ not assigned to any variable in \mathcal{X}' . We can thus assign this value to x without violating any constraint in \mathcal{C} , and so if $(\mathcal{X}', \mathcal{U}', \mathcal{D}', \mathcal{C}')$ is satisfiable so is $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$. The other direction is trivial.

Using this rule iteratively until it can no longer be applied, we obtain a MAD-CSP instance where each variable has a domain size which is bounded by the total number of remaining variables in the instance. This immediately implies that the problem has a kernel of size $O(n^2)$.

Theorem 3.3. *The MAD-CSP problem has a kernel of size $O(n^2)$, independent of the underlying convexity assumption.*

4. Fixed number of values

In this section we consider the MAD-CSP problem when parameterized by the size of the universe m of the instance. We will show that MAD-CSP with convex constraints is FPT in this parameter. In contrast, as previously shown in Lemma 2.1, MAD-CSP with convex domains is para-NPC.

Before presenting our parameterized algorithm, we show that MAD-CSP with convex constraints is equivalent to the classical LIST COLORING PROPER INTERVAL GRAPHS (LCPIG) problem. A graph $G := (V, E)$ is an *interval graph* if there exists a set S of intervals on the real line and a bijection $I: V \rightarrow S$ such that $\{u, v\} \in E \iff I(u) \cap I(v) \neq \emptyset$ (see e.g. [20]). If no interval is properly contained in another in S , then G is a *proper interval graph*. In LCPIG, we are given a proper interval graph G , where each vertex v has a list of colors $L(v)$, and the goal is to assign a color to each vertex from its list such that adjacent vertices are assigned different colors.

To show that MAD-CSP with convex constraints is equivalent to LCPIG, we use an equivalent definition of proper interval graphs due to Roberts [39], who showed that G is a proper interval graph iff the adjacency matrix of G has the *consecutive 1's property*; that is, if there is an ordering of the vertices of G such that the adjacency matrix of G under this ordering has 1's appearing consecutively in each row and column. Due to the classical Booth and Lueker [8] algorithm, this ordering can be computed in linear time. It is easy to verify that the adjacency matrix of the constraint graph associated with a MAD-CSP instance with convex constraints has the consecutive 1's property. Conversely, any graph whose adjacency matrix has the consecutive 1's property can be interpreted as a constraint graph of a MAD-CSP instance with convex constraints. We therefore obtain the following theorem, which will also be useful in Section 5.

Theorem 4.1. *MAD-CSP with convex constraints is equivalent to the LCPIG problem.*

Now using the fact that LCPIG is NP-complete even in the case where each list is an interval of consecutive colors [5,7], we immediately obtain from Theorem 4.1 the following corollary.

Corollary 4.2. *MAD-CSP is NP-complete even in the bi-convex case.*

We next proceed to describe our parameterized algorithm for the parameter m . We first observe the following simple lemma that follows directly from the definition of the MAD-CSP problem.

Lemma 4.3. *If there is any constraint $C \in \mathcal{C}$ with scope greater than m , then the given MAD-CSP instance cannot be satisfied.*

Proof. Any satisfying assignment of the given MAD-CSP instance must assign all variables in C to different values, which is impossible, since there are $m < |C|$ different values altogether.

We next use Lemma 4.3 above to bound the treewidth of our constraint graph. For this, we will use a slightly different graph invariant called *pathwidth*. The pathwidth of a given graph is always bounded from below by its treewidth, and it is well known that the graph has pathwidth k iff the smallest clique number of an interval graph that contains it as a subgraph has maximum clique-size $k + 1$ (see e.g. [13]). This result together with Lemma 4.3 above and the fact that the constraint graph of a MAD-CSP instance with convex constraints is an interval graph, immediately gives us the following lemma.

Lemma 4.4. *If each constraint scope in a given MAD-CSP instance with convex constraints is of size at most k , then the constraint graph of the instance has pathwidth at most $k - 1$.*

We now can directly use Theorem 2.2 of Gottlob et al. [21] to obtain the main result of this section.

Theorem 4.5. *MAD-CSP with convex constraints can be solved in $m^m \cdot n^{O(1)}$ time.*

5. Fixed maximum domain size

We next consider the parameter $\ell := \max_{x \in X} |\mathcal{D}(x)|$ which measures the maximum number of values in any variable domain. Recall that by Lemma 2.1 we know that MAD-CSP with convex domains is para-NPC for this parameter. The following lemma shows that this also holds for convex constraints.

Lemma 5.1. *MAD-CSP with convex constraints is NP-complete for instances with $\ell \geq 3$.*

Proof. The problem is clearly in NP. To show NP-hardness, we use a result of Jansen [22] who proved that the list coloring problem for graphs that are the union of two cliques is NP-complete, even when the list size of each vertex is at most 3. Since a graph which is a union of two cliques is a proper interval graph with two maximal cliques, this implies the lemma by the equivalence proven in Theorem 4.1. \square

The proof of Lemma 5.1 uses a result of Jansen [22]. It also implies that MAD-CSP with convex constraints is para-NPC even when parameterized by both the maximum domain size and the number of constraints. We next show that for these two parameters, MAD-CSP turns out to be in FPT when the domains are convex. The complexity of the bi-convex case for the single parameter $\ell := \max_{x \in X} |\mathcal{D}(x)|$ remains open. We begin with the following lemma, which we will also need in Section 6. Its proof is based on Theorem 2.2.

Lemma 5.2. *MAD-CSP can be solved in $\ell^{t+1} \cdot n^{O(1)}$, where t is the treewidth of the constraint graph, independent of the underlying convexity assumption.*

Proof. Let $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ be a given instance of MAD-CSP with the treewidth of the constraint graph bounded by some integer t . To prove the theorem, we use [Theorem 2.2](#) by reducing $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ to a general CSP instance $(\mathcal{X}, \mathcal{U}', \mathcal{C}')$ with $|\mathcal{U}'| = \ell$, and where the treewidth of the constraint graph remains bounded by t .

In order to reduce the universe size to ℓ , we change the domain of every variable $x \in \mathcal{X}$ to a subset of $\mathcal{U}' := \{1, 2, \dots, \ell\}$. For every variable $x \in \mathcal{X}$, we choose some bijective mapping δ_x from its original domain $\mathcal{D}(x)$ to its new domain $\mathcal{D}'(x) := \{1, 2, \dots, |\mathcal{D}(x)|\}$. The new domain $\mathcal{D}'(x)$ can be seen as containing only the indices of the entries of the original domain. To adjust the constraints appropriately, recall that every edge in the constraint graph indicates that the two adjacent vertices must be assigned different values. Hence for an edge (x, y) in the constraint graph of $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$, this can be formulated as a constraint $C \in \mathcal{C}'$ in the general CSP, which requires us to assign value $\delta_x(i)$ to the variable x and value $\delta_y(j)$ to the variable y , for $i \in \mathcal{D}(x)$ and $j \in \mathcal{D}(y)$, iff $i \neq j$. Of course, we only add a constraint if $\mathcal{D}(x) \cap \mathcal{D}(y) \neq \emptyset$.

It is clear that every solution to the general CSP $(\mathcal{X}, \mathcal{D}', \mathcal{C}')$ instance corresponds to a solution of the original MAD-CSP instance $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ and vice versa. Furthermore, the constraint graphs of the two instances are isomorphic. Since this reduction can be carried out in polynomial time, applying [Theorem 2.2](#) finishes the proof. \square

We next introduce a graph of MAD-CSP instances which is defined by considering the relationships between the domains of different variables. Following this, we will prove an analogous result to [Lemma 5.2](#) for the case where the treewidth of the so-called domain graph is bounded.

Definition 5.1. The domain graph of a MAD-CSP instance $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ is a graph with vertex set \mathcal{X} where two variables $x, y \in \mathcal{X}$ are adjacent iff $\mathcal{D}(x) \cap \mathcal{D}(y) \neq \emptyset$.

Lemma 5.3. *MAD-CSP can be solved in $\ell^{t+1} \cdot n^{O(1)}$, where t is the treewidth of the domain graph, independent of the underlying convexity assumption.*

Proof. To prove the theorem, we reduce $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ to a general CSP instance $(\mathcal{X}, \mathcal{U}', \mathcal{C}')$ using the same construction as in the proof of [Lemma 5.2](#). It is not difficult to see that the constraint graph of the new instance is a subgraph of the domain graph of the original instance. Since all subgraphs of a graph with treewidth bounded by t also have treewidth bounded by t , the treewidth of the constraint graph of $(\mathcal{X}, \mathcal{U}', \mathcal{C}')$ is also at most t . The theorem therefore follows, again, by applying [Theorem 2.2](#). \square

Our goal now is to use [Lemma 5.3](#) above to show that MAD-CSP is FPT when parameterized by both ℓ and $|\mathcal{C}|$. To do this, we show that unless our given MAD-CSP instance is unsatisfiable, the domain graph of this instance has treewidth bounded by a function of both these parameters. The following lemma gives the first step in this direction.

Lemma 5.4. *Let $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ be an instance of MAD-CSP, and let D be a domain of some variable in \mathcal{X} . If there are more than $\ell \cdot |\mathcal{C}|$ variables $x \in \mathcal{X}$ with $\mathcal{D}(x) = D$, then $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ cannot be satisfied.*

Proof. Let $C \in \mathcal{C}$ be an arbitrary constraint. If there are more than ℓ variables $x \in C$ with $\mathcal{D}(x) = D$ then the given MAD-CSP instance cannot be satisfied, since all these variables must take different values, and $|D| \leq \ell$. Now if there are more than $\ell \cdot |\mathcal{C}|$ variables $x \in \mathcal{X}$ with $\mathcal{D}(x) = D$, by the pigeonhole principle more than ℓ of these belong to same constraint, and the lemma follows.

Theorem 5.5. *MAD-CSP with convex domains can be solved in $\ell^{2c\ell^2} \cdot n^{O(1)}$ time, where $c := |\mathcal{C}|$ is the number of constraints in the given instance.*

Proof. We prove the theorem by applying [Lemma 5.3](#). To do this, we first observe that in the case of convex domains, the domain graph G of a given MAD-CSP instance is an interval graph. This can be seen by assigning each variable its interval domain. Our algorithm begins by checking whether there is a domain D in our instance where the number of variables with this domain is more than $c\ell$. If this is the case, we report that the instance is unsatisfiable by [Lemma 5.4](#). Otherwise, recalling that no domain is contained in another, the number of interval domains which start or end at a certain value is at most $c\ell$. In this case, the number of interval domains which start or end at a certain value s , which is exactly the number of interval domains that are contained inside the interval $[s, s + \ell]$ or $[s - \ell, s]$, is at most $c\ell$. From this it follows that any domain intersects fewer than $2c\ell^2$ other domains, and therefore G has maximum clique size less than $2c\ell^2$. Thus, G has treewidth at most $2c\ell^2 - 1$, and we can apply [Lemma 5.3](#) to complete the proof. \square

6. Fixed maximum size of constraint scope

In this section we examine the parameter k , which measures the maximum size of constraint scopes, that is, $k = \max_{C \in \mathcal{C}} |C|$. [Lemma 2.1](#) showed that MAD-CSP is para-NPC without convexity of the constraints. To show that convex constraints make the problem fixed-parameter tractable with respect to k , we first argue that variables with domain size larger than k^2 can be safely removed.

Lemma 6.1. Let $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ be an instance of MAD-CSP where all constraints $C \in \mathcal{C}$ are convex and all constraint scopes are of size at most k . Also, let $x \in \mathcal{X}$ be a variable with $\mathcal{D}(x) > k^2$. Then $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ is satisfiable iff $(\mathcal{X}', \mathcal{U}', \mathcal{D}', \mathcal{C}')$ is satisfiable, where $\mathcal{X}' := \mathcal{X} \setminus \{x\}$, $\mathcal{U}' := \bigcup_{x \in \mathcal{X}'} \mathcal{D}(x)$, $\mathcal{D}' := \mathcal{D}|_{\mathcal{X}'}$, and $\mathcal{C}' := \{C \setminus \{x\} : C \in \mathcal{C} \text{ and } C \neq \{x\}\}$.

Proof. Since no constraint is a proper subset of another constraint, and since all constraints are convex, we know that for every variable $x \in \mathcal{X}$ there is at most one constraint which has x as its first variable. This in turn implies that a variable can occur only in at most k constraints. Now suppose there is an assignment satisfying the instance $(\mathcal{X}', \mathcal{U}', \mathcal{D}', \mathcal{C}')$. As x is a member of at most k constraints, and each constraint has at most k variables, the total number of disallowed values for x is at most k^2 . Since $\mathcal{D}(x) > k^2$, we can always assign a value to x without violating any constraint in \mathcal{C} , and thus $(\mathcal{X}, \mathcal{U}, \mathcal{D}, \mathcal{C})$ is also satisfiable. The other direction in the lemma is trivial. \square

This reduction rule gives us an FPT algorithm for convex constraints as follows.

Theorem 6.2. MAD-CSP with convex constraints can be solved in $k^{2k} \cdot n^{O(1)}$ time.

Proof. Given a MAD-CSP instance with maximum size of constraint scopes k , we know from Lemma 4.4 that the pathwidth, and therefore also the treewidth, of the constraint graph is at most $k - 1$. By above Lemma 6.1, we can reduce, in polynomial time, the given MAD-CSP instance to a MAD-CSP instance with maximum domain size k^2 . As the treewidth of the constraint graph remains bounded by $k - 1$, applying Lemma 5.2 finishes the proof. \square

7. Conclusions and open problems

We have studied the complexity of MAD-CSPs with structural restrictions under several parameterizations and have shown a number of cases that are fixed-parameter tractable. The most prominent open problem arising directly from this investigation is whether in the bi-convex case, the problem is fixed-parameter tractable when parameterized by the domain size. We conjecture that the problem is FPT, but have only been able to show this for the aggregate parameter: domain size and number of constraints.

More generally, it would be interesting to similarly investigate CSPs for other kinds of structural restrictions, such as almost convex variables and domains, and for other constraints or sets of constraints besides Multiple AllDiff. Because of the wide applicability of CSPs, sometimes these parameterizations are relevant because the parameters may be small for the natural input distributions faced in practice, but the study of the parameterized complexity of these problems is also motivated by the observation that during the backtracking search of a constraint solver, relevant parameters may become small, and thus FPT algorithms may have high practical relevance.

References

- [1] E. Althaus, S. Canzar, K.M. Elbassioni, A. Karrenbauer, J. Mestre, Approximating the interval constrained coloring problem, in: 11th SWAT, 2008, pp. 210–221.
- [2] C. Bessière, E. Hebrard, B. Hnich, Z. Kiziltan, T. Walsh, Filtering algorithms for the NValue constraint, in: 2nd CPAIOR, 2005, pp. 79–93.
- [3] C. Bessière, E. Hebrard, B. Hnich, Z. Kiziltan, T. Walsh, SLIDE: a useful special case of the CARDPATH constraint, in: 18th ECAI, 2008, pp. 475–479.
- [4] C. Bessière, G. Katsirelos, N. Narodytska, C.-G. Quimper, T. Walsh, Propagating conjunctions of AllDifferent constraints, in: 24th AAAI, 2010, pp. 27–32.
- [5] M. Biró, M. Hujter, Z. Tuza, Precoloring extension. I: interval graphs, Discrete Math. 100 (1992) 267–279.
- [6] H.L. Bodlaender, R.G. Downey, M.R. Fellows, D. Hermelin, On problems without polynomial kernels, J. Comput. System Sci. 75 (2009) 423–434.
- [7] F. Bonomo, G. Durán, J. Marengo, Exploring the complexity boundary between coloring and list-coloring, Ann. Oper. Res. 169 (2009) 3–16.
- [8] K.S. Booth, G.S. Lueker, Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms, J. Comput. System Sci. 13 (1976) 335–379.
- [9] A. Bulatov, A. Krokhin, P. Jeavons, The complexity of maximal constraint languages, in: 33rd STOC, 2001, pp. 667–674.
- [10] A.A. Bulatov, A dichotomy theorem for constraints on a three-element set, in: 43rd FOCS, 2002, p. 649.
- [11] J. Byrka, A. Karrenbauer, L. Sanita, The interval constrained 3-coloring problem, in: 9th LATIN, 2010, pp. 591–602.
- [12] R. Dechter, J. Pearl, Tree clustering for constraint networks, Artif. Intell. 38 (1989) 353–366.
- [13] R. Diestel, Graph Theory, Springer-Verlag, 2000.
- [14] R. Downey, M. Fellows, Parameterized Complexity, Springer-Verlag, 1999.
- [15] T. Feder, M.Y. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory, SIAM J. Comput. 28 (1998) 57–104.
- [16] M.R. Fellows, T. Friedrich, D. Hermelin, N. Narodytska, F.A. Rosamond, Constraint satisfaction problems: convexity makes AllDifferent constraints tractable, in: 22nd IJCAI, 2011, pp. 522–527.
- [17] J. Flum, M. Grohe, Parameterized Complexity Theory, Springer-Verlag, 2006.
- [18] E.C. Freuder, A sufficient condition for backtrack-bounded search, J. ACM 32 (1985) 755–761.
- [19] M. Garey, D. Johnson, Computers and intractability: a guide to the theory of NP-completeness, W. H. Freeman, 1979.
- [20] M. Golumbic, Algorithmic Graph Theory and Perfect Graphs, Academic Press, 1980.
- [21] G. Gottlob, F. Scarcello, M. Sideri, Fixed-parameter complexity in AI and nonmonotonic reasoning, Artif. Intell. 138 (2002) 55–86.
- [22] K. Jansen, The optimum cost chromatic partition problem, in: 3rd CIAC, 1997, pp. 25–36.
- [23] P. Jeavons, D.A. Cohen, M. Gyssens, Closure properties of constraints, J. ACM 44 (1997) 527–548.
- [24] C. Jefferson, S. Kadioglu, K.E. Petrie, M. Sellmann, S. Zivný, Same-relation constraints, in: 15th CP, 2009, pp. 470–485.
- [25] G. Katsirelos, N. Narodytska, T. Walsh, Combining symmetry breaking and global constraints, in: 13th CSCLP, 2008, pp. 84–98.
- [26] C. Komusiewicz, R. Niedermeier, J. Uhlmann, Deconstructing intractability: a case study for interval constrained coloring, in: 20th CPM, 2009, pp. 207–220.
- [27] M. Kutz, K.M. Elbassioni, I. Katriel, M. Mahajan, Simultaneous matchings: hardness and approximation, J. Comput. System Sci. 74 (2008) 884–897.
- [28] F. Lardeux, E. Monfroy, F. Saubion, B. Crawford, C. Castro, SAT encoding and CSP reduction for interconnected AllDiff constraints, in: 8th MICAI, 2009, pp. 360–371.
- [29] J.-L. Lauriere, ALICE: a language and a program for stating and solving combinatorial problems, Artif. Intell. 10 (1978) 29–127.

- [30] M. Leconte, A bounds-based reduction scheme for constraints of difference, in: 2nd CONSTRAINT, 1996, pp. 19–28.
- [31] A. López-Ortiz, C.-G. Quimper, J. Tromp, P. van Beek, A fast and simple algorithm for bounds consistency of the alldifferent constraint, in: 18th IJCAI, 2003, pp. 245–250.
- [32] D. Marx, Can you beat treewidth? in: 48th FOCS, 2007, pp. 169–179.
- [33] D. Marx, Tractable hypergraph properties for constraint satisfaction and conjunctive queries, in: 42nd STOC, 2010, pp. 735–744.
- [34] K. Mehlhorn, Constraint programming and graph algorithms, in: 27th ICALP, 2000, pp. 571–575.
- [35] K. Mehlhorn, S. Thiel, Faster algorithms for bound-consistency of the sortedness and the alldifferent constraint, in: 6th CP, 2000, pp. 306–319.
- [36] R. Niedermeier, Invitation to Fixed-Parameter Algorithms, Oxford University Press, 2006.
- [37] J.-F. Puget, A fast algorithm for the bound consistency of alldiff constraints, in: 13th AAI, 1998, pp. 359–366.
- [38] J.-C. Régin, A filtering algorithm for constraints of difference in CSPs, in: 9th AAI, 1994, pp. 362–367.
- [39] F. Roberts, Indifference graphs, in: Proof Techniques in Graph Theory, 1969, pp. 139–146.
- [40] M. Samer, S. Szeider, Constraint satisfaction with bounded treewidth revisited, J. Comput. System Sci. 76 (2010) 103–114.
- [41] T.J. Schaefer, The complexity of satisfiability problems, in: 10th STOC, 1978, pp. 216–226.
- [42] C. Schulte, P.J. Stuckey, Efficient constraint propagation engines, ACM Trans. Program. Lang. Syst. 31 (2008) 1–43.
- [43] B.M. Smith, Modelling, in: F. Rossi, P. van Beek, T. Walsh (Eds.), Handbook of Constraint Programming, chapter 11, 2006.
- [44] K. Stergiou, Heuristics for dynamically adapting propagation, in: 18th ECAI, 2008, pp. 485–489.
- [45] W.-J. van Hove, The AllDifferent constraint: a survey, in: 6th CSCP, 2001.
- [46] W.-J. van Hove, I. Katriel, Global constraints, in: F. Rossi, P. van Beek, T. Walsh (Eds.), Handbook of Constraint Programming, chapter 6, Elsevier, 2006.
- [47] J. Yip, P.V. Hentenryck, Evaluation of length-lex set variables, in: 15th CP, 2009, pp. 817–832.