

On Backdoors To Tractable Constraint Languages^{*}

Clement Carbonnel^{1,3}, Martin C. Cooper², and Emmanuel Hebrard¹

¹ CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France
{carbonnel,hebrard}@laas.fr

² IRIT, University of Toulouse III, 31062 Toulouse, France
cooper@irit.fr

³ University of Toulouse, INP Toulouse, LAAS, F-31400 Toulouse, France

Abstract. In the context of CSPs, a strong backdoor is a subset of variables such that every complete assignment yields a residual instance guaranteed to have a specified property. If the property allows efficient solving, then a small strong backdoor provides a reasonable decomposition of the original instance into easy instances. An important challenge is the design of algorithms that can find quickly a small strong backdoor if one exists. We present a systematic study of the parameterized complexity of backdoor detection when the target property is a restricted type of constraint language defined by means of a family of polymorphisms. In particular, we show that under the weak assumption that the polymorphisms are idempotent, the problem is unlikely to be FPT when the parameter is either r (the constraint arity) or k (the size of the backdoor) unless $P = NP$ or $FPT = W[2]$. When the parameter is $k + r$, however, we are able to identify large classes of languages for which the problem of finding a small backdoor is FPT.

1 Introduction

Unless $P=NP$, the constraint satisfaction problem (CSP) is in general intractable. However, one can empirically observe that solution methods scale well beyond what a worst-case complexity analysis would suggest.

In order to explain this gap, Williams, Gomes and Selman introduced the notion of *backdoor* [1]. A strong backdoor is a set of variables whose complete assignments all yield an easy residual problem. When it is small, it therefore corresponds to a weak spot of the problem through which it can be attacked. Indeed, by branching first on the variables of a backdoor of size k , we ensure that the depth of the search tree is bounded by k . There exists a similar notion of weak backdoor, ensuring that at least one assignment yields an easy problem, however, we shall focus on strong backdoors and omit the adjective “strong”.

Finding small backdoors is then extremely valuable in order to efficiently solve constraint problems, however, it is very likely to be itself intractable. In

^{*} supported by ANR Project ANR-10-BLAN-0210.

order to study the computational complexity of this problem, we usually consider backdoors with respect to a given tractable class T , i.e., such that all residual problems fall into the class T . In Boolean Satisfiability (SAT), it was shown that finding a minimum backdoor with respect to HORNSAT, 2-SAT and their disjunction is fixed-parameter tractable with respect to the backdoor size [2][3]. It is significantly harder, however, to do so with respect to bounded treewidth formulas [4]. In this paper we study the computational complexity of finding a strong backdoor to a *semantic* tractable class of CSP, in the same spirit as a very recent work by Gaspers et al. [3]. Assuming that $P \neq NP$, semantic tractable classes are characterized by unions and intersections of languages of constraints closed by some operations. We make the following three main contributions:

- We first consider the case where these operations are *idempotent*, and show that computing a k -backdoor with respect to such a class is NP-hard even on bounded arity CSPs, and W[2]-hard for the parameter k if the arity is not bounded. Observe that the scope of this result is extremely wide, as most tractable classes of interest are idempotent.
- Then, we characterize another large category of tractable classes, that we call *Helly*, and for which finding k -backdoors is fixed-parameter tractable in $k + r$ where r is the maximum arity of a constraint.
- Lastly, we show that finding k -backdoors with respect to many semantic tractable classes that are not Helly is W[2]-hard for $k + r$ (and remains W[2]-hard for k if r is fixed). However, we do not prove a strict dichotomy since a few other conditions must be met besides not being Helly.

The paper is organized as follows: After introducing the necessary technical background in Section 2, we study idempotent tractable classes in Section 3 and Helly classes as well as a family of non-Helly classes in Section 4.

2 Preliminaries

Constraint satisfaction problems A constraint satisfaction problem (CSP) is a triplet (X, D, C) where X is a set of variables, D is a domain of values, and C is a set of constraints. For simplicity, we assume D to be a finite subset of \mathbb{N} . A constraint is a pair (S, R) where $S \subseteq X$ is the scope of the constraint and R is an $|S|$ -ary relation on D , i.e. a subset of $D^{|S|}$ representing the possible assignments to S . A solution is an assignment $X \rightarrow D$ that satisfies every constraint. The goal is to decide whether a solution exists.

A *constraint language* is a set of relations. The domain of a constraint language Γ is denoted by $D(\Gamma)$ and contains all the values that appear in the tuples of the relations in Γ . Given a constraint language Γ , $\text{CSP}(\Gamma)$ is the restriction of the generic CSP to instances whose constraints are relations from Γ . The *Dichotomy Conjecture* by Feder and Vardi says that for every finite Γ , $\text{CSP}(\Gamma)$ is either in P or NP-complete [5]. Since the conjecture is still open, the complexity of constraint languages is a very active research area (see e.g. [6][7][8]).

It is known that the complexity of a language is determined by its set of closure operations [9]. Specifically, an operation $f : D(\Gamma)^a \rightarrow D(\Gamma)$ of arity a is a *polymorphism* of Γ if for every $R \in \Gamma$ of arity r and $t_1, \dots, t_a \in R$, $f(t_1, \dots, t_a) = (f(t_1[1], \dots, t_a[1]), \dots, f(t_1[r], \dots, t_a[r])) \in R$. A polymorphism f is *idempotent* if $\forall x \in D, f(x, x, \dots, x) = x$. We denote by $\text{Pol}(\Gamma)$ (resp. $\text{IdPol}(\Gamma)$) the set of all polymorphisms (resp. idempotent polymorphisms) of Γ . Given two languages Γ_1, Γ_2 with $D(\Gamma_2) \subseteq D(\Gamma_1)$, we write $\text{Pol}(\Gamma_1) \subseteq \text{Pol}(\Gamma_2)$ if the restriction to $D(\Gamma_2)$ of every $f \in \text{Pol}(\Gamma_1)$ is in $\text{Pol}(\Gamma_2)$. A wide range of operations have been shown to induce polynomial-time solvability of any language they preserve: these include near-unanimity operations [10], edges [11], semilattices [9], 2-semilattices [12] and totally symmetric operations of all arities [13].

Composite classes A semantic class is a set of languages. A semantic class T is tractable if $\text{CSP}(\Gamma) \in P$ for every $\Gamma \in T$, and recognizable in polynomial time if the membership problem ‘Does $\Gamma \in T$?’ is in P . We say that a semantic class T is *atomic* if there exists an operation $f : \mathbb{N}^a \rightarrow \mathbb{N}$ such that $\Gamma \in T$ if and only if $f|_{D(\Gamma)} \in \text{Pol}(\Gamma)$, where $f|_{D(\Gamma)}$ denotes the restriction of f to $D(\Gamma)$. We sometimes denote such a class by T_f and say that f *induces* T_f . We call a semantic class T *simple* if there exists a set \mathcal{T} of atomic classes such that $T = \bigcap_{T_f \in \mathcal{T}} T_f$. Finally, a semantic class T is *composite* if there exists a set \mathcal{T} of simple classes such that $T = \bigcup_{T_s \in \mathcal{T}} T_s$. In both cases, the set \mathcal{T} is allowed to be infinite. Using the distributivity of intersection over union, it is easy to see that any class derived from atomic classes through any combination of intersections and unions is composite. We say that an atomic class T_f is *idempotent* if f is idempotent. By extension, a composite class is idempotent if can be obtained by intersections and unions of idempotent atomic classes.

Example 1. Consider the class of max-closed constraints, introduced in [14]. This class is tractable as any CSP instance over a max-closed constraint language can be solved by establishing (generalised) arc-consistency. Using our terminology, this class is exactly the atomic class induced by the operation $\max(\cdot, \cdot)$, and thus it is composite. Max-closed constraints have been generalized to any language that admits a semilattice polymorphism, i.e. a binary operation f such that $f(x, x) = x$, $f(x, y) = f(y, x)$ and $f(f(x, y), z) = f(x, f(y, z))$ for any $x, y, z \in D$ [9]. If we denote by Sml the set of all possible semilattice operations on \mathbb{N} , this larger class corresponds to $\bigcup_{f \in \text{Sml}} T_f$, which is composite but not atomic.

Example 2. For a given language Γ , let $\bar{\Gamma}$ be the language obtained from Γ by adding all possible unary relations over $D(\Gamma)$ with a single tuple. Consider the very large class T_{BW} of languages Γ such that $\text{CSP}(\bar{\Gamma})$ (and thus $\text{CSP}(\Gamma)$) can be solved by achieving k -consistency for some k that only depends on Γ . This property is equivalent to the existence of two idempotent polymorphisms f and g such that for every $x, y \in D(\Gamma)$ [15][16],

- (i) $g(y, x, x, x) = g(x, y, x, x) = g(x, x, y, x) = g(x, x, x, y)$
- (ii) $f(y, x, x) = f(x, y, x) = f(x, x, y)$
- (iii) $f(x, x, y) = g(x, x, x, y)$

If we denote by FGBW the set of all pairs of operations (f, g) on \mathbb{N} satisfying these three conditions, the class T_{BW} is composite and idempotent since it can be written as $T_{BW} = \cup_{(f,g) \in \text{FGBW}} (T_f \cap T_g)$.

The choice to study composite classes shows multiple advantages. First, they are general enough to capture most natural semantic tractable classes defined in the literature, and they also allow us to group together tractable languages that are solved by the same algorithm (such as arc consistency or Gaussian elimination). Second, membership in these classes is hereditary: if $\Gamma \in T$ and $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma')$, then $\Gamma' \in T$. In particular, any sublanguage of a language in T is in T , and every composite class contains the empty language.

Strong backdoors Given an instance (X, D, C) of $\text{CSP}(\Gamma)$, assigning a variable $x \in X$ to a value $d \in D$ is done by removing the tuples inconsistent with $x \leftarrow d$ from the constraints whose scope include x , and then removing the variable x from the instance (thus effectively reducing the arity of the neighbouring constraints by one). A *strong backdoor* to a semantic class T is a subset $S \subseteq X$ such that every complete assignment of the variables from S yields an instance whose language is in T . Note that assigning a variable involves no further inference (e.g., arc consistency); indeed doing so has been shown to make backdoors potentially much harder to detect [17]. There exist alternative forms of backdoors, such as weak backdoors [1] and partition backdoors [18], but we only consider strong backdoors throughout this paper so we may omit the word “strong” in proofs. The goal of this work is to study how the properties of the target semantic class T affect the (parameterized) complexity of the following problem.

STRONG T -BACKDOOR: Given a CSP instance I and an integer k , does I have a strong backdoor to T of size at most k ?

Parameterized complexity A problem is *parameterized* if each instance x is coupled with a nonnegative integer k called the *parameter*. A parameterized problem is *fixed-parameter tractable* (FPT) if it can be solved in time $O(f(k)|x|^{O(1)})$, where f is any computable function. For instance, VERTEX COVER parameterized with the size k of the cover is FPT as it can be solved in time $O(1.2738^k + kn)$ [19], where n is the number of vertices of the input graph. The class XP contains the parameterized problems that can be solved in time $O(f(k)|x|^{g(k)})$ for some computable functions (f, g) . FPT is known to be a proper subset of XP [20]. Between these extremes lies the *Weft Hierarchy*:

$$\text{FPT} = \text{W}[0] \subseteq \text{W}[1] \subseteq \text{W}[2] \subseteq \dots \subseteq \text{XP}$$

where for every t , $\text{W}[t+1]$ is believed to be strictly larger than $\text{W}[t]$. These classes are closed under FPT-reductions, which map an instance (x, k) of a problem L_1 to an instance (x', k') of a problem L_2 such that:

- (x', k') can be built in time $O(f(k)|x|^{O(1)})$ for some computable function f
- (x', k') is a yes-instance if and only if (x, k) is
- $k' \leq g(k)$ for some computable function g

For instance, k -CLIQUE is $W[1]$ -complete [20] when the parameter is k . Note that if considering the parameter as a constant yields an NP-hard problem, the parameterized version is not in XP (and thus not FPT) unless $P = NP$.

3 General Hardness

We consider two parameters: k , the size of the backdoor, and r , the maximum arity of the constraint network. Under the very weak assumption that T is composite and idempotent, we show that STRONG T -BACKDOOR is unlikely to be FPT for either of the parameters taken separately, assuming that $P \neq NP$ and $FPT \neq W[2]$, as we shall do throughout the paper. In both cases, we show that our results extend to the class of Boolean CSPs with minor modifications.

Our hardness results will be obtained by reductions from various forms of the p -HITTING SET problem: given a universe U , a collection $S = \{S_i \mid i = 1..n\}$ of subsets of U with $|S_i| = p$ and an integer k , does there exist a subset $H \subseteq U$ such that $|H| \leq k$ and $\forall i, H \cap S_i \neq \emptyset$? This problem is NP-complete for every fixed $p \geq 2$ [21], and $W[2]$ -complete when the parameter is k and p is unbounded [20]. The special case $p = 2$ is called VERTEX COVER, and the input is typically given in the form of a graph $G = (U, S)$ and an integer k .

We will make use of two elementary properties of idempotent composite classes. First, any relation with a single tuple is closed by every idempotent operation. Thus, adding such a relation to a language does not affect its membership in idempotent classes. The second property is slightly more general. Given a relation R of arity r , let M_R be the matrix whose rows are the tuples of R sorted by lexicographic order (so that M_R is unique). We say that a relation R is an *extension* of a relation R' if M_R has all the columns of $M_{R'}$, plus extra columns that are either constant (i.e. every value in that column is the same) or copies of some columns of $M_{R'}$. In that case, since $\text{IdPol}(\{R\}) = \text{IdPol}(\{R'\})$, $\{R\} \in T$ if and only if $\{R'\} \in T$, for every idempotent composite class T .

For the rest of the document, we represent relations as lists of tuples delimited by square brackets (e.g. $R = [t_1, \dots, t_n]$), while tuples are delimited by parentheses (e.g. $t_1 = (d_1, \dots, d_r)$).

3.1 Hardness on bounded arity CSPs

Theorem 1. STRONG T -BACKDOOR is NP-hard for every idempotent composite tractable class T , even for binary CSPs.

Proof. We reduce from VERTEX COVER. Let $I = (G, k)$ be an instance of VERTEX COVER. We consider two cases. First, suppose that $\Gamma = \{[(1), (2)], [(2), (3)], [(1), (3)]\} \in T$. We create a CSP with one variable per vertex in G , and if two variables correspond to adjacent vertices we add the constraint $\neq_{1,2,3}$ (inequality over the 3-element domain) between them. Since $\text{CSP}(\{\neq_{1,2,3}\})$ is NP-hard and T is tractable, a valid backdoor of size at most k must correspond to a vertex cover on G . Conversely, the variables corresponding to a vertex cover form a

backdoor: after every complete assignment to these variables, the language of the reduced instance is a subset of Γ , which is in T since T is composite and hence hereditary. Now, suppose that $\Gamma \notin T$. We duplicate the column of each relation in Γ to obtain the binary language $\Gamma' = \{R_1, R_2, R_3\}$ with $R_1 = [(1, 1), (2, 2)]$, $R_2 = [(2, 2), (3, 3)]$ and $R_3 = [(1, 1), (3, 3)]$. Since Γ' is an extension of Γ and T is idempotent, Γ' is not in T . Then, we follow the same reduction as in the first case, except that we add the three constraints R_1, R_2, R_3 instead of $\neq_{1,2,3}$ between two variables associated with adjacent vertices. By construction, a backdoor must be a vertex cover. Conversely, if we have a vertex cover, after any assignment of the corresponding variables we are left with at most one tuple per constraint, and the resulting language is in T by idempotency. \square

In the case of Boolean CSPs, Theorem 1 cannot apply verbatim. This is due to the fact that every binary Boolean language is a special case of 2-SAT and is therefore tractable. Thus, a binary Boolean CSP has always a backdoor of size 0 to any class that is large enough to contain 2-SAT, and the minimum backdoor problem is trivial. The next proposition shows that this is the only case for which STRONG T -BACKDOOR is not NP-hard under the idempotency condition. Note that looking for a strong backdoor in a binary Boolean CSP has no practical interest; however this case is considered for completeness.

Proposition 1. *On Boolean CSPs with arity at most r , STRONG T -BACKDOOR is NP-hard for every idempotent composite tractable class T if $r \geq 3$. For $r = 2$, STRONG T -BACKDOOR is either trivial (if every binary Boolean language is in T) or NP-hard.*

Proof. The proof is essentially identical to that of Theorem 1, only with a larger number of cases to examine. First, suppose that every binary Boolean language is in T . Then, if $r = 2$, STRONG T -BACKDOOR is trivial. If $r \geq 3$, we reduce from 3-HITTING SET. Let $I = (U, S, k)$ be an instance of 3-HITTING SET. We create a CSP instance with one variable per element in U , and for each $(u_i, u_j, u_l) \in S$, if $R_1 = [(1, 0), (0, 1)] \in T$ we add the constraint $R_2 = [(1, 0, 0), (0, 1, 0), (0, 0, 1)]$ on the corresponding variables and $R_3 = [(1, 0, 0), (0, 1, 1)]$ otherwise. Observe that, since $\text{CSP}(\{R_2\})$ is known to be NP-hard (by a reduction from (1 in 3)-SAT) and R_3 is an extension of R_1 , in either case the added constraint does not belong to T . Thus, if a backdoor of size at most k exists, then the corresponding subset of U must be a hitting set. Conversely, let B denote the CSP variables associated with a hitting set of size at most k . If we used the constraint R_2 in the reduction, then after a complete assignment of B the remaining constraints are a subset of $\{R_1, [(0)], [(1)], [(0, 0)]\}$ which is in T since T is idempotent and $R_1 \in T$. If the reduction was done using R_3 , after any assignment we are left with a CSP instance whose language is a subset of $\{[(0)], [(1)]\}$ which is in T by idempotency. Finally, in both cases, B is a backdoor of size k , which completes this part of the reduction. Now, suppose that there exists a binary Boolean language Γ that is not in T . Once more, we reduce from VERTEX COVER. Let (G, k) be an instance of VERTEX COVER and $\Gamma_1 = \{[(0, 1)]\}$. We create a CSP instance with one variable per vertex in G . If two vertices are adjacent, we add between them the

constraint of T_1 if $T_1 \notin T$ and all constraints in T otherwise. By construction any T-backdoor of size at most k must be a vertex cover, and the same line of reasoning as in the proof of Theorem 1 gives us that the variables corresponding to any hitting set of size at most k is a T-backdoor, which concludes the proof. \square

3.2 Hardness when the parameter is the size of the backdoor

In general, a large strong backdoor is not of great computational interest as the associated decomposition of the original instance is very impractical. Thus, it makes sense to design algorithms that are FPT when the parameter is the size of the backdoor. In this section we show that in the case of idempotent composite classes such algorithms cannot exist unless $\text{FPT} = \text{W}[2]$. Furthermore, we establish this result under the very restrictive condition that the input CSP has a single constraint, which highlights the fact that **STRONG T-BACKDOOR** is more than a simple pseudo-HITTING SET on the constraints outside T .

For any natural numbers m, e ($m \geq 3$), we denote by $R_3^m(e)$ the relation obtained by duplicating the last column of $[(e+1, e, e), (e, e+1, e), (e, e, e+1)]$ until the total arity becomes m . It is straightforward to see that $\text{CSP}(\{R_3^m(e)\})$ is NP-hard for every m, e by a reduction from 1-in-3-SAT. In a similar fashion, we define $R_2^m(e)$ as an extension of $[(e+1, e), (e, e+1)]$ of arity m .

Theorem 2. *STRONG T-BACKDOOR is $W[2]$ -hard for every idempotent composite tractable class T when the parameter is the size of the backdoor, even if the CSP has a single constraint.*

Proof. The proof is an FPT-reduction from p -HITTING SET parameterized with solution size k . Let (p, U, S) be an instance of p -HITTING SET, where U is the universe ($|U| = n$) and $S = \{S_i \mid i = 1..s\}$ is the collection of p -sets. We assume without loss of generality that $p \geq 3$ (if this is not the case we pad each set with unique elements). We build an n -ary relation R , where each column is associated with a value from U , as follows. For every $S_i \in S$, we consider two cases. If $\{R_2^2(2i)\} \notin T$, we add two tuples t_1, t_2 to R such that the restriction of $[t_1, t_2]$ to the columns corresponding to the values appearing in S_i form the relation $R_2^p(2i)$, and the other columns are constant with value $2i$. If $\{R_2^2(2i)\} \in T$, we add 3 tuples t_1, t_2, t_3 such that the restriction of t_1, t_2, t_3 to the columns corresponding to S_i form $R_3^p(2i)$, and the remaining columns are constant with value $2i$. Once the relation is complete, we apply it to n variables to obtain an instance of our backdoor problem. See Figure 1 for an example of the construction.

Suppose we have a backdoor of size at most k , and suppose there exists a set S_i such that none of the corresponding variables belong to the backdoor. Then, if we assign every variable in the backdoor to $2i$, the reduced constraint must belong to T . By idempotency, we can further assign every remaining variable outside of S_i to the value $2i$ and the resulting constraint must still be in T . The reduced constraint becomes either $R_2^p(2i)$ if $\{R_2^2(2i)\} \notin T$, or $R_3^p(2i)$, which is not in T since T is tractable and we assume $\text{P} \neq \text{NP}$. In both cases, this




Fig. 1: Example of reduction from a 3-HITTING SET instance to the problem of finding a backdoor to the class T . The reduction produces a single constraint C .

constraint does not belong to T , and we have a contradiction. Therefore, if there is a backdoor of size at most k , we also have a hitting set of size at most k .

Conversely, suppose we have a hitting set of size at most k . We prove that the associated set of variables form a backdoor. Observe that two blocks (i.e. pairs/triples) of tuples of the constraint C associated with different sets do not share any common value; hence, after assigning the variables corresponding to the hitting set to any values, the resulting constraint is either empty or a sub-relation of a single block associated with the set S_i . The latter case yields two possibilities. If T does not contain $\{R_2^2(2i)\}$, then the block i must have been reduced to a single tuple, since the two initial tuples t_1, t_2 satisfy $t_1[x_j] \neq t_2[x_j]$ for all x_j associated with a value in S_i . Thus, by idempotency, the resulting constraint is in T . Now, if T contains $\{R_2^2(2i)\}$, the resulting constraint has at most two tuples (same argument as above), which can only happen if all the variables are assigned to the value $2i$. If we are in this situation, the new constraint must be an extension of $R_2^2(2i)$ and hence is in T . Therefore, our hitting set provides a strong backdoor in our CSP instance, which concludes the reduction. \square

This theorem still holds on Boolean CSPs if we allow multiple constraints in the target instance, even if these constraints are all the same relation.

Proposition 2. *On Boolean CSPs, STRONG T -BACKDOOR is $W[2]$ -hard for every idempotent composite tractable class T when the parameter is the size of the backdoor, even if the CSP has a single type of constraint.*

Proof. The proof is similar to the non-Boolean case, but more straightforward since we are allowed multiple constraints. We FPT-reduce from p -HITTING SET parameterized by solution size. Let (p, U, S) be an instance of p -HITTING SET, where U is the universe ($|U| = n$) and $S = \{S_i \mid i = 1..s\}$ is the collection of p -sets. We assume $p \geq 3$, as we did for Theorem 2. We create a CSP instance with one variable per element in U . For each set $S_i \in S$, if $\{R_2^p(0)\} \in T$ we add the constraint $R_3^p(0)$ on the variables corresponding to the values in S_i , and $R_2^p(0)$ otherwise. By construction, the language of the instance only contains constraints outside T , so a strong backdoor of size k must intersect every constraint and hence corresponds to a hitting set of (p, U, S) . Conversely, the set of variables B corresponding to a hitting set of size at most k form a backdoor: After every assignment of B , at least one variable in each constraint is assigned, so

the language is either formed of relations with a most one tuple (if $\{R_2^p(0)\} \notin T$) or a collection of extensions of $R_2^p(0)$ ($\{R_2^p(0)\} \notin T$) plus relations with at most one tuple. In either case, the resulting language is in T by idempotency, which concludes the proof. \square

Remark 1. Partition Backdoors is an alternative form of backdoors recently introduced by Bessiere et al [18]. Such backdoors are especially interesting in the case of conservative classes (conservativity is more restrictive than idempotency, since each polymorphism is required to satisfy $f(x_1, \dots, x_a) \in \{x_1, \dots, x_a\}$). The authors argue that, given a partition of the constraints $C = \{C_1, C_2\}$ such that the language of C_1 is in a conservative class T , the vertex cover of the primal graph of C_2 is a strong backdoor to T . The minimum-size *partition backdoor* is then the best such backdoor over every possible partition of the constraints. Computing the minimum-size partition backdoor is FPT in the parameter $k + l$, where l is the size of the constraint language; our results show that computing the actual minimum strong backdoor is a much harder problem as it is still W[2]-hard for the larger parameters $k + m$ (Theorem 2) and $d + k + l$ (Proposition 2), where m is the number of constraints and d the size of the domain.

4 Combined parameters: Helly classes and limits

We have shown in sections 3.1 and 3.2 that considering independently the maximum constraint arity r and the size of the backdoor k as parameters is unlikely to yield FPT tractability. We now consider the combined parameter $k + r$ and show that FPT tractability ensues for numerous tractable composite classes.

In order to design an algorithm for STRONG T -BACKDOOR that is FPT for $k + r$, it is important to have a procedure to *check* whether a subset of variables of size at most k is a strong backdoor to T . The natural algorithm for this task runs in time $O(mrt d^k P(\Gamma))$ (where m is the number of constraints, t the maximum number of tuples and $P(\Gamma)$ the complexity of the membership problem of a language Γ in T) by checking independently each of the d^k possible assignments of B . In our case this approach is not satisfactory: since d is not a parameter, the term d^k is problematic for the prospect of an algorithm FPT in $k + r$. The next lemma presents an alternative algorithm for the “backdoor check” problem that is only exponential in the number of constraints m . Although it may seem impractical at first sight (as m is typically much larger than k), we will show that it can be exploited for many tractable classes.

Lemma 1. *Let T be a composite class recognizable in time $P(\Gamma)$. Let $I = (X, D, C)$ be a CSP instance with m constraints of arity at most r and containing at most t tuples, and $B \subseteq X$. It is possible to decide whether B is a strong backdoor to T in time $O(mrt^2 + m^2r(2t)^m P(\Gamma))$.*

Proof. We first focus on a single constraint (S, R) . Let $B_S = B \cap S$. Observe that at most t different assignments of B_S can leave R nonempty, since the subrelations of R obtained with each assignment are pairwise disjoint and their

union is R . To compute these assignments in polynomial time, one can explore a search tree. Starting from a node labelled R , we pick a nonfixed variable $v \in B_S$ and for every $d \in D(v)$ such that the subrelation $R_{v=d}$ is not empty we create a child node labelled with $R_{v=d}$. Applying this rule recursively, we obtain a tree of depth at most r and with no more than t leaves, so it has at most rt nodes. The time spent at each node is $O(t)$, so computing all leaves can be done in time $O(rt^2)$. Now, suppose that for each constraint $c = (R, S) \in C$ we have computed this set ϕ_c of all the locally consistent assignments of B_S and stored the resulting subrelation. For every $\phi \in \prod_{c \in C} \phi_c$ and every possible subset C' of the constraints, we check if the restriction ϕ_s of ϕ to the constraints of C' is a consistent assignment (i.e., no variable is assigned multiple values). If ϕ_s is consistent, we temporarily remove from the instance the constraints outside C' , we apply the assignment ϕ_s and we check whether the language of the resulting instance is in T . The algorithm returns that B is a backdoor if and only if each membership test in T is successful. To prove the correctness of the algorithm, suppose that ψ is an assignment of B such that the resulting language is not in T . Then, at least one subset of the constraints have degraded into non-empty subrelations. For each of these constraints, the restricted assignment ψ_R is consistent with the others, so the algorithm must have checked membership of the resulting language in T and concluded that B is not a strong backdoor. Conversely, if B is a strong backdoor, every complete assignment of B yields an instance in T . In particular, if we consider only a subset of the constraints after each assignment, the language obtained is also in T since T is composite (and hence hereditary). Thus, none of the membership tests performed by the algorithm will fail. The complexity of the algorithm is $O(mrt^2 + m^2rt^m 2^m P(\Gamma))$. \square

We say that a composite class T is h -Helly if it holds that for any language Γ , if every $\Gamma_h \subseteq \Gamma$ of size at most h is in T then Γ is in T . This property is analogous to the well-studied Helly properties for set systems. We call Helly number of T the minimum positive integer h such that T is h -Helly. Being characteristic of a class defined exclusively in terms of polymorphisms over \mathbb{N} , the Helly number is independant from parameters like the domain size or the arity of the languages. The next theorem is the motivation for the study of such classes, and is the main result of this section.

Theorem 3. *For every fixed composite class T recognizable in polynomial time, if T has a finite Helly number then STRONG T -BACKDOOR is FPT when the parameter is $k + r$, where k is the size of the backdoor and r is the maximum constraint arity.*

Proof. Let h denote the Helly number of T . The algorithm is a bounded search tree that proceeds as follows. Each node is labelled by a subset of variables B . The root of the tree is labelled with the empty set. At each node, we examine every possible combination of h constraints and check if B is a strong backdoor for the subset in time $O(hrt^2 + h^2r(2t)^h P(\Gamma))$ (where $P(\Gamma)$ is the polynomial complexity of deciding the membership of a language Γ in T) using Lemma 1. Suppose that

B is a strong backdoor for every h -subset. Then, for any possible assignment of B , each h -subset of the constraints of the resulting instance must be in T : otherwise, B would not be a strong backdoor for the h original constraints that generated them. Since T is h -Helly, we can conclude that B is a valid strong backdoor for the whole instance. Now suppose that we have found a h -subset for which B is not a strong backdoor. For every variable x in the union of the scopes of the constraints in this subset that is not already in B (there are at most rh such variables x), we create a child node labelled with $B \cup \{x\}$. At each step we are guaranteed to add at least one variable to B , so we stop creating child nodes when we reach depth k . The algorithm returns ‘YES’ at the first node visited that corresponds to a strong backdoor, and ‘NO’ if no such node is found.

If no strong backdoor of size at most k exists, it is clear that the algorithm correctly returns ‘NO’. Now suppose that a strong backdoor \mathbf{B} , $|\mathbf{B}| \leq k$, exists. Observe that if a node is labelled with $B \subset \mathbf{B}$ and B is not a backdoor for some h -subset of constraints, then \mathbf{B} contains at least one more variable within this subset. Since the algorithm creates one child per variable that can be added and the root is labelled with a subset of \mathbf{B} , by induction there must be a path from the root to a node labelled with \mathbf{B} and the algorithm returns ‘YES’.

The complexity of the procedure is $O((rh)^k m^h (hrt^2 + h^2 r (2t)^h P(\Gamma))) = O(f(k+r)m^h(ht^2 + h^2(2t)^h P(\Gamma)))$. \square

In contrast to the previous hardness results, the target tractable class is not required to be idempotent. However, the target class must have a finite Helly number, which may seem restrictive. The following series of results aims to identify composite classes with this particular property.

Lemma 2. *A composite class T is simple if and only if it is 1-Helly.*

Proof. Let T be a simple class, i.e. an intersection of atomic classes $T = \bigcap_{f \in \mathcal{F}} T_f$. Let Γ be a constraint language such that each $\{R\} \subseteq \Gamma$ is in T . Then, every $f \in \mathcal{F}$ preserves every relation in Γ and thus preserves Γ , so $\Gamma \in T$ and T is 1-Helly. Conversely, let T be a 1-Helly composite class. Let $\mathcal{F} = \{f \mid f \text{ preserves every } \{R\} \in T\}$. Every $\Gamma \in T$ admits as polymorphism every $f \in \mathcal{F}$ (as each $\{R\} \subseteq \Gamma$ is in T and thus is preserved by f), so $T \subseteq \bigcap_{f \in \mathcal{F}} T_f$. The other way round, a language Γ in $\bigcap_{f \in \mathcal{F}} T_f$ is preserved by every $f \in \mathcal{F}$ and thus must be a sublanguge of $\Gamma_\infty = \bigcup_{\{R\} \in T} \{R\}$, which is in T since T is 1-Helly, so $\Gamma \in T$ and $\bigcap_{f \in \mathcal{F}} T_f \subseteq T$. Finally, $T = \bigcap_{f \in \mathcal{F}} T_f$ and so T is simple. \square

Proposition 3. *Let h be a positive integer and \mathcal{T} be a set of simple classes. Then, $T = \{\Gamma \mid \Gamma \text{ belongs to every } T_i \in \mathcal{T} \text{ except at most } h\}$ is a $(h+1)$ -Helly composite class.*

Proof. T is composite since it is the union of every possible intersection of all but h classes from \mathcal{T} and any class derived from atomic classes through any combination of intersections and unions is composite. We write $\mathcal{T} = \{T_i \mid i \in I\}$. Let Γ be a language such that every sublanguge of size at most $h+1$ is in T . For each $R \in \Gamma$ we define $S(R) = \{T_i \mid \{R\} \notin T_i\}$. By Lemma 2, simple classes are

1-Helly so $\Gamma \notin T_i \Leftrightarrow (\exists R \in \Gamma \text{ such that } \{R\} \notin T_i) \Leftrightarrow T_i \in \cup_{R \in \Gamma} S(R)$. So $\Gamma \in T$ if and only if $|\cup_{R \in \Gamma} S(R)| \leq h$. We discard from Γ every relation R such that $|S(R)| = 0$ as they have no influence on the membership of Γ in T . If that process leaves Γ empty, then it belongs to T . Otherwise, let s_j denote the maximum size of $\cup_{R \in \Gamma_j} S(R)$ over all size- j subsets Γ_j of Γ . Since each sublanguange Γ_j of size $j \leq h + 1$ is in T , from the argument above we have $1 \leq s_1 \leq \dots \leq s_{h+1} \leq h$, thus there exists $j < h + 1$ such that $s_j = s_{j+1}$. Let $\Gamma_j \subseteq \Gamma$ denote a set of j relations such that $|\cup_{R \in \Gamma_j} S(R)| = s_j$. Suppose there exists $R_0 \in \Gamma$ such that $S(R_0) \not\subseteq \cup_{R \in \Gamma_j} S(R)$. Then, $|\cup_{R \in \Gamma_j \cup \{R_0\}} S(R)| > s_j = s_{j+1}$, and we get a contradiction. So $\cup_{R \in \Gamma} S(R) \subseteq \cup_{R \in \Gamma_j} S(R)$, hence $|\cup_{R \in \Gamma} S(R)| \leq h$ and Γ is in T . Therefore, T is $(h + 1)$ -Helly. \square

In the particular case where \mathcal{T} is finite and $h = |\mathcal{T}| - 1$, we get the following nice corollary. Recall that a composite class is any union of simple classes.

Corollary 1. *Any union of h simple classes is h -Helly.*

Example 3. Let $T = \{\Gamma \mid \Gamma \text{ is either min-closed, max-closed or 0/1/all}\}$. T is the union of 3 well-known tractable semantic classes. By definition, min-closed and max-closed constraints are respectively the languages that admit $\min(\dots)$ and $\max(\dots)$ as polymorphisms. Likewise, 0/1/all constraints have been shown to be exactly the languages that admit as polymorphism the majority operation [9]

$$f(x, y, z) = \begin{cases} y & \text{if } y = z \\ x & \text{otherwise} \end{cases}$$

Thus, T is the union of 3 atomic classes and hence is 3-Helly by Corollary 1. Since T is also recognizable in polynomial time, by Theorem 3 STRONG T -BACKDOOR is FPT when parameterized by backdoor size and maximum arity.

In the light of these results, it would be very interesting to show a dichotomy. Is STRONG T -BACKDOOR with parameter $k + r$ at least W[1]-hard for every tractable composite class T that does not have a finite Helly number? While we leave most of this question unanswered, we have identified generic sufficient conditions for W[2]-hardness when r is fixed and the parameter is k .

Given a bijection $\phi : D_1 \rightarrow D_2$, we denote by R_ϕ the relation $[(d, \phi(d)), d \in D_1]$. Given a language Γ , a subdomain D_1 of $D(\Gamma)$ is said to be *conservative* if every $f \in \text{Pol}(\Gamma)$ satisfies $f(x_1, \dots, x_m) \in D_1$ whenever $\{x_1, \dots, x_m\} \subseteq D_1$. For instance, $D(\Gamma')$ is conservative for every $\Gamma' \subseteq \Gamma$, and for every column of some $R \in \Gamma$ the set of values that appear in that column is conservative. Then, we say that a class T is *value-renamable* if for every $\Gamma \in T$ and $\phi : D_1 \rightarrow D_2$, where D_1 is a conservative subdomain of Γ and $D_2 \cap D(\Gamma) = \emptyset$, $\Gamma \cup \{R_\phi\}$ is in T . For instance, the class of 0/1/all constraints introduced in Example 3 is value-renamable, but max-closed constraints are not (as they rely on a fixed order on \mathbb{N}). We also say that a composite class T is *domain-decomposable* if for each pair of languages $\Gamma_1 \in T$ and $\Gamma_2 \in T$, $D(\Gamma_1) \cap D(\Gamma_2) = \emptyset$ implies $\Gamma_1 \cup \Gamma_2 \in T$. Value-renamability and domain-decomposability are natural properties of any

class that is large enough to be invariant under minor (from the algorithmic viewpoint) modifications of the constraint languages.

Given a language Γ and a bijection $\phi : D(\Gamma) \rightarrow D'$, we denote by $\phi(\Gamma)$ the language over D' obtained by replacing every tuple $t = (d_1, \dots, d_r)$ in every relation in Γ by $\phi(t) = (\phi(d_1), \dots, \phi(d_r))$.

Lemma 3. *Let $\Gamma = \Gamma_1 \cup \Gamma_2 \cup \{R_\phi\}$ where ϕ is a bijection from $D(\Gamma_2)$ to some domain D_1 . Then, $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma_1 \cup \phi(\Gamma_2))$.*

Proof. Let $f \in \text{Pol}(\Gamma)$ of arity a . We only need to show that f preserves $\phi(\Gamma_2)$, as f already preserves Γ_1 . Since f preserves R_ϕ , for each $(d_1, \phi(d_1)), \dots, (d_a, \phi(d_a)) \in R_\phi$ we have $(f(d_1, \dots, d_a), f(\phi(d_1), \dots, \phi(d_a))) \in R_\phi$, so $f(\phi(d_1), \dots, \phi(d_a)) = \phi(f(d_1, \dots, d_a))$ for every $d_1, \dots, d_a \in D(\Gamma_2)$. Then, given a tuples $\phi(t_1), \dots, \phi(t_a)$ of $\phi(\Gamma_2)$, $f(\phi(t_1), \dots, \phi(t_a)) = \phi(f(t_1, \dots, t_a)) \in \phi(\Gamma_2)$ since $f(t_1, \dots, t_a) \in \Gamma_2$. Therefore, f is a polymorphism of $\phi(\Gamma_2)$ and $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma_1 \cup \phi(\Gamma_2))$.

Theorem 4. *On CSPs with arity at most r , if T is a composite class that is*

- idempotent
- not 1-Helly for constraints of arity at most r
- value-renamable
- domain-decomposable

then STRONG T -BACKDOOR is $W[2]$ -hard when the parameter is k .

Proof. Since T is not 1-Helly for constraints of arity at most r , there exists a language $\Gamma_m = \{R_i \mid i \in 1..l_m\}$ (of arity $r_m \leq r$ and over a domain $D_m, |D_m| = d_m$) such that $l_m > 1$ and every sublanguange of Γ_m is in T but Γ_m is not. Since T is fixed, we shall consider that Γ_m is fixed as well and hence has constant size. We assume for simplicity of presentation that every $R \in \Gamma_m$ has arity r_m .

We perform an FPT-reduction from p -HITTING SET parameterized with solution size as follows. Let (p, U, S) be an instance of p -HITTING SET, with $S = \{S_1, \dots, S_s\}$ and $U = \{u_1, \dots, u_n\}$. For every $u_i \in U$, we associate a unique variable x_i . For every set $S_j = (u_{\sigma_j(1)}, \dots, u_{\sigma_j(p)})$, we add $2r_m$ new variables $y_j^1, \dots, y_j^{r_m}, z_j^1, \dots, z_j^{r_m}$ and we create $p+2$ new disjoint domains D_j^i , $i \in [0 \dots p+1]$ of size d_m . Then, we pick a chain of $p+1$ bijections $\psi_j^i : D_j^i \rightarrow D_j^{i+1}$, $i \in [0 \dots p]$ and we add a chain of constraints $R_{\psi_j^i}$ between the $p+2$ variables $(y_j^{r_m}, x_{\sigma_j(1)}, \dots, x_{\sigma_j(p)}, z_j^1)$. Afterwards, we pick a bijection $\phi_j : D_m \rightarrow D_j^0$ and we apply $\phi_j(R_1)$ to $y_j^1, \dots, y_j^{r_m}$. In the same fashion, if we denote by ψ_j the bijection from D_j^0 to D_j^{p+1} obtained by composition of all the ψ_j^i , we apply every constraint in $(\psi_j \circ \phi_j)(\Gamma_m \setminus \{R_1\})$ to the variables $z_j^1, \dots, z_j^{r_m}$. The main idea behind the construction is that both $\Gamma_m \setminus \{R_1\}$ and $\{R_1\}$ are in T but Γ_m is not: by adding $\phi_j(R_1)$ on the variables y , $\psi_j \circ \phi_j(\Gamma_m \setminus \{R_1\})$ on the variables z and the chain of bijections $R_{\psi_j^i}$ of the x variables, we have a language that is not in T but assigning any value to x yields a residual language in T (the proof can be found below). We use this property to encode a HITTING SET instance. See Figure 2 for an example of the reduction.

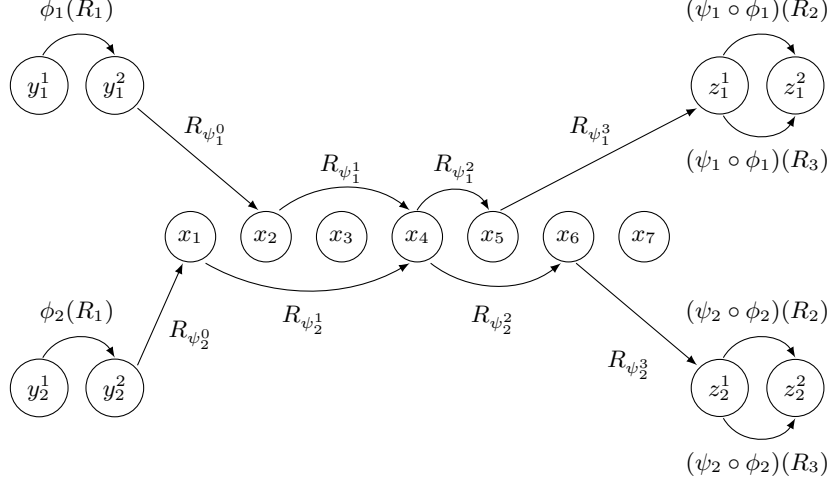


Fig. 2: Example of the construction for $U = (u_1, \dots, u_7)$, two sets $S_1 = (u_2, u_4, u_5)$, $S_2 = (u_1, u_4, u_6)$, $\Gamma_m = \{R_1, R_2, R_3\}$ and $r_m = 2$. Each arrow is a (binary) constraint. The upper part of the instance is constructed from S_1 and the lower part from S_2 .

Suppose we have a backdoor to T of size at most k . Then, for each set S_j , at least one variable from $(y_j^1, \dots, y_j^{r_m}, x_{\sigma_j(1)}, \dots, x_{\sigma_j(p)}, z_j^1, \dots, z_j^{r_m})$ must belong to the backdoor. Suppose this is not the case. Then, the language Γ of any reduced instance would contain the relations of $\{\phi_j(R_1), (\psi_j \circ \phi_j)(\Gamma_m \setminus \{R_1\})\}$ plus the relations $R_{\psi_j^i}$. Applying Lemma 3 $p+1$ times, we get $\text{Pol}(\Gamma) \subseteq \text{Pol}((\psi_j \circ \phi_j)(R_1) \cup (\psi_j \circ \phi_j)(\Gamma_m \setminus \{R_1\})) = \text{Pol}((\psi_j \circ \phi_j)(\Gamma_m))$. Thus, if Γ is in T , then so is $(\psi_j \circ \phi_j)(\Gamma_m)$. Then, by value-renamability $\{(\psi_j \circ \phi_j)(\Gamma_m) \cup R_{(\psi_j \circ \phi_j)^{-1}}\}$ is also in T and using Lemma 3 again, Γ_m is in T , which is a contradiction. Therefore, a hitting set of size at most k can be constructed by including every value u_i such that x_i is in the backdoor, and if any variable from $y_j^1, \dots, y_j^{r_m}, z_j^1, \dots, z_j^{r_m}$ belongs to the backdoor for some j , we also include $u_{\sigma_j(1)}$.

Conversely, a hitting set forms a backdoor. After every complete assignment of the variables from the hitting set, the set of constraints associated with any set S_j can be partitioned into sublanguages whose domains have an empty intersection (see Figure 2). The sublanguages are either:

- $\phi_j(R_1)$ together with some constraints $R_{\psi_j^i}$ and a residual unary constraint with a single tuple. This language is in T by Lemma 3, value-renamability and idempotency.
- $(\psi_j \circ \phi_j)(\Gamma_m \setminus \{R_1\})$ together with some constraints $R_{\psi_j^i}$ and a residual unary constraint with a single tuple. This case is symmetric.
- A (possibly empty) chain of constraints $R_{\psi_j^i}$ plus unary constraints with a single tuple, which is again in T since T is idempotent, value-renamable and contains the language $\{\emptyset\}$.

Furthermore, the sublanguages associated with different sets S_j also have an empty domain intersection. Since T is domain-decomposable, the resulting language is in T . \square

Note that this result does not conflict with Theorem 3, since any class T that is domain-decomposable, value-renamable and not 1-Helly cannot have a finite Helly number (part of the proof of Theorem 4 amounts to showing that one can build in polynomial time arbitrarily large languages Γ such that every sublanguange is in T but Γ is not). While the proof may seem technical, the theorem is actually easy to use and applies almost immediately to many known tractable classes: to prove W[2]-hardness of STRONG T -BACKDOOR on CSPs of arity bounded by r , one only has to prove value-renamability, domain-decomposability (which is usually straightforward) and exhibit a language Γ such that each $\{R\} \subset \Gamma$ is in T but Γ is not.

Example 4. An idempotent operation f is totally symmetric (TSI) if it satisfies $f(x_1, \dots, x_a) = f(y_1, \dots, y_a)$ whenever $\{x_1, \dots, x_a\} = \{y_1, \dots, y_a\}$. Using the same notations as in Example 2, it has been shown in [13] that $\text{CSP}(\Gamma)$ is solved by arc-consistency if and only if Γ has TSI polymorphisms of all arities. We show that this class of languages (which we denote by T_{TSI}) falls in the scope of Theorem 4 even for binary relations. First, this class is composite and idempotent: If we denote by $\text{TS}(a)$ the set of all possible TSI operations on \mathbb{N} of arity a and $\text{ATS} = \prod_{a \in \mathbb{N}^*} \text{TS}(a)$, we have $T_{\text{TSI}} = \cup_{\mathcal{F} \in \text{ATS}} (\cap_{f \in \mathcal{F}} T_f)$. To prove domain-decomposability and value-renamability, we will use the equivalent and more convenient characterization that Γ is in T_{TSI} if and only if Γ has a TSI of arity $|D(\Gamma)|$. Without loss of generality, we consider TSI polymorphisms as set functions and write $f(x_1, \dots, x_a) = f(\{x_1, \dots, x_a\})$.

- Domain-decomposability: Let $\Gamma_1, \Gamma_2 \in T$ be constraint languages with respective TSI polymorphisms f_1, f_2 (of respective arities $|D(\Gamma_1)|, |D(\Gamma_2)|$), and $D(\Gamma_1) \cap D(\Gamma_2) = \emptyset$. Let f be the operation on $D(\Gamma_1) \cup D(\Gamma_2)$ of arity $|D(\Gamma_1) \cup D(\Gamma_2)|$ defined as follows:

$$f(x_1, \dots, x_m) = \begin{cases} f_1(\{x_1, \dots, x_m\}) & \text{if } \{x_1, \dots, x_m\} \subseteq D(\Gamma_1) \\ f_2(\{x_1, \dots, x_m\}) & \text{if } \{x_1, \dots, x_m\} \subseteq D(\Gamma_2) \\ \max(x_1, \dots, x_m) & \text{otherwise} \end{cases}$$

f is totally symmetric and preserves both Γ_1 and Γ_2 , so f is a polymorphism of $\Gamma_1 \cup \Gamma_2$ and $\Gamma_1 \cup \Gamma_2 \in T_{\text{TSI}}$. Therefore, T_{TSI} is domain-decomposable.

- Value-renamability: Let $\Gamma \in T_{\text{TSI}}$ be a language with a TSI polymorphism f_1 of arity $|D(\Gamma)|$. Let $\phi : D_1 \rightarrow D_2$ be a bijection, where D_1 is a conservative subdomain of $D(\Gamma)$ and $D_2 \cap D(\Gamma) = \emptyset$. Then, the operation of arity $|D(\Gamma) \cup D_2|$ defined as

$$f(x_1, \dots, x_m) = \begin{cases} f_1(\{x_1, \dots, x_m\}) & \text{if } \{x_1, \dots, x_m\} \subseteq D(\Gamma) \\ \phi(f_2(\{\phi^{-1}(x_1), \dots, \phi^{-1}(x_m)\})) & \text{if } \{x_1, \dots, x_m\} \subseteq D_2 \\ \max(x_1, \dots, x_m) & \text{otherwise} \end{cases}$$

is a TSI and preserves both Γ and R_ϕ (the proof is straightforward using the fact that D_1 is a conservative subdomain), so $\Gamma \cup \{R_\phi\} \in T_{\text{TSI}}$ and T_{TSI} is value-renamable.

- Not 1-Helly: Let $R_1 = [(0, 0), (0, 1), (1, 0)]$ and $R_2 = [(1, 1), (0, 1), (1, 0)]$. Both $\{R_1\}$ and $\{R_2\}$ are in T_{TSI} (as they are respectively closed by \min and \max , which are 2-ary TSIs), but $\{R_1, R_2\}$ is not (R_1 forces $f(0, 1) = 0$ and R_2 forces $f(0, 1) = 1$ for every TSI polymorphism f), so T_{TSI} is not 1-Helly.

Finally, we conclude that Theorem 4 applies to T_{TSI} even for binary constraints. The same reasoning also applies to many other tractable classes, such as languages preserved by a near-unanimity ($f(y, x, \dots, x) = f(x, y, x, \dots, x) = \dots = f(x, \dots, x, y) = x$) or a Mal'tsev ($f(x, x, y) = f(y, x, x) = y$) polymorphism.

5 Related work

A very recent paper by Gaspers et al. [3] has independently investigated the same topic (parameterized complexity of strong backdoor detection for tractable semantic classes) and some of their results seem close to ours. In particular, one of their theorems (Theorem 5) is similar to our Proposition 2, but is less general as they assume the target class to be an union of atomic classes. They also study the case where r is bounded and k is the parameter, as we do, but their result (Theorem 6) is more specific and can be shown to be implied by our Theorem 4.

6 Conclusion

We have shown that finding small strong backdoors to tractable constraint languages is often hard. In particular, if the tractable class is a set of languages closed by an idempotent operation, or can be defined by arbitrary conjunctions and disjunctions of such languages, then finding a backdoor to this class is NP-hard even when all constraints have a fixed arity. Moreover, it is $\text{W}[2]$ -hard with respect to the backdoor size k .

When considering the larger parameter $k + r$, however, we have shown that strong backdoor detection is FPT provided that the target class is h -Helly for a constant h , that is, membership in this class can be decided by checking all h -tuples of relations. We then give a complete characterization of 1-Helly classes, and we use this result to show that any finite union of 1-Helly classes induces a backdoor problem FPT in $k + r$. Finally, we characterize another large family of tractable classes for which backdoor detection is $\text{W}[2]$ -hard for the parameter k even if r is fixed. This result can be used to derive hardness of backdoor detection for many known large tractable classes, provided they have certain natural properties (which we call value-renamability and domain-decomposability).

References

1. Williams, R., Gomes, C.P., Selman, B.: Backdoors to typical case complexity. In: Proceedings of the 18th International Joint Conference on Artificial Intelligence. IJ-CAI'03, San Francisco, CA, USA, Morgan Kaufmann Publishers Inc. (2003) 1173–1178
2. Nishimura, N., Ragde, P., Szeider, S.: Detecting Backdoor Sets with Respect to Horn and Binary Clauses. In: SAT. (2004)
3. Gaspers, S., Misra, N., Ordyniak, S., Szeider, S., Živný, S.: Backdoors into heterogeneous classes of sat and csp. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence. AAAI '14 (2014)
4. Gaspers, S., Szeider, S.: Strong Backdoors to Bounded Treewidth SAT. In: FOCS. (2013) 489–498
5. Feder, T., Vardi, M.Y.: The computational structure of monotone monadic smp and constraint satisfaction: A study through datalog and group theory. *SIAM Journal on Computing* **28**(1) (1998) 57–104
6. Bulatov, A.A.: Complexity of conservative constraint satisfaction problems. *ACM Trans. Comput. Logic* **12**(4) (July 2011) 24:1–24:66
7. Barto, L., Bulín, J.: Csp dichotomy for special polyads. *IJAC* **23**(5) (2013) 1151–1174
8. Barto, L., Kozik, M.: Constraint satisfaction problems solvable by local consistency methods. *J. ACM* **61**(1) (January 2014) 3:1–3:19
9. Jeavons, P., Cohen, D., Gyssens, M.: Closure properties of constraints. *J. ACM* **44**(4) (July 1997) 527–548
10. Jeavons, P., Cohen, D., Cooper, M.: Constraints, consistency, and closure. *Artificial Intelligence* **101** (1998) 101–1
11. Idziak, P., Markovic, P., McKenzie, R., Valeriote, M., Willard, R.: Tractability and learnability arising from algebras with few subpowers. *SIAM Journal on Computing* **39**(7) (2010) 3023–3037
12. Bulatov, A.A.: Combinatorial problems raised from 2-semilattices. *Journal of Algebra* **298**(2) (2006) 321 – 339
13. Dalmau, V., Pearson, J.: Closure functions and width 1 problems. In: Principles and Practice of Constraint Programming–CP99, Springer (1999) 159–173
14. Jeavons, P., Cooper, M.: Tractable constraints on ordered domains. *Artificial Intelligence* **79** (1995) 327–339
15. Barto, L., Kozik, M.: Constraint satisfaction problems solvable by local consistency methods. *Journal of the ACM (JACM)* **61**(1) (2014) 3
16. Kozik, M., Krokhin, A., Valeriote, M., Willard, R.: Characterizations of several maltsev conditions. preprint (2013)
17. Dilkina, B., Gomes, C.P., Sabharwal, A.: Tradeoffs in the complexity of backdoors to satisfiability: dynamic sub-solvers and learning during search. *Annals of Mathematics and Artificial Intelligence* (2014) 1–33
18. Bessière, C., Carbonnel, C., Hebrard, E., Katsirelos, G., Walsh, T.: Detecting and exploiting subproblem tractability. In: International Joint Conference on Artificial Intelligence (IJCAI), Beijing, China (August 2013)
19. Chen, J., Kanj, I.A., Xia, G.: Improved upper bounds for vertex cover. *Theor. Comput. Sci.* **411**(40–42) (September 2010) 3736–3756
20. Downey, R.G., Fellows, M.R.: Parameterized Complexity. Springer-Verlag (1999) 530 pp.
21. Karp, R.M.: Reducibility among combinatorial problems. In: Complexity of Computer Computations. (1972) 85–103