

Using a min-cut generalisation to go beyond Boolean surjective VCSPs[†]

Gregor Matl

Department of Informatics, Technical University of Munich, Germany
`matlg@in.tum.de`

Stanislav Živný*

Department of Computer Science, University of Oxford, UK
`standa.zivny@cs.ox.ac.uk`

Abstract

In this work, we first study a natural generalisation of the Min-Cut problem, where a graph is augmented by a superadditive set function defined on its vertex subsets. The goal is to select a vertex subset such that the weight of the induced cut plus the set function value are minimised. In addition, a lower and upper bound is imposed on the solution size. We present a polynomial-time algorithm for enumerating all near-optimal solutions of this Bounded Generalised Min-Cut problem.

Second, we apply this novel algorithm to surjective general-valued constraint satisfaction problems (VCSPs), i.e., VCSPs in which each label has to be used at least once. On the Boolean domain, Fulla, Uppman, and Živný [ACM ToCT'18] have recently established a complete classification of surjective VCSPs based on an unbounded version of the Generalised Min-Cut problem. Their result features the discovery of a new non-trivial tractable case called EDS that does not appear in the non-surjective setting.

As our main result, we extend the class EDS to arbitrary finite domains and provide a conditional complexity classification for surjective VCSPs of this type based on a reduction to smaller domains. On three-element domains, this leads to a complete classification of such VCSPs.

1 Introduction

Constraint satisfaction problems (CSPs) are fundamental computer science problems studied in artificial intelligence, logic (as model checking of the positive primitive fragment of first-order logic), graph theory (as homomorphisms between relational structures), and databases (as conjunctive queries) [15]. A vast generalisation of CSPs is that of general-valued CSPs (VCSPs) [25], see also [7]. Recent years have seen some remarkable progress on our understanding

[†]An extended abstract of this work appeared in *Proceedings of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019)* [22]. Stanislav Živný was supported by a Royal Society University Research Fellowship. The work was done while Gregor Matl was at the University of Oxford. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714532). The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

*Corresponding author.

of the computational complexity of CSPs and VCSPs, as will be discussed later in related work. We start with a few definitions to state existing as well as our new results.

We consider regular, surjective as well as lower-bounded VCSPs on the extended rationals $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$. An *instance* $I = (V, D, \phi_I)$ of either of these problems is given by a finite set of variables $V = \{x_1, \dots, x_n\}$, a finite set of labels D called the *domain*, and an *objective function* $\phi_I : D^n \rightarrow \overline{\mathbb{Q}}$. The objective function is of the form

$$\phi_I(x_1, \dots, x_n) = \sum_{i=1}^t w_i \cdot \gamma_i(\mathbf{x}_i),$$

where $t \in \mathbb{N}$ and, for each $1 \leq i \leq t$, $\gamma_i : D^{\text{ar}(\gamma_i)} \rightarrow \overline{\mathbb{Q}}$ is a *weighted relation* of arity $\text{ar}(\gamma_i) \in \mathbb{N}$, $w_i \in \mathbb{Q}_{\geq 0}$ is a *weight* and $\mathbf{x}_i \in V^{\text{ar}(\gamma_i)}$ is a tuple of variables from V called the *scope* of γ_i .

Regular, surjective and lower-bounded VCSPs differ only in their solution space, although this makes a big difference in complexity. If I is an instance of a regular VCSP, an *assignment* is a map $s : V \rightarrow D$ assigning a label from D to each variable. In the surjective setting, only a surjective map $s : V \rightarrow D$ is an assignment. For lower-bounded VCSPs, a fixed lower bound $l : D \rightarrow \mathbb{N}_0$ is provided and an assignment is a map $s : V \rightarrow D$ such that $|s^{-1}(d)| \geq l(d)$ for every label $d \in D$. In other words, a lower bound $l(d)$ on the number of occurrences of each label $d \in D$ is imposed. This is a generalisation of surjective VCSPs where the lower bound is always 1. (We are not aware of any previous work on lower-bounded VCSPs, which we introduce in this work.) The *value* of an assignment s is given by $\phi_I(s(x_1), \dots, s(x_n))$. An assignment is called *feasible* if its value is finite, and is called *optimal* if it is of minimal value among all assignments for the instance. The objective is to obtain an optimal assignment.

While finding an optimal assignment is NP-hard in general, valued constraint languages impose a natural restriction on the types of instances that are allowed. A *valued constraint language*, or simply a *language*, is a possibly infinite set of weighted relations. In this paper, we only consider languages of bounded arity, that is languages admitting a fixed upper bound on the arity of all weighted relations contained in them. Weighted relations in any VCSP instance will be stored explicitly.

We denote the class of regular VCSP instances with objective functions using only weighted relations from a language Γ by $\text{VCSP}(\Gamma)$. Similarly, $\text{VCSP}_s(\Gamma)$ is the class of surjective VCSP instances with weighted relations from Γ and, for some lower bound $l : D \rightarrow \mathbb{N}_0$, $\text{VCSP}_l(\Gamma)$ is the class of lower-bounded VCSP instances over Γ with bound l .

A language Γ is *globally tractable* if there is a polynomial-time algorithm for solving each instance of $\text{VCSP}(\Gamma)$, or *globally intractable* if $\text{VCSP}(\Gamma)$ is NP-hard. Analogously, Γ is *globally s -tractable* if there is a polynomial-time algorithm for $\text{VCSP}_s(\Gamma)$, or *globally s -intractable* if $\text{VCSP}_s(\Gamma)$ is NP-hard. And Γ is *globally l -tractable* if $\text{VCSP}_l(\Gamma)$ is solvable in polynomial time for every fixed lower bound $l : D \rightarrow \mathbb{N}_0$, or *globally l -intractable* if $\text{VCSP}_l(\Gamma)$ is NP-hard for at least one fixed lower bound $l : D \rightarrow \mathbb{N}_0$. Thus, global l -tractability implies global s -tractability, and global s -intractability implies global l -intractability.

The following examples show how well-studied variants of the MIN-CUT problem can be modelled in the VCSP frameworks we have defined.

Example 1 (*r -TERMINAL MIN-CUT*). Given a graph $G = (V, E)$ with non-negative edge weights $w : E \rightarrow \mathbb{Q}_{\geq 0}$ and designated terminal vertices $s_1, \dots, s_r \in V$, the *r -TERMINAL MIN-CUT* problem asks to partition V into subsets X_1, \dots, X_r such that $s_d \in X_d$ for all $d \in [r] := \{1, \dots, r\}$, while the accumulated weight of all edges going between distinct sets X_i and X_j is minimised. For $r = 2$, this problem is also known as the (s, t) -MIN-CUT problem.

We show how this problem can be represented as a regular VCSP. Let $\gamma_{r\text{-cut}}$ denote the binary weighted relation defined for $x, y \in [r]$ by $\gamma_{r\text{-cut}}(x, y) = 0$ if $x = y$ and $\gamma_{r\text{-cut}}(x, y) = 1$ otherwise. Furthermore, for each label $d \in [r]$, let ρ_d denote the constant relation given by $\rho_d(d) = 0$ and $\rho_d(x) = \infty$ for $d \neq x \in [r]$. Let $\Gamma_{r\text{-cut}} = \{\gamma_{r\text{-cut}}, \rho_1, \dots, \rho_r\}$.

Finding an optimal r -terminal cut is equivalent to solving the VCSP($\Gamma_{r\text{-cut}}$) instance $I = (V, [r], \phi)$ with objective function

$$\phi(x_1, \dots, x_n) = \rho_1(s_1) + \dots + \rho_r(s_r) + \sum_{(u,v) \in E} w(u, v) \cdot \gamma_{r\text{-cut}}(u, v).$$

To see this, observe that there is a correspondence between feasible assignments $s : V \rightarrow [r]$ and r -terminal cuts X_1, \dots, X_r by setting $X_d = \{v \in V : s(v) = d\}$. Hence, an optimal assignment induces an optimal cut.

The r -TERMINAL MIN-CUT problem can be solved in polynomial time if $r = 2$, but it is NP-hard for any $r \geq 3$ [9]. Since every VCSP($\Gamma_{r\text{-cut}}$) instance can also be reduced to an instance of the r -TERMINAL MIN-CUT problem, the language $\Gamma_{r\text{-cut}}$ is globally tractable if $r = 2$ and globally intractable for $r \geq 3$. ♣

Example 2 (r -WAY MIN-CUT). Without setting out any terminals, the r -WAY MIN-CUT problem asks to partition V into non-empty subsets X_1, \dots, X_r such that weight of the induced cut is minimised. Finding an optimal r -way min-cut is equivalent to solving the VCSP $_s(\{\gamma_{r\text{-cut}}\})$ instance $I = (V, [r], \phi)$ with objective function

$$\phi(x_1, \dots, x_r) = \sum_{(u,v) \in E} w(u, v) \cdot \gamma_{r\text{-cut}}(u, v).$$

The r -WAY MIN-CUT problem can be solved in polynomial time for every fixed integer r [13]. Since every VCSP $_s(\{\gamma_{r\text{-cut}}\})$ instance can be reduced to an r -WAY MIN-CUT problem as well, the language $\{\gamma_{r\text{-cut}}\}$ is globally s -tractable. ♣

For a fixed $l : D \rightarrow V$, VCSP $_l(\{\gamma_{r\text{-cut}}\})$ allows to model a generalisation of the r -WAY MIN-CUT problem where a partition X_1, \dots, X_r of V minimising the induced cut is sought under the condition that $|X_d| \geq l(d)$ for every $d \in D$. As far as we know, the complexity of both VCSP $_l(\{\gamma_{r\text{-cut}}\})$ and the lower-bounded r -WAY MIN-CUT problem is unknown.

Related Work Early results on CSPs include the fundamental results of Schaefer on Boolean CSPs [24] and of Hell and Nešetřil on graph CSPs [14]. The computational complexity of CSPs has drawn a lot of attention following the seminal paper of Feder and Vardi [11]. Using the algebraic approach [17, 4], the complexity of CSPs on finite domains was resolved in two independent papers by Bulatov [5] and Zhuk [29]. The computational complexity of the problem of minimising the number of unsatisfied constraints (and more generally rational-valued weighted relations) was obtained by Thapper and Živný in [28]. Finally, the computational complexity of general-valued CSPs on finite domains was obtained by the work of Kozik and Ochremiak [21] and Kolmogorov, Krokhin, and Rolínek [19].

Many constraint solvers allow not only constraints that apply locally to the variables specified as arguments, but also some sort of global constraints. In fact, the latter are the default representations in most constraint solvers [23]. Among VCSPs with global constraints studied from the complexity point of view are CSPs with global cardinality constraints, or

CCSPs, where it is specified how often exactly each label has to occur in an assignment. A dichotomy theorem for CCSPs on finite domains was established by Bulatov and Marx [6].

Surjective VCSPs, which can be seen as imposing a global constraint, have been studied by Fulla, Uppman, and Živný [12], following earlier results on CSPs by Creignou and Hébrard [8] and Bodirsky, Kára, and Martin [2]. Unfortunately, the algebraic approach that has proved pivotal in the understanding of the computational complexity of regular CSPs and VCSPs is not applicable in the surjective setting.

The following two facts are easy to show (see, e.g. [12]): (i) intractable languages are also s -intractable; (ii) a tractable language Γ is also s -tractable if Γ includes all constant relations. Consequently, new s -tractable languages can only occur (if at all) as subsets of tractable languages that do not contain all constant relations. The first example of such languages have been presented in [12]. In particular, the authors of [12] have identified languages on the Boolean domain that are *essentially a downset*, or EDS, as a new class of efficiently solvable problems and, in doing so, have provided a complexity classification of surjective VCSPs on the Boolean domain. Informally, a weighted relation γ is EDS if both the set of feasible tuples of γ and the set of optimal tuples of γ are essentially downsets. Here a relation is called essentially a downset if it can be written as a conjunction of downsets and binary equality relations. Equivalently, a relation is essentially a downset if it admits a binary polymorphism $\text{sub}(x, y) = \min(x, 1 - y)$. A finite language is EDS if every weighted relation in Γ is EDS. The definition for infinite languages is more complicated. We give a formal definition of the EDS class in Section 3 and refer the reader to [12] for more details.

The tractability result of EDS languages is based on the Generalised Min-Cut (GMC) problem for graphs, also introduced in [12]. In a GMC instance, the goal is to find a non-trivial subset of the vertices such that the weight of the induced cut and a superadditive set function are minimised simultaneously. In particular, the following has been shown in [12]. Firstly, the objective function of surjective VCSPs that are EDS can be approximated by an instance of the GMC problem. Secondly, there is a polynomial-time algorithm to enumerate all solutions to the GMC problem that are optimal up to a constant factor. These two together give an efficient algorithm for surjective VCSPs that are EDS.

Contributions This paper extends the class EDS to arbitrary finite domains. We introduce a class SIM of languages that exhibit properties *similar to a Boolean language*. Based on this class, we define the class SEDS of languages *similar to EDS* as a natural extension of EDS and classify languages from this extension based on two criteria. Firstly, we give a subclass SDS, or *similar to a downset*, of SEDS that guarantees global ℓ -tractability without additional requirements. Secondly, we prove that the complexity of lower-bounded VCSPs over any remaining SEDS languages is equivalent to the complexity over a particular language on a smaller domain, which can be constructed by including all possible ways to assign a certain label. This is illustrated in Figure 1 (left), where we use the notation $\text{fix}(\Gamma)$, formally defined in Section 3.2. Informally, for a language Γ defined on domain D that includes the label 0, $\text{fix}(\Gamma)$ is the language on domain $D \setminus \{0\}$ obtained by including, for every weighted relation $\gamma \in \Gamma$ of arity n and a subset U of the arguments of γ , the weighted relation $\text{fix}_U[\gamma]$, which is a weighted relation on $D \setminus \{0\}$ of arity $n - |U|$ defined as the restriction of γ that fixes the label 0 to all arguments in U .

One implication of our results is a dichotomy theorem for lower-bounded VCSPs on the Boolean domain; every Boolean language is either globally ℓ -tractable or globally ℓ -intractable.

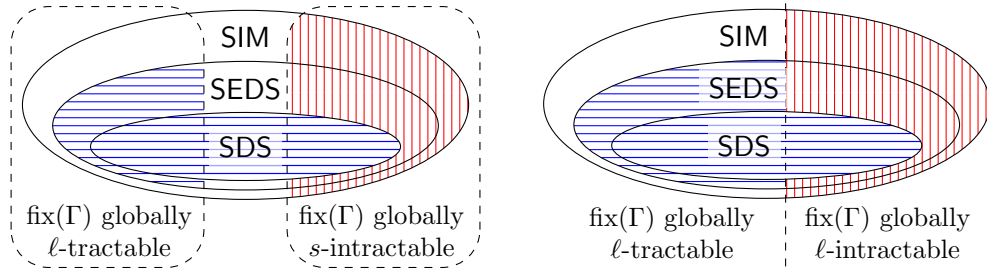


Figure 1: Classification of SEDS languages on arbitrary finite domains (left) and on three-element domains (right). A language Γ is globally ℓ -tractable when marked by horizontal (blue) lines and globally s -intractable when marked by vertical (red) lines, depending on the language $\text{fix}(\Gamma)$ on a smaller domain. (Recall that global s -intractability implies global ℓ -intractability.) In case of three-element domains, the Boolean language $\text{fix}(\Gamma)$ is either globally ℓ -tractable or globally ℓ -intractable, while this is not known for larger domains.

Although lower-bounded VCSPs are more general than surjective VCSPs, this classification coincides with the dichotomy theorem for surjective VCSPs given by [12].

In addition, combining our reduction of SEDS languages to a smaller domain and the dichotomy theorem for the Boolean domain leads to a classification of all SEDS languages on three-element domains with respect to ℓ -tractability, which is featured on the right-hand side of Figure 1.

The foundation of our results is an extension of the Generalised Min-Cut problem that might be of independent interest. Given integers $p, q \in \mathbb{N}_0$, a graph with non-negative edge weights and a superadditive set function defined on its vertices, the goal in the Bounded Generalised Min-Cut problem is, just like in the GMC problem, to find a subset of the vertices such that the sum of the induced cut and the superadditive set function evaluated on it are minimal among all possible solutions. The solution space, however, is restricted to subsets containing at least q and at most all but p vertices.

If an optimal solution has value 0, there can be exponentially many optimal solution, e.g. when there are no edges and the superadditive function always evaluates to 0. Our main algorithmic result is that, for all instances with non-zero optimal value and for any constant bounds $p, q \in \mathbb{N}_0$, all solutions that are optimal up to a constant factor can be enumerated in polynomial time (and thus, in particular, there are only polynomially many of them).

We finish with two remarks on, as far as we can tell, unrelated work. First, it is natural to consider Karger’s elegant (randomised) min-cut algorithm [18], which also allows to enumerate (polynomially many) near-optimal cuts, and try to adapt it to the newly introduced Bounded Generalised Min-Cut problem. Despite trying, we do not see any way of doing it.¹ Moreover, we only know how to establish our tractability results on surjective VCSPs by a reduction to the Bounded Generalised Min-Cut problem that includes that superadditive function, but that one fails many properties required by Karger’s algorithm. (For instance, superadditive functions are not necessarily submodular.) Second, it is notationally convenient to go back and forth between weighted relations (on a domain of size $k + 1$) and k -set functions, as we

¹It appears to be an issue that the superadditive set function is evaluated only for the solution set, while the set of remaining vertices may exhibit an excessively large set function value even in an optimal solution. That makes it implausible to think a local criterion for edge contractions could incorporate the superadditive set function in a suitable manner, i.e. somehow preventing the set function value from getting too large.

will explain in Section 3 and use throughout the paper. We do not see a connection (suggested by an anonymous reviewer of the extended abstract of this work [22]) to the characterisation of arc consistency via set polymorphisms [11, 10], which are properties of (weighted) relations but not their equivalent description. More generally, we do not know whether our tractability result could be established using recent work on consistency methods for CSPs [1] or LP relaxations for VCSPs [20, 27].

Organisation We will proceed in the following manner. Section 2 gives a polynomial-time algorithm for enumerating all near-optimal optimal solutions of the Bounded Generalised Min-Cut problem. In Section 3, we extend the notion of EDS to larger domains. A classification of languages from this extension is presented in Section 4. Section 5 provides a dichotomy theorem for lower-bounded VCSPs on the Boolean domain.

2 The Bounded Generalised Min-Cut Problem

We begin by presenting our algorithm for the Bounded Generalised Min-Cut problem. The problem is based on the notion of superadditive set functions, which we define first.

Definition 3. A *set function* on a finite set V is a function $f : 2^V \rightarrow \overline{\mathbb{Q}}$ defined on subsets of V ; it is *normalised* if it satisfies $f(\emptyset) = 0$ and $f(X) \geq 0$ for all $X \subseteq V$.

A set function f on V is *increasing* if it is normalised and $f(X) \leq f(Y)$ for all $X \subseteq Y \subseteq V$. It is *superadditive* if it is normalised and, for all disjoint $X, Y \subseteq V$, it holds that

$$f(X) + f(Y) \leq f(X \cup Y). \quad (\text{SUP})$$

Since equation (SUP) implies that $f(X) \leq f(X) + f(Y \setminus X) \leq f(Y)$ for all $X \subseteq Y \subseteq V$, every superadditive set function must also be increasing.

Definition 4. For $p, q \in \mathbb{N}_0$, the *Bounded Generalised Min-Cut* problem with lower bound q and upper bound p is denoted by GMC_q^p .

A GMC_q^p *instance* h is given by an undirected graph $G = (V, E)$ with edge weights $w : E \rightarrow \mathbb{Q}_{\geq 0} \cup \{\infty\}$ and an oracle defining a superadditive set function f on V . For $X \subseteq V$, let $w(X) = \sum_{|\{u,v\} \cap X|=1} w(\{u,v\})$ denote the weight of the cut induced by X .

A *solution* of instance h is any set $X \subseteq V$ such that $|X| \geq q$ and $|X| \leq |V| - p$. The objective is to minimise the value $h(X) = f(X) + w(X)$. A solution X is *optimal* if the value $h(X)$ is minimal among all solutions for this instance. We denote the value of an optimal solution by λ . For any $\alpha \geq 1$, a solution X is α -*optimal* if $h(X) \leq \alpha\lambda$.

The Generalised Min-Cut problem, simply denoted by GMC, is the Bounded Generalised Min-Cut problem with lower and upper bound 1. All α -optimal solutions of a GMC instance can be enumerated in polynomial time according to [12, Theorem 5.11], which we restate here.

Theorem 5 ([12]). *For any instance h of the GMC problem on n vertices with optimal value $0 < \lambda < \infty$ and any constant $\alpha \in \mathbb{N}$, the number of α -optimal solutions is at most $n^{20\alpha-15}$. There is an algorithm that finds all of them in polynomial time.*

We will assume that all edges are strictly positive-valued, as they can be ignored otherwise. Similarly to [12, Lemma 53] for the GMC problem, we can easily detect and solve the problem when $\lambda = 0$ or $\lambda = \infty$.

Lemma 6. *For any $p, q \in \mathbb{N}_0$, where q is a constant, a polynomial-time algorithm can determine whether the optimal value of a GMC_q^p instance h on a graph $G = (V, E)$ is $\lambda = 0$, $1 < \lambda < \infty$ or $\lambda = \infty$. In case $\lambda = 0$, it can provide an optimal solution.*

Proof. First, we assume $\lambda = 0$. Consider some optimal solution $X \subseteq V$. Then $h(X) = 0$ implies that X cannot cut any edges and, hence, must be a union of connected components $C_1, \dots, C_k \subseteq V$ from G for some $k \in \mathbb{N}$. The union $Y = C_1 \cup \dots \cup C_{\min(k, q)}$ of up to q of those components must still satisfy $q \leq |Y| \leq |X| \leq |V| - p$ and $h(Y) = 0$, because the superadditive set function f is increasing and $Y \subseteq X$. Consequently, an algorithm can check all $O(n^q)$ combinations of up to q components from G in order to find the solution Y . And vice versa, if no such solution of value 0 is found, it can be concluded that $\lambda > 0$.

Similarly, to probe whether $\lambda = \infty$, we consider those vertices that are connected by infinite-weight edges as components, because any finite-valued solution cannot cut those edges. It is then sufficient to check all $O(n^q)$ combinations of up to q components to see whether a finite-valued solution exists. Otherwise, if all these candidates have infinite value when they are comprised of q or more vertices, any solution that does not cut any infinity-edges must be a superset of one of these candidates and therefore have infinite value as well due to the increasing nature of f . \square

Consequently, our goal is to provide a polynomial-time algorithm for enumerating near-optimal solutions in the case that the optimal value is both positive and finite. Before doing so, we give two auxiliary lemmas based on [12, Lemma 5.6] and [12, Lemma 5.10].

Lemma 7. *For any $p, q \in \mathbb{N}_0$, any GMC_q^p instance h on a graph $G = (V, E)$ and any subset $V' \subseteq V$, there is a GMC_q^p instance h' on the induced subgraph $G[V']$ that preserves the objective value of all solutions $X \subseteq V'$. In particular, any α -optimal solution X of h such that $X \subseteq V'$ is α -optimal for h' as well.*

Proof. Edges with exactly one endpoint in V' need to be taken into account separately because they do not appear in the induced subgraph. We accomplish that by defining the new set function f' by

$$f'(X) = f(X) + \sum_{u \in X} \sum_{v \in V \setminus V'} w(u, v)$$

for all $X \subseteq V'$. By the construction, f' is superadditive, and the objective value $h'(X)$ for any solution $X \subseteq V'$ equals $h(X)$.

Note that the minimum objective value for h' is greater than or equal to the minimum objective value for h . Therefore, any solution $X \subseteq V'$ that is α -optimal for h is also α -optimal for h' . \square

When a solution of some bounded GMC instance is split into two parts, the next lemma gives a bound on the values of these parts based on edges involved in the split.

Lemma 8. *Let h be a GMC_q^p instance over vertices V with optimal value λ and let $X, Y \subseteq V$ such that $h(X) \leq \alpha\lambda$ and $w(Y) \leq \beta\lambda$ for some $\alpha \geq 1$ and $\beta \geq 0$. Then it holds*

$$h(X \setminus Y) + h(X \cap Y) \leq (\alpha + 2\beta)\lambda.$$

Proof. It is well-known and can easily be verified that the cut function w is posimodular, meaning that $w(A) + w(B) \geq w(A \setminus B) + w(B \setminus A)$ for all $A, B \subseteq V$.

As a consequence, we have

$$\begin{aligned} w(X) + w(Y) &\geq w(X \setminus Y) + w(Y \setminus X) \\ w(Y) + w(Y \setminus X) &\geq w(X \cap Y) + w(\emptyset), \end{aligned}$$

and hence,

$$w(X) + 2w(Y) \geq w(X \setminus Y) + w(X \cap Y).$$

By superadditivity of f , it holds $f(X) \geq f(X \setminus Y) + f(X \cap Y)$. The claim then follows from the fact that $f(X) + w(X) + 2w(Y) \leq (\alpha + 2\beta)\lambda$. \square

With these preparations on hand, we now proceed with our main algorithmic result.

Theorem 9. *For some constant $q \geq 2$, let h be a GMC_q^1 instance on a graph $G = (V, E)$ of size $n = |V|$ with optimal value $0 < \lambda < \infty$. Let $Y \cup Z = V$ be a partition of V and let $Y_1 \cup \dots \cup Y_k = Y$ for some $k \in \mathbb{N}_0$ be a partition of Y satisfying $0 < |Y_i| < q$ and $h(Y_i) \leq \frac{\lambda}{3q}$ for all $1 \leq i \leq k$.*

Then for every constant $\alpha \geq 1$, at most $\frac{|Z|}{n} \cdot n^{\tau(q, \alpha)}$ α -optimal solutions $X \subseteq V$ of h satisfy $|X \cap Y| < q$, where $\tau(q, \alpha) = 60q\alpha + 41q + 7$. These solutions can all be enumerated in polynomial time.

Note that with $Y = \emptyset$ and $Z = V$, this theorem states for any GMC_q^1 instance that the number of α -optimal solutions is bounded by $n^{\tau(q, \alpha)}$.

Proof. Proof by induction over $n + \frac{|Z|}{n+1}$; that is, induction primarily over n and, for equal values of n , also over $|Z|$. For $n \leq q$ or $Z = \emptyset$, there are no solutions of the described form and hence, the statement holds.

Now, fix some $n > q$, some GMC_q^1 instance h on a graph $G = (V, E)$ of size n with optimum value $0 < \lambda < \infty$ and partitions $Y \cup Z = V$ and $Y_1 \cup \dots \cup Y_k = Y$ as described. By the induction hypothesis, we can assume that the theorem holds for every graph of size $n' < n$ as well as for every partition $\tilde{Y} \cup \tilde{Z} = V$ of graph G satisfying $|\tilde{Z}| < |Z|$.

To simplify matters, we can replace any infinite edge weights in G with a large value ($\alpha \cdot (1 + f(V) + \sum_{w(u,v) < \infty} w(u, v))$ works) without affecting any of our assumptions or the set of α -optimal solutions we are looking for. Thus, we will subsequently assume that all edges are finite-valued.

According to Lemma 7, there exists a GMC_q^1 instance h_Z on the induced subgraph $G[Z]$ that preserves the objective value of every solution $X' \subseteq Z$ with respect to h . In the following, we treat h_Z as a GMC instance (i.e. with lower bound 1). Let λ_Z denote the optimal value of h_Z . We can assume $\lambda_Z < \infty$ because otherwise, due to the absence of infinite-weight edges in G and the superadditivity of f , no finite-valued solution $X \subseteq V$ of h can satisfy $X \cap Z \neq \emptyset$. Let $Y_{k+1} \subseteq Z$ be an optimal solution of h_Z , i.e. $h_Z(Y_{k+1}) = \lambda_Z$.

If $h(Y_{k+1})$ is sufficiently large, we show that it is essentially sufficient to enumerate GMC solutions of $G[Z]$ up to a constant factor. For small $h(Y_{k+1})$, our strategy will be to reduce the problem to the partition $Y' \cup Z' = V$, where $Y' = Y_1 \cup \dots \cup Y_k \cup Y_{k+1}$ and $Z' = Z \setminus Y_{k+1}$. This approach is outlined in Figure 2.

Consider any α -optimal solution $X \subseteq V$ of h satisfying $|X \cap Y| < q$. For some integer t , let i_1, \dots, i_t denote indices such that $X \cap Y = X \cap (Y_{i_1} \cup \dots \cup Y_{i_t})$, i.e. such that X has vertices only in Y_{i_1}, \dots, Y_{i_t} and Z . Since $|X \cap Y| < q$, we require that $t < q$. Let $U = Y_{i_1} \cup \dots \cup Y_{i_t}$.




Figure 2: Given a partition $V = Y \cup Z$ with $Y = Y_1 \cup \dots \cup Y_k$ of the vertices of a GMC_q^1 instance h , we want to find every solution X such that $h(X) \leq \alpha\lambda$ and $|X \cap Y| < q$. Consider the GMC instance h_Z on $G[Z]$ with optimal solution Y_{k+1} . If $h(Y_{k+1}) \geq \frac{\lambda}{3q}$, $X \cap Z$ must be a near-optimal solution of h (left, Case 1). Otherwise, we apply the induction hypothesis either on the subgraph $G[Z']$, where $Z' = Z \setminus Y_{k+1}$ (right, Case 2a), or on the partition $V = Z' \cup (Y_1 \cup \dots \cup Y_{k+1})$ (Case 2b).

Case 1: If $\lambda_Z \geq \frac{\lambda}{3q}$, we aim to bound the value $h(X \cap Z)$ relative to λ_Z . Since $w(Y_i) \leq h(Y_i) \leq \frac{\lambda}{3q}$ for every $1 \leq i \leq k$ by assumption, it must hold that

$$\begin{aligned} w(U) &= \sum_{|\{u,v\} \cap U|=1} w(\{u,v\}) \leq \sum_{j=1}^t \left(\sum_{|\{u,v\} \cap Y_{i_j}|=1} w(\{u,v\}) \right) \\ &= \sum_{j=1}^t w(Y_{i_j}) \leq t \cdot \frac{\lambda}{3q} < q \cdot \frac{\lambda}{3q} = \frac{\lambda}{3}. \end{aligned}$$

According to Lemma 8 with $\beta = \frac{1}{3}$, it follows that

$$h(X \setminus U) + h(X \cap U) \leq \left(\alpha + \frac{2}{3} \right) \lambda,$$

and in particular, since $X \cap Z = X \setminus U$, we have

$$h(X \cap Z) \leq \left(\alpha + \frac{2}{3} \right) \lambda.$$

Assuming $\lambda_Z \geq \frac{\lambda}{3q}$, we can limit the value $h(X \cap Z)$ relative to λ_Z by

$$\left(\alpha + \frac{2}{3} \right) \lambda \leq \left(\alpha + \frac{2}{3} \right) \cdot 3q\lambda_Z = (3q\alpha + 2q)\lambda_Z.$$

Given that $X \cap Z \neq \emptyset$, the above equation implies that if $X \cap Z \subsetneq Z$, then $X \cap Z$ is a $(3q\alpha + 2q)$ -optimal solution of the GMC instance h_Z . According to Theorem 5, there are at most

$$n^{20\lceil 3q\alpha + 2q \rceil - 15} \leq n^{20(3q\alpha + 2q + 1) - 15} = n^{60q\alpha + 40q + 5}$$

$(3q\alpha + 2q)$ -optimal solutions of GMC instance h_Z , which can all be enumerated in polynomial time. Pairing up these choices for $X \cap Z$, in addition to the possibility $X = Z$, with the at most $\sum_{i=0}^{q-1} \binom{n}{i} \leq \sum_{i=0}^{q-1} n^i \leq \sum_{i=0}^{q-1} \left(\frac{1}{2}\right)^{q-i} n^q \leq n^q$ sets of up to $q - 1$ vertices from Y gives at most

$$\left(n^{60q\alpha + 40q + 5} + 1 \right) \cdot n^q \leq n^{60q\alpha + 41q + 6} = \frac{1}{n} \cdot n^{\tau(q,\alpha)} \leq \frac{|Z|}{n} \cdot n^{\tau(q,\alpha)} \quad (\text{Case 1})$$

overall choices for X in this case, as required.

Case 2a: Now, let's assume that $\lambda_Z \leq \frac{\lambda}{3q}$ and furthermore that $|X \cap Y'| \geq q$, where $Y' = Y \cup Y_{k+1}$. Then it must hold $w(Y_{k+1}) \leq \lambda_Z \leq \frac{\lambda}{3q}$. Let $U' = Y_{i_1} \cup \dots \cup Y_{i_t} \cup Y_{k+1}$ so that it holds $X \cap Y' \subseteq U'$. Similar to the previous case, we can bound $w(U')$ by

$$w(U') \leq w(Y_{i_1}) + \dots + w(Y_{i_t}) + w(Y_{k+1}) \leq (t+1) \cdot \frac{\lambda}{3q} \leq q \cdot \frac{\lambda}{3q} = \frac{\lambda}{3}.$$

According to Lemma 8 with $\beta = \frac{1}{3}$, it must then hold that

$$h(X \setminus U') + h(X \cap U') \leq \left(\alpha + \frac{2}{3}\right) \lambda.$$

Assuming that $|X \cap Y'| \geq q$, the set $X \cap U' = X \cap Y'$ is a solution of h and must have value $h(X \cap U') \geq \lambda$. For $Z' = Z \setminus Y_{k+1}$, it therefore holds that

$$h(X \cap Z') = h(X \setminus U') \leq \left(\alpha + \frac{2}{3}\right) \lambda - h(X \cap U') \leq \left(\alpha - \frac{1}{3}\right) \lambda.$$

Let $h_{Z'}$ denote the GMC_q^1 instance on the induced subgraph $G[Z']$ that preserves the value of h as detailed in Lemma 7. Unless $|X \cap Z'| < q$ or $X \cap Z' = Z'$, the set $X \cap Z'$ is an $(\alpha - \frac{1}{3})$ -optimal solution of $h_{Z'}$ (in particular, this case can be ignored when $\alpha < \frac{4}{3}$). By applying the induction hypothesis on $h_{Z'}$ with the trivial partition $\emptyset \cup Z' = Z'$, it follows that the number of $(\alpha - \frac{1}{3})$ -optimal solutions is at most

$$\frac{|Z'|}{|Z|} \cdot (|Z'|)^{\tau(q, \alpha - \frac{1}{3})} \leq n^{\tau(q, \alpha - \frac{1}{3})}.$$

In addition, there are at most $\sum_{i=0}^{q-1} \binom{n}{i} \leq n^q$ subsets of Z' that have size less than q . Accounting also for the possibility $X \cap Z' = Z'$, there are at most

$$n^{\tau(q, \alpha - \frac{1}{3})} + n^q + 1 \leq 3n^{\tau(q, \alpha - \frac{1}{3})} \leq n^{\tau(q, \alpha - \frac{1}{3}) + 1}$$

choices for $X \cap Z'$ in this case.

Next, we limit the number of choices for $X \cap Y'$. Since X contains at most $q-1$ vertices from Y (less than n^q choices) and since Y_{k+1} contains at most $q-1$ vertices (at most 2^{q-1} choices), the number of possible choices for $X \cap Y'$ is limited by

$$n^q \cdot 2^{q-1} \leq n^{2q}.$$

Pairing up each possible choice for $X \cap Z'$ with each choice for $X \cap Y'$ gives a total of at most

$$n^{\tau(q, \alpha - \frac{1}{3}) + 1} \cdot n^{2q} = n^{\tau(q, \alpha - \frac{1}{3}) + 2q + 1} \leq \frac{1}{n} \cdot n^{\tau(q, \alpha)} \quad (\text{Case 2a})$$

solutions, where the last inequality follows from the fact that

$$\tau(q, \alpha) - \tau\left(q, \alpha - \frac{1}{3}\right) = 60q \cdot \frac{1}{3} \geq 2q + 2.$$

Case 2b: Finally, let's assume that $\lambda_Z \leq \frac{\lambda}{3q}$ and that $|X \cap Y'| < q$. Since $h_Z(Y_{k+1}) = \lambda_Z < \lambda$ implies $|Y_{k+1}| < q$, we can apply the induction hypothesis for instance h with the partition $Y' \cup Z' = V$ to limit the number of choices for X . Consequently, this number is at most

$$\frac{|Z'|}{n} \cdot n^{\tau(q,\alpha)} \leq \frac{|Z| - 1}{n} \cdot n^{\tau(q,\alpha)}. \quad (\text{Case 2b})$$

Summing up the bounds for Case 2a and Case 2b, the overall number of choices for X if $\lambda_Z \leq \frac{\lambda}{3q}$ is bounded by

$$\frac{1}{n} \cdot n^{\tau(q,\alpha)} + \frac{|Z| - 1}{n} \cdot n^{\tau(q,\alpha)} = \frac{|Z|}{n} \cdot n^{\tau(q,\alpha)}.$$

This proves the upper bound of $\frac{|Z|}{n} \cdot n^{\tau(q,\alpha)}$ solutions of the described form.

A polynomial-time algorithm to enumerate all such solutions follows almost immediately from these calculations. Given that λ might not be known beforehand, we simply check both Case 1 and Case 2.

Note that the induction hypothesis is used only in Case 2a, where all $(\alpha - \frac{1}{3})$ -optimal solutions of GMC_q^1 instance $h_{Z'}$ with partition $\emptyset \cup Z'$ need to be computed, and in Case 2b, where all α -optimal solutions of GMC_q^1 instance h_Z with partition $Y' \cup Z'$ need to be computed. It is straightforward to verify that the algorithmic complexity of all required operations except for these two recursive calls can be bounded by some polynomial $\text{poly}(n)$. We show by induction that $T_\alpha(n, Z) = 3\alpha n^{3\alpha} \cdot |Z| \cdot \text{poly}(n)$ is an upper bound on the overall complexity.

$$\begin{aligned} T_\alpha(n, Z) &\leq \text{poly}(n) + T_{\alpha - \frac{1}{3}}(|Z'|, Z') + T_\alpha(n, Z') \\ &\leq \text{poly}(n) + (3\alpha - 1)n^{3\alpha} \cdot \text{poly}(n) + 3\alpha n^{3\alpha} \cdot (|Z| - 1) \cdot \text{poly}(n) \\ &\leq 3\alpha n^{3\alpha} \cdot |Z| \cdot \text{poly}(n) \quad \square \end{aligned}$$

Corollary 10. *For any $p, q \in \mathbb{N}_0$ and $\alpha \geq 1$, where q and α are constants, and for any GMC_q^p instance h with optimal value $0 < \lambda < \infty$, all α -optimal solutions can be enumerated in polynomial time.*

Proof. Let $h = f + w$ be a GMC_q^p instance with $0 < \lambda < \infty$. First, we assume that $p \geq 1$ and $q \geq 2$. The superadditive set function

$$f'(X) = \begin{cases} \infty & \text{if } |X| > |V| - p \\ f(X) & \text{otherwise} \end{cases}$$

defines a GMC_q^1 instance $h' = f' + w$ where every solution $X \subseteq V$ of size $|X| > |V| - p$ has infinite value so that the set of finite-valued solutions and their values are identical for h and h' . Therefore, it is sufficient to enumerate all α -optimal solutions of h' , which can be accomplished in polynomial time according to Theorem 9

If $p = 0$ or $q < 2$, there are up to $|V| + 2$ additional solutions that can all be checked in polynomial time. \square

3 Extending EDS to Larger Domains

In this section, we formally introduce the classes SIM, SEDS and SDS. In order to simplify our notation, we will subsequently always consider the $(k + 1)$ -element domain $D = \{0, 1, \dots, k\}$ for some integer k . Any other domain of size $k + 1$ can simply be relabelled without affecting its properties. One label from the domain will play a special role; without loss of generality (due to relabellings), it will be 0.

3.1 k -Set Functions

It will be convenient to go back and forth between weighted relations and k -set functions, which is, subject to a minor technical assumption, always possible.

Definition 11. Let $k \in \mathbb{N}$ and let V be a finite set. A k -set function on V is a function $f : (k + 1)^V \rightarrow \overline{\mathbb{Q}}$ defined on k -tuples of pairwise disjoint subsets of V . A k -set function f over V is *normalised* if it satisfies $f(\emptyset, \dots, \emptyset) = 0$ and $f(X_1, \dots, X_k) \geq 0$ for all disjoint $X_1, \dots, X_k \subseteq V$.

Note that a 1-set function is simply a *set function* as defined in Section 2. The correspondence between weighted relations and k -set functions is formalised by the next definition.

Definition 12. Let γ be an n -ary weighted relation on the $(k + 1)$ -element domain $D = \{0, 1, \dots, k\}$, and let f be the k -set function on $V = [n]$ that is defined for disjoint sets $X_1, \dots, X_k \subseteq V$ by $f(X_1, \dots, X_k) = \gamma(\mathbf{x})$, where the i -th coordinate of \mathbf{x} is given by $x_i = d$ if $i \in X_d$ for some $0 \neq d \in D$ and $x_i = 0$ otherwise. Then γ *corresponds* to f .

Furthermore, we say that γ *corresponds under normalisation* to a k -set function if $\gamma(\mathbf{0}^n) < \infty$ and $\gamma(\mathbf{0}^n) \leq \gamma(\mathbf{x})$ for all $\mathbf{x} \in D^n$. In this case, the k -set function corresponding under normalisation to γ is the normalised k -set function corresponding to $\gamma - \gamma(\mathbf{0}^n)$, i.e. the weighted relation with offset such that the assignment $\mathbf{0}^n$ evaluates to 0.

According to this definition, there is a unique k -set function corresponding to every weighted relation on the $(k + 1)$ -element domain, and vice versa. Furthermore, assuming that $\gamma(\mathbf{0}^n) < \infty$, a weighted relation γ corresponds under normalisation to a k -set function precisely if it admits multimorphism $\langle c_0 \rangle$, which we will formally define in Section 5.

The next definition states when a k -set function is approximated by a (1-)set function. This approximation will serve as central tool in order to bring the structure of languages from larger domains essentially down to a Boolean domain.

Definition 13. Let f be a k -set function and g a set function on V . We say that g *α -approximates* f if, for all disjoint $X_1, \dots, X_k \subseteq V$, it holds that

$$g(X_1 \cup \dots \cup X_k) \leq f(X_1, \dots, X_k) \leq \alpha \cdot g(X_1 \cup \dots \cup X_k).$$

3.2 Fixing a Label: Reduced Languages

Reducing a language to a smaller domain by fixing the occurrences of label 0, as defined subsequently, will become a central tool in our classification.

Definition 14. Let γ be a weighted relation on domain D of arity n and let $U \subseteq [n]$. Then $\text{fix}_U[\gamma]$ is the weighted relation on domain $D^* = D \setminus \{0\}$ of arity $m = n - |U|$ defined for $x_1, \dots, x_m \in D^*$ by

$$\text{fix}_U[\gamma](x_1, \dots, x_m) = \gamma(y_1, \dots, y_n), \quad \text{where } y_i = \begin{cases} 0 & \text{if } i \in U \\ x_{|[i] \setminus U|} & \text{otherwise.} \end{cases}$$

In other words, $\text{fix}_U[\gamma]$ takes an assignment from domain D^* to all variables except for those with index in U , and evaluates it through γ by assigning label 0 to the remaining variables. In Definition 15, we generalise this concept in order to express the language that is generated by fixing every possible assignment of label 0.

Definition 15. Let Γ be a language on domain D . For any $\gamma \in \Gamma$, let $\text{fix}(\gamma)$ denote the set $\{\text{fix}_U[\gamma] : U \subseteq [\text{ar}(\gamma)]\}$ generated by fixing any possible subset of variables to label 0. We define the language $\text{fix}(\Gamma)$ on domain $D^* = D \setminus \{0\}$ by $\text{fix}(\Gamma) = \bigcup_{\gamma \in \Gamma} \text{fix}(\gamma)$.

3.3 Extending EDS to Larger Domains

The class EDS, or *essentially a downset*, has been introduced in [12] for the Boolean domain.

Definition 16. For any $\alpha \geq 1$, a normalised set function f on V is α -EDS if, for all $X, Y \subseteq V$, it holds that

$$f(X \setminus Y) \leq \alpha \cdot (f(X) + f(Y)). \quad (\text{EDS})$$

A weighted relation is α -EDS if it corresponds under normalisation to a set function that is α -EDS. Moreover, a language Γ is EDS if there is some $\alpha \geq 1$ such that every weighted relation $\gamma \in \Gamma$ is α -EDS.

Fulla et al. showed [12] that EDS languages are globally s -tractable. We improve upon this result by proving that such languages are in fact globally ℓ -tractable, and we extend the idea of being essentially a downset to larger domains through the classes SIM, SEDS and SDS.

Intuitively, a language is SIM, or *similar to a Boolean language*, if it can be approximated by a language over the Boolean domain using Definition 13. More precisely, for each weighted relation, the value of any two assignments that assign label 0 to the same set of variables must be equal up to a constant factor. This way, when disregarding constant factors, all non-zero labels can be treated as a single one, leading us essentially to the Boolean domain.

Definition 17. Let f be a normalised k -set function on set V . For any $\alpha \geq 1$, f is called α -SIM if, for all disjoint $X_1, \dots, X_k \subseteq V$ and all disjoint $Y_1, \dots, Y_k \subseteq V$ such that $X_1 \cup \dots \cup X_k = Y_1 \cup \dots \cup Y_k$, it holds that

$$f(X_1, \dots, X_k) \leq \alpha \cdot f(Y_1, \dots, Y_k). \quad (\text{SIM})$$

A weighted relation is α -SIM if it corresponds under normalisation to a k -set function that is α -SIM. Moreover, a language Γ is SIM if there is some $\alpha \geq 1$ such that every weighted relation $\gamma \in \Gamma$ is α -SIM.

Note that every normalised set function is 1-SIM. Hence, EDS is a subclass of SIM. Going beyond the Boolean domain, the class SEDS of languages *similar to EDS* arises as a natural generalisation of EDS. Intuitively, SEDS contains precisely those languages that can be approximated by EDS languages.

Definition 18. For any $\alpha \geq 1$, a normalised k -set function f on V is α -SEDS if it is α -SIM and, for all disjoint $X_1, \dots, X_k \subseteq V$ and all disjoint $Y_1, \dots, Y_k \subseteq V$, it holds that

$$f(X_1 \setminus Y_1, \dots, X_k \setminus Y_k) \leq \alpha \cdot (f(X_1, \dots, X_k) + f(Y_1, \dots, Y_k)). \quad (\text{SEDS})$$

A weighted relation is α -SEDS if it corresponds under normalisation to a k -set function that is α -SEDS. Moreover, a language Γ is SEDS if there is some $\alpha \geq 1$ such that every weighted relation $\gamma \in \Gamma$ is α -SEDS.

The class SDS, or *similar to a downset*, imposes a stricter requirement than SEDS. When any arguments of a weighted relation are changed to label 0, the value should decrease, stay equal or increase by at most a constant factor. Intuitively, weighted relations of these languages can be approximated by increasing set functions.

Definition 19. For any $\alpha \geq 1$, a normalised k -set function f on V is α -SDS if it is α -SIM and in addition, for all disjoint $X_1, \dots, X_k, Y_1, \dots, Y_k \subseteq V$, it holds that

$$f(X_1, \dots, X_k) \leq \alpha \cdot f(X_1 \cup Y_1, \dots, X_k \cup Y_k). \quad (\text{SDS})$$

A weighted relation is α -SDS if it corresponds under normalisation to a k -set function that is α -SDS, and a language Γ is SDS if there is some $\alpha \geq 1$ such that every weighted relation $\gamma \in \Gamma$ is α -SDS.

Note that SDS is a subclass of SEDS. To see this, consider any α -SDS k -set function f on V . Then it holds for all disjoint $X_1, \dots, X_k \subseteq V$ and all disjoint $Y_1, \dots, Y_k \subseteq V$ that

$$f(X_1 \setminus Y_1, \dots, X_k \setminus Y_k) \leq \alpha \cdot f(X_1, \dots, X_k) \leq \alpha \cdot (f(X_1, \dots, X_k) + f(Y_1, \dots, Y_k)),$$

proving that f is α -SEDS.

4 Classifying SEDS and SDS Languages

In this section, we first show that a SEDS language Γ is globally ℓ -tractable if it is SDS or if the reduced language $\text{fix}(\Gamma)$ is globally ℓ -tractable. Afterwards, we prove global s -intractability of the remaining SEDS languages conditioned on global s -intractability of $\text{fix}(\Gamma)$.

We begin by restating [12, Theorem 5.17] concerning EDS languages and then devise similar approximations for SEDS and SDS languages.

Theorem 20 ([12]). *For any α -EDS set function f on V , there exists a GMC instance h that $\alpha^{n+2} (n^3 + 2n)$ -approximates f , where $n = |V|$.*

Lemma 21. *For any α -SEDS k -set function f on V , there exists an α -EDS set function g that α^2 -approximates f .*

Proof. We define the set function g on V by $g(X) = \frac{1}{\alpha} f(X, \emptyset, \dots, \emptyset)$. Observe that, since f is normalised, it holds $g(\emptyset) = f(\emptyset, \dots, \emptyset) = 0$ and $g(X) = \frac{1}{\alpha} f(X, \emptyset, \dots, \emptyset) \geq 0$ for every $X \subseteq V$. Thus, g is normalised as well. In addition, for all $X, Y \subseteq V$, it holds that

$$\alpha \cdot (g(X) + g(Y)) = f(X, \emptyset, \dots, \emptyset) + f(Y, \emptyset, \dots, \emptyset) \geq \frac{1}{\alpha} \cdot f(X \setminus Y, \emptyset, \dots, \emptyset) = g(X \setminus Y),$$

where the second step uses equation (SEDS). Hence, g is α -EDS.

It remains to show that g α^2 -approximates f . For this purpose, consider any disjoint $X_1, \dots, X_k \subseteq V$ and let $X = \bigcup_{i=1}^k X_i$ denote their union. Since f is α -SIM, it holds that

$$g(X) = \frac{1}{\alpha} f(X, \emptyset, \dots, \emptyset) \leq f(X_1, \dots, X_k) \leq \alpha \cdot f(X, \emptyset, \dots, \emptyset) = \alpha^2 \cdot g(X). \quad \square$$

By combining Lemma 21 and Theorem 20, we can deduce the following result.

Theorem 22. *For any α -SEDS k -set function f on V , there exists a GMC instance h that $\alpha^{n+4} (n^3 + 2n)$ -approximates f , where $n = |V|$.*

Proof. Let f be an α -SEDS k -set function defined on V . According to Lemma 21, there exists an α -EDS set function g that α^2 -approximates f , meaning that, for all disjoint $X_1, \dots, X_k \subseteq V$, it holds

$$g\left(\bigcup_{i=1}^k X_i\right) \leq f(X_1, \dots, X_k) \leq \alpha^2 \cdot g\left(\bigcup_{i=1}^k X_i\right). \quad (1)$$

According to Theorem 20, as an α -EDS set function, g is $\alpha^{n+2} (n^3 + 2n)$ -approximable by some GMC instance h , meaning that, for every $X \subseteq V$, it holds

$$h(X) \leq g(X) \leq \alpha^{n+2} (n^3 + 2n) \cdot h(X). \quad (2)$$

By combining (1) and (2), it follows that, for all disjoint $X_1, \dots, X_k \subseteq V$, we have

$$h\left(\bigcup_{i=1}^k X_i\right) \leq f(X_1, \dots, X_k) \leq \alpha^{n+4} (n^3 + 2n) \cdot h\left(\bigcup_{i=1}^k X_i\right),$$

proving that h $\alpha^{n+4} (n^3 + 2n)$ -approximates f . \square

There is a more restrictive approximation of SDS languages through superadditive set functions, which can be thought of as GMC instances without edges.

Theorem 23. *For any α -SDS k -set function f on V , there exists a superadditive set function g that $n\alpha^{n+1}$ -approximates f , where $n = |V|$.*

Proof. Let the set function g on V be given by

$$g(X) = \frac{\alpha^{|X|-n-1} |X|}{n} \cdot f(X, \emptyset, \dots, \emptyset). \quad (3)$$

Observe that since f is normalised, it holds $g(\emptyset) = f(\emptyset, \dots, \emptyset) = 0$ and $g(X) \geq \frac{1}{n\alpha^{n+1}} \cdot f(X, \emptyset, \dots, \emptyset) \geq 0$ for every $X \subseteq V$. Thus, g is normalised as well. Moreover, g is superadditive, because for all disjoint $\emptyset \neq X, Y \subseteq V$, it holds that

$$\begin{aligned} g(X) + g(Y) &= \frac{\alpha^{|X|-n-1} |X|}{n} \cdot f(X, \emptyset, \dots, \emptyset) + \frac{\alpha^{|Y|-n-1} |Y|}{n} \cdot f(Y, \emptyset, \dots, \emptyset) \\ &\stackrel{\text{(SDS)}}{\leq} \frac{\alpha^{|X|-n-1} |X|}{n} \cdot \alpha \cdot f(X \cup Y, \emptyset, \dots, \emptyset) + \frac{\alpha^{|Y|-n-1} |Y|}{n} \cdot \alpha \cdot f(X \cup Y, \emptyset, \dots, \emptyset) \\ &\leq \frac{\alpha^{|X|+|Y|-n-1} |X|}{n} \cdot f(X \cup Y, \emptyset, \dots, \emptyset) + \frac{\alpha^{|X|+|Y|-n-1} |Y|}{n} \cdot f(X \cup Y, \emptyset, \dots, \emptyset) \\ &\stackrel{X \cap Y = \emptyset}{=} \frac{\alpha^{|X \cup Y|-n-1} |X \cup Y|}{n} \cdot f(X \cup Y, \emptyset, \dots, \emptyset) \quad (\text{since } X \cap Y = \emptyset) \\ &= g(X \cup Y). \end{aligned}$$

It remains to show that g $n\alpha^{n+1}$ -approximates f . Consider any disjoint $X_1, \dots, X_k \subseteq V$ and let $X = \bigcup_{i=1}^k X_i$. If $X = \emptyset$, it holds $g(X) = f(X_1, \dots, X_k) = 0$. Otherwise, it holds on the one hand that

$$g(X) = \frac{\alpha^{|X|-n-1} |X|}{n} \cdot f(X, \emptyset, \dots, \emptyset) \leq \frac{1}{\alpha} f(X, \emptyset, \dots, \emptyset) \stackrel{\text{(SIM)}}{\leq} f(X_1, \dots, X_k)$$

and on the other hand that

$$n\alpha^{n+1} \cdot g(X) = \alpha^{|X|} \cdot |X| \cdot f(X, \emptyset, \dots, \emptyset) \geq \alpha \cdot f(X, \emptyset, \dots, \emptyset) \stackrel{\text{(SIM)}}{\geq} f(X_1, \dots, X_k). \quad \square$$

Based on these approximations, we now show our main tractability theorem, which in places closely follows the proof of [12, Theorem 5.18].

Theorem 24. *Let Γ be a SEDS language. Then Γ is globally ℓ -tractable if it is SDS or if the reduced language $\text{fix}(\Gamma)$ is globally ℓ -tractable.*

Proof. Let Γ be an SEDS language on domain D . Then every weighted relation $\gamma \in \Gamma$ corresponds under normalisation to a k -set function f_γ . Furthermore, weighted relations in Γ are of bounded arity. If Γ is SDS, Theorem 23 implies that for some $\alpha \in \mathbb{N}$, every such k -set function f_γ can be α -approximated by a superadditive set function h_γ . In the following, we treat h_γ as a GMC instance without any edge weights. If Γ is not SDS, we can still α -approximate every k -set function f_γ by a GMC instance h_γ according to Theorem 22, but there is no restriction on the edge weights.

Let $l : D \rightarrow \mathbb{N}_0$ be a fixed lower bound and consider any $\text{VCSP}_l(\Gamma)$ instance I with objective function

$$\phi_I(x_1, \dots, x_n) = \sum_{i=1}^t w_i \cdot \gamma_i(\mathbf{x}^i).$$

Let f_I be the k -set function corresponding under normalisation to the objective function ϕ_I . We construct a GMC instance h that α -approximates f_I .

For $i \in [t]$, we relabel the vertices of h_{γ_i} to match the variables in the scope \mathbf{x}^i of the i -th constraint (i.e., vertex j is relabelled to x_j^i) and identify vertices in case of repeated variables. As the constraint is weighted by a non-negative factor w_i , we also scale the weights of the edges of h_{γ_i} and the superadditive set function by w_i (note that non-negative scaling preserves superadditivity). Instance h is then obtained by adding up GMC instances h_{γ_i} for all $i \in [t]$. In the following, we treat h as a $\text{GMC}_{l^*}^{l(0)}$ instance, where $l^* = \sum_{i=1}^k l(i)$. Note that if Γ is SDS, h has zero edge weights.

Let X_0, \dots, X_k be a partition of $[n]$ such that $f_I(X_1, \dots, X_k)$ is minimal among all partitions satisfying $|X_d| \geq l(d)$ for all $d \in D$. In other words, X_0, \dots, X_k corresponds to an optimal assignment for instance I . Let $X = \bigcup_{d=1}^k X_d$ denote all indices with non-zero labels. In addition, let $Y \subseteq [n]$ denote an optimal solution of the $\text{GMC}_{l^*}^{l(0)}$ instance h and let $\lambda = h(Y)$.

Since $|Y| \geq l^*$, there must exist some partition Y_1, \dots, Y_k of Y such that $|Y_d| \geq l(d)$ for all $1 \leq d \leq k$. Because h α -approximates f_I , it holds that

$$\lambda \leq h(X) \leq f_I(X_1, \dots, X_k) \leq f_I(Y_1, \dots, Y_k) \leq \alpha \cdot h(Y) = \alpha \cdot \lambda.$$

Hence, X is an α -optimal solution of h .

According to Lemma 6, it can be determined in polynomial time whether $\lambda = 0$, $\lambda = \infty$ or $0 < \lambda < \infty$. Furthermore, in case $\lambda = 0$, a solution Z such that $h(Z) = 0$ can be found. We explore this case first. Because Z must have size $|Z| \geq l^*$ as a solution of $\text{GMC}_{l^*}^{l(0)}$ instance h , we can select some partition Z_1, \dots, Z_k of Z such that $|Z_d| \geq l(d)$ for all $1 \leq d \leq k$. Since h α -approximates f_I , it must hold $f_I(Z_1, \dots, Z_k) \leq \alpha \cdot h(Z) = 0$, meaning that Z_1, \dots, Z_k represents an optimal assignment for instance I .

If $\lambda = \infty$, then obviously there are no feasible assignments.

Otherwise, it holds $0 < \lambda < \infty$. In this case, we distinguish whether Γ is SDS or $\text{fix}(\Gamma)$ is globally ℓ -tractable.

First, we assume that Γ is SDS and hence, that h has zero edge weights. We claim that the size of X is bounded by a constant. For the sake of contradiction, assume that $|X| \geq (\alpha + 1)l^*$. Then there are disjoint subsets $Z_1, Z_2, \dots, Z_{\alpha+1} \subseteq X$ such that $|Z_i| \geq l^*$ for all $1 \leq i \leq \alpha + 1$. Being a solution of h , every Z_i must have value at least $h(Z_i) \geq \lambda$. Based on the superadditivity of h , we arrive at the contradiction

$$(\alpha + 1) \cdot \lambda \leq h(Z_1) + \dots + h(Z_{\alpha+1}) \leq h(X) \leq \alpha \cdot \lambda.$$

Thus, it must hold $|X| < (\alpha + 1)l^*$. This leaves less than $O(n^{(\alpha+1)l^*})$ possible choices for X , each of which admits at most $O(k^{(\alpha+1)l^*})$ partitions of the form $X_1 \cup \dots \cup X_k = X$. By checking all of these, we can retrieve the sets X_1, \dots, X_k in polynomial time.

Now, we assume that $\text{fix}(\Gamma)$ is globally ℓ -tractable. According to Corollary 10, there are only polynomially many α -optimal solutions of h , all of which can be computed in polynomial time. X must be among those solutions. By repeating the following process for all of them, we can assume that X is known, and so is $X_0 = [n] \setminus X$.

Let $D^* = D \setminus \{0\}$ and let $l|_{D^*} : D^* \rightarrow \mathbb{N}$ denote the restriction of l to D^* . We consider the $\text{VCSP}_{l|_{D^*}}(\text{fix}(\Gamma))$ instance $I_X = (X, D^*, \phi_X)$, where objective function ϕ_X is constructed from ϕ_I by fixing label 0 to the variables in X_0 . This construction can be realised by replacing every weighted relation γ_i from ϕ_I with $\text{fix}_{U_i}[\gamma_i]$ instead, where U_i are the indices of the variables from X_0 in the scope of γ_i , and the remaining variables in the scope of γ_i form the new scope for $\text{fix}_{U_i}[\gamma_i]$. According to this construction, by assigning label 0 to the variables in X_0 , every assignment for I_X induces an assignment for I with the same objective value. This includes the assignment for I_X represented by the sets X_1, \dots, X_k . Thus, an optimal assignment for I_X , which can be obtained efficiently when $\text{fix}(\Gamma)$ is globally ℓ -tractable, induces an optimal assignment for I . \square

Remark 25. In fact, the algorithm presented in Theorem 24 can, for every fixed lower bound $l : D \rightarrow \mathbb{N}_0$ and every $\text{VCSP}_l(\Gamma)$ instance I with optimal value $0 < \lambda < \infty$, enumerate all optimal assignments for I in polynomial time if either

- (i) Γ is SDS, or
- (ii) Γ is SEDS, $\text{fix}(\Gamma)$ is globally ℓ -tractable and for every $\text{VCSP}_l(\text{fix}(\Gamma))$ instance with optimal value $0 < \lambda < \infty$, all optimal assignments can be enumerated in polynomial time.

To complete our analysis of SEDS languages, we will now focus on the case that a language is not SDS and that $\text{fix}(\Gamma)$ is globally s -intractable. Going even beyond SEDS, our main hardness result is that SIM languages are globally s -intractable under those circumstances.

Theorem 26. *Let Γ be a valued constraint language over domain D that is SIM, but not SDS, and let $\text{fix}(\Gamma)$ be globally s -intractable. Then Γ is globally s -intractable.*

Proof. Since $\text{fix}(\Gamma)$ is globally s -intractable, the domain D must have at least three elements. Let $\alpha \geq 1$ be such that Γ is α -SIM. We show that $\text{VCSP}_s(\text{fix}(\Gamma))$ is reducible to $\text{VCSP}_s(\Gamma)$.

For this purpose, let $I = (V, D^*, \phi_I)$ be any $\text{VCSP}_s(\text{fix}(\Gamma))$ instance on domain $D^* = D \setminus \{0\}$ with objective function $\phi_I(\mathbf{x}) = \sum_{i=1}^t w_i \gamma_i(\mathbf{x}_i)$.

Every constraint γ_i must be of the form $\gamma_i = \text{fix}_{U_i}[\sigma_i]$ for some weighted relation $\sigma_i \in \Gamma$ and some set $U_i \subseteq [\text{ar}(\sigma_i)]$. Let σ'_i denote the identification of the weighted relation σ_i at the coordinates in U_i , i.e. such that $\sigma'_i(\mathbf{x}_i, 0) = \gamma_i(\mathbf{x}_i)$ for every $\mathbf{x}_i \in (D^*)^{\text{ar}(\gamma_i)}$. Here and later on in the proof, the notation $\sigma'_i(\mathbf{x}_i, 0)$ is shorthand for $\sigma'_i(x_{i,1}, \dots, x_{i,\text{ar}(\gamma_i)}, 0)$. Note that σ'_i is expressible over Γ . We will utilise these relations later in the proof in order to express the objective function ϕ_I over Γ .

Let $\varepsilon > 0$ be a lower bound for the smallest positive difference between the values of any two assignments for instance I . In other words, we select ε sufficiently small so that if the objective value of some assignment is κ , then there is no other assignment with objective value in $(\kappa - \varepsilon, \kappa)$ or $(\kappa, \kappa + \varepsilon)$. Note that ε can be calculated efficiently by multiplying the denominators of all values that the constraints can obtain and of all weights that occur in ϕ_I .

Similarly, let ω denote an upper bound for the largest finite value that any assignment for instance I can obtain.

If Γ is not SDS, in particular not $\left(\frac{2|V|^2 \cdot \omega}{\varepsilon} \cdot \alpha^4\right)$ -SDS, then there must exist a weighted relation $\gamma \in \Gamma$ of some arity r and disjoint $X_1, \dots, X_k, Y_1, \dots, Y_k \subseteq [r]$ such that, in violation of equation (SDS), the k -set function f corresponding under normalisation to γ satisfies

$$f(X_1, \dots, X_k) > \frac{2|V|^2 \cdot \omega}{\varepsilon} \cdot \alpha^4 \cdot f(X_1 \cup Y_1, \dots, X_k \cup Y_k). \quad (4)$$

Let $X = \bigcup_{d=1}^k X_d$ and $Y = \bigcup_{d=1}^k Y_d$. Since f is α -SIM, we can transform (4) to

$$f(X, \emptyset, \dots, \emptyset) > \frac{2|V|^2 \cdot \omega}{\varepsilon} \cdot \alpha^2 \cdot f(X \cup Y, \emptyset, \dots, \emptyset). \quad (5)$$

Without loss of generality, we can assume that $\gamma(\mathbf{0}^r) = 0$ so that γ and f are interchangeable. In order to simplify notation, we first define the 3-ary weighted relation γ^* by $\gamma^*(x, y, z) = \gamma(\mathbf{s}_{x,y,z})$ for $x, y, z \in D$, where the i -th coordinate s_i of $\mathbf{s}_{x,y,z}$ is $s_i = x$ if $i \in X$, $s_i = y$ if $i \in Y$ and $s_i = z$ otherwise.

According to the (5), it holds that

$$\gamma^*(1, 0, 0) > \frac{2|V|^2 \cdot \omega}{\varepsilon} \cdot \alpha^2 \cdot \gamma(1, 1, 0).$$

Since γ is α -SIM, this implies for all $x, y, z \in D^*$ that

$$\gamma^*(x, 0, 0) > \frac{2|V|^2 \cdot \omega}{\varepsilon} \cdot \gamma(y, z, 0). \quad (6)$$

Finally, let $\nu > 0$ be a sufficiently large value so that, for all $x, y, z \in D$ such that $\gamma^*(x, y, z) > 0$, it holds that

$$\nu \cdot \gamma^*(x, y, z) > \omega. \quad (7)$$

Based on these definitions, we can now complete the proof. We distinguish two cases.

Case 1: First, assume that $\gamma^*(1, 1, 1) = 0$.

We construct the $\text{VCSP}_s(\Gamma)$ instance $I' = (V \cup \{z\}, D, \phi_{I'})$ with objective function

$$\phi_{I'}(\mathbf{x}, z) = \sum_{x, y \in V} \nu \cdot \gamma^*(x, y, y) + \sum_{i=1}^t w_i \sigma'_i(\mathbf{x}_i, z).$$

From $\gamma^*(1, 1, 1) = 0$ and the fact that Γ is α -SIM, it follows that $\gamma^*(x, y, y) = 0$ for all $x, y \in D^*$. We focus on assignments for I' of the form $\mathbf{x} \in (D^*)^{|V|}$ and $z = 0$. For every such assignment, it must hold

$$\phi_{I'}(\mathbf{x}, z) = 0 + \sum_{i=1}^t w_i \sigma'_i(\mathbf{x}_i, z) = \sum_{i=1}^t w_i \gamma_i(\mathbf{x}_i) = \phi_I(\mathbf{x}).$$

Hence, every assignment for I' of the form $\mathbf{x} \in (D^*)^{|V|}$ and $z = 0$ induces an assignment $\mathbf{x} \in (D^*)^{|V|}$ for I with the same objective value, and vice versa. In particular, if I is feasible, then there is an assignment for I' of value at most ω . To show that an optimal assignment for I can be derived from an optimal assignment for I' , it remains to show that every minimal assignment for I' must be of the described form, which we do by showing that every assignment violating this form must have value greater than ω .

Consider any surjective assignment for \mathbf{x} and z . Since $|D| \geq 3$, there must be some variable $x \in V$ such that $x \neq 0$. If there was any $y \in V$ with assigned label $y = 0$, then it would hold $\gamma^*(x, y, y) > 0$ according to (6) and therefore

$$\phi_{I'}(\mathbf{x}, z) \geq \nu \cdot \gamma^*(x, y, y) \stackrel{(7)}{>} \omega.$$

Thus, we can assume $\mathbf{x} \in (D^*)^{|V|}$ in every minimal assignment. By the surjectivity of the assignment, that implies $z = 0$ and completes the reduction proof in this case.

Case 2: Now, assume that $\gamma^*(1, 1, 1) > 0$. In this case, we construct the $\text{VCSP}_s(\Gamma)$ instance $I^* = (V \cup \{z\}, D, \phi_{I^*})$ with objective function

$$\phi_{I^*}(\mathbf{x}, z) = \nu \cdot \gamma^*(z, z, z) + \sum_{x, y \in V} \frac{\varepsilon \cdot \gamma^*(x, y, z)}{2|V|^2 \max_{a, b \in D^*} \gamma^*(a, b, 0)} + \sum_{i=1}^t w_i \sigma'_i(\mathbf{x}_i, z).$$

An assignment of the form $\mathbf{x} \in (D^*)^{|V|}$ and $z = 0$ satisfies on the one hand that

$$\phi_{I^*}(\mathbf{x}, z) \leq 0 + \frac{\varepsilon}{2} + \sum_{i=1}^t w_i \gamma_i(\mathbf{x}_i) = \frac{\varepsilon}{2} + \phi_I(\mathbf{x}),$$

and on the other hand that

$$\phi_{I^*}(\mathbf{x}, z) \geq \sum_{i=1}^t w_i \gamma_i(\mathbf{x}_i) = \phi_I(\mathbf{x}).$$

Hence, an assignment for I^* of the form $\mathbf{x} \in (D^*)^{|V|}$ and $z = 0$ induces an assignment $\mathbf{x} \in (D^*)^{|V|}$ for I of similar value, and vice versa. It remains to show that every minimal

assignment for I^* must be of this form. This completes the reduction proof, because by our choice of ε , a minimal assignment for I^* of this form must then induce a minimal assignment for I .

By the assumption $\gamma^*(1, 1, 1) > 0$ and since Γ is SIM, we have $\gamma^*(z, z, z) > 0$ for every $z \in D^*$. Thus, every assignment of the form $\mathbf{x} \in D^{|V|}$ and $z \in D^*$ must have objective value

$$\phi_{I^*}(\mathbf{x}, z) \geq \nu \cdot \gamma^*(z, z, z) \stackrel{(7)}{>} \omega$$

and thereby cannot be optimal.

Otherwise, when $z = 0$, there must be some $x \in V$ in every surjective assignment such that $x = 1$. If there was any variable $y \in V$ such that $y = 0$, then, for the summand $\gamma^*(x, y, z)$ in the second part of ϕ_{I^*} , it would hold that

$$\gamma^*(x, y, z) = \gamma^*(1, 0, 0) > \frac{2|V|^2 \cdot \omega}{\varepsilon} \cdot \max_{0 \neq a, b} \gamma^*(a, b, 0),$$

and hence,

$$\phi_{I^*}(\mathbf{x}, z) \geq \sum_{x, y \in V} \frac{\varepsilon \cdot \gamma^*(x, y, z)}{2|V|^2 \max_{a, b \in D^*} \gamma^*(a, b, 0)} > \omega.$$

Thus, in addition to $z = 0$, it must also hold $\mathbf{x} \in (D^*)^{|V|}$ in every minimal assignment. This reduces $\text{VCSP}_s(\text{fix}(\Gamma))$ to $\text{VCSP}_s(\Gamma)$ in this case as well and thereby completes our proof. \square

5 Lower-Bounded VCSPs on the Boolean Domain

In this final section, we prove our dichotomy theorem for lower-bounded VCSPs on the Boolean domain and, in the end, extend this result to SEDS languages on three-element domains. A classification of Boolean surjective VCSPs has been given by [12] based on polymorphisms and multimorphisms [17, 7], which we introduce in the following.

Definition 27. Let r and s be positive integers and let γ be a r -ary weighted relation on domain D . An operation $o : D^s \rightarrow D$ is a *polymorphism* of γ (and γ *admits* polymorphism o) if, for all $\mathbf{x}_1, \dots, \mathbf{x}_s \in D^r$ such that $\gamma(\mathbf{x}_1), \dots, \gamma(\mathbf{x}_s) < \infty$, it holds $\gamma(o(\mathbf{x}_1, \dots, \mathbf{x}_s)) < \infty$, where o is applied componentwise as

$$o(\mathbf{x}_1, \dots, \mathbf{x}_s) = (o(x_{1,1}, \dots, x_{s,1}), \dots, o(x_{1,r}, \dots, x_{s,r})).$$

A language Γ admits polymorphism o if o is a polymorphism of every $\gamma \in \Gamma$.

Definition 28. Let r and s be positive integers and let γ be a r -ary weighted relation on domain D . A list $\langle o_1, \dots, o_s \rangle$ of s -ary polymorphisms of γ is a *multimorphism* of γ (and γ *admits* multimorphism $\langle o_1, \dots, o_s \rangle$) if, for all $\mathbf{x}_1, \dots, \mathbf{x}_s \in D^r$, it holds that

$$\sum_{i=1}^s \gamma(o_i(\mathbf{x}_1, \dots, \mathbf{x}_s)) \leq \sum_{i=1}^s \gamma(\mathbf{x}_i).$$

A language Γ admits multimorphism $\langle o_1, \dots, o_s \rangle$ if every $\gamma \in \Gamma$ admits $\langle o_1, \dots, o_s \rangle$.

For $d \in D$, the constant operation $c_d : D \rightarrow D$ is defined by $c_d(x) = d$ for every $x \in D$. According to this definition, a language Γ admits multimorphism $\langle c_d \rangle$ for some $d \in D$ if every weighted relation $\gamma \in \Gamma$ satisfies $\gamma(d, d, \dots, d) \leq \gamma(\mathbf{x})$ for all $\mathbf{x} \in D^{\text{ar}(\gamma)}$. Such a language is always tractable, but it may not be s -tractable or ℓ -tractable. Note that the class **SIM** and all subclasses only contain languages that admit multimorphism $\langle c_0 \rangle$, because this is a requirement for corresponding under normalisation to a k -set function.

In addition, the following operations for the Boolean domain $D = \{0, 1\}$, which were initially given by [7], will be relevant for us.

- The binary operation \min (\max) returns the smaller (larger) of its two arguments with respect to the order $0 < 1$.
- The minority operation $\text{Mn} : D^3 \rightarrow D$ is defined for $x, y \in D$ by $\text{Mn}(x, x, y) = \text{Mn}(x, y, x) = \text{Mn}(y, x, x) = y$.
- Similarly, the majority operation $\text{Mj} : D^3 \rightarrow D$ is given for $x, y \in D$ by $\text{Mj}(x, x, y) = \text{Mj}(x, y, x) = \text{Mj}(y, x, x) = x$.

Furthermore, given a Boolean language Γ , let $\neg(\Gamma)$ denote the language where labels 0 and 1 are flipped. This can be seen as relabelling the domain so that VCSPs over Γ and over $\neg(\Gamma)$ have the same complexity.

Based on these operations, [12, Theorem 3.2] gives a classification of Boolean $\overline{\mathbb{Q}}$ -valued languages with respect to global s -tractability, which we restate here.

Theorem 29 ([12]). *Let Γ be a Boolean language. Then Γ is globally s -tractable if Γ is EDS, if $\neg(\Gamma)$ is EDS or if Γ admits any of the following multimorphisms: $\langle \min, \min \rangle$, $\langle \max, \max \rangle$, $\langle \min, \max \rangle$, $\langle \text{Mn}, \text{Mn}, \text{Mn} \rangle$, $\langle \text{Mj}, \text{Mj}, \text{Mj} \rangle$, $\langle \text{Mj}, \text{Mj}, \text{Mn} \rangle$. Otherwise, Γ is globally s -intractable.*

Note that if $\text{P} \neq \text{NP}$, global s -tractability and global s -intractability are mutually exclusive. In order to extend Theorem 29 to lower-bounded VCSPs, we rely on the results from Section 4 as well as the following two auxiliary lemmas.

Lemma 30. *Let Γ be a Boolean language and let $\alpha \geq 1$. Then Γ is α -SEDS if and only if it is α -EDS.*

Proof. As a Boolean language, Γ is α -SIM if every weighted relation $\gamma \in \Gamma$ corresponds under normalisation to a set function. This is the case if Γ is α -EDS.

The remainder of the definitions of EDS and SEDS from pages 13 and 14 are equivalent for the Boolean domain, showing the statement. \square

Recall that for a label $d \in D$, the constant relation ρ_d is defined by $\rho_d(d) = 0$ and $\rho_d(x) = \infty$ for $d \neq x \in D$. Let $\mathcal{C}_D = \{\rho_d \mid d \in D\}$ denote the set of constant unary relations.

Lemma 31. *Let Γ be a language on domain D such that $\Gamma \cup \mathcal{C}_D$ is globally tractable. Then Γ is globally ℓ -tractable.*

Proof. Let $l : D \rightarrow \mathbb{N}_0$ be a fixed lower bound, let $l^* = \sum_{d \in D} l(d)$ and consider any VCSP $_l(\Gamma)$ instance $I = (V, D, \phi_I)$. There are only $O(|V|^{l^*})$, i.e. polynomially many, ways to select disjoint sets $V_d \subseteq V$ of size $|V_d| = l(d)$ for all $d \in D$. For each such choice, we construct a

VCSP($\Gamma \cup \mathcal{C}_D$) instance $I' = (V, D, \phi'_I)$, where ϕ'_I is constructed from ϕ_I by adding a constraint $\rho_d(x)$ for every $d \in D$ and every $x \in V_d$. These additional constraints guarantee that only those assignments for I' are feasible that respect lower bound l .

Conversely, every assignment for I that respects lower bound l is an assignment for some instance I' constructed from some disjoint sets $V_d \subseteq V$ of the described form. Therefore, an assignment that is minimal among all optimal assignments for instances I' must be an optimal assignment for I . \square

Theorem 32. *Let Γ be a Boolean language. Then Γ is globally ℓ -tractable if and only if it is globally s -tractable. Otherwise, Γ is globally ℓ -intractable.*

Proof. We assume that $P \neq NP$, because otherwise every language is globally ℓ -tractable and the statement trivially holds true. If Γ is globally s -tractable, it must satisfy at least one of the properties listed in Theorem 29.

First, we assume that Γ admits any of the multimorphisms $\langle \min, \min \rangle$, $\langle \max, \max \rangle$, $\langle \min, \max \rangle$, $\langle \text{Mn}, \text{Mn}, \text{Mn} \rangle$, $\langle \text{Mj}, \text{Mj}, \text{Mj} \rangle$ or $\langle \text{Mj}, \text{Mj}, \text{Mn} \rangle$. Then $\Gamma \cup \{\rho_0, \rho_1\}$ must be tractable as well, because the constant relations ρ_0 and ρ_1 both admit all of these multimorphisms. This implies global ℓ -tractability of Γ according to Lemma 31.

If Γ is EDS, it must also be SEDS according to Lemma 30. Furthermore, the reduced language $\text{fix}(\Gamma)$ is trivial in this case and, in particular, globally ℓ -tractable. Hence, Γ must be globally ℓ -tractable by Theorem 24. The same applies if $\neg(\Gamma)$ is EDS.

Otherwise, Γ must be globally s -intractable according to Theorem 29. That immediately implies global ℓ -intractability. \square

Hence, the classification from Theorem 29 is also valid for lower-bounded VCSPs. For \mathbb{Q} -valued and $\{0, \infty\}$ -valued languages, a tighter classification of Boolean surjective VCSPs is provided in [12], which can in the same way be lifted to lower-bounded VCSPs by Theorem 32. In particular, a Boolean \mathbb{Q} -valued language Γ is globally ℓ -tractable precisely if it is EDS, if $\neg(\Gamma)$ is EDS or if Γ is submodular.

While our focus so far has been on global s -tractability and global ℓ -tractability, there is a further distinction for infinite languages. A language Γ is *tractable* if every finite subset $\Gamma' \subseteq \Gamma$ is globally tractable, and *intractable* if some finite subset is globally intractable. The terms *s-tractability* and *ℓ -tractability* are defined analogously for surjective and lower-bounded VCSPs. [12, Remark 2] outlines a dichotomy theorem for Boolean languages with respect to s -tractability. We lift this result to lower-bounded VCSPs.

Corollary 33. *Let Γ be a Boolean language. Then Γ is ℓ -tractable if and only if it is s -tractable. Otherwise, Γ is ℓ -intractable.*

Proof. If Γ is s -tractable, every finite subset $\Gamma' \subseteq \Gamma$ is s -tractable. Since s -tractability and global s -tractability coincide for finite languages, every finite $\Gamma' \subseteq \Gamma$ must be globally s -tractable. By Theorem 32, every finite $\Gamma' \subseteq \Gamma$ is then globally ℓ -tractable and therefore, Γ is ℓ -tractable.

Otherwise, if Γ is not s -tractable, there must be some finite subset $\Gamma' \subseteq \Gamma$ that is not s -tractable. In this case, Γ' cannot be globally s -tractable and must instead be globally ℓ -intractable by Theorem 32. Hence, Γ is ℓ -intractable. \square

Moreover, we can now classify lower-bounded VCSPs over SEDS languages on three-element domains.

Theorem 34. *Let Γ be a SEDS language on domain $D = \{0, 1, 2\}$. Then Γ is globally ℓ -tractable if it is SDS or if $\text{fix}(\Gamma)$ is globally ℓ -tractable, and globally ℓ -intractable otherwise.*

Proof. If Γ is SDS or $\text{fix}(\Gamma)$ globally ℓ -tractable, the statement follows from Theorem 24. Otherwise, $\text{fix}(\Gamma)$ must be globally s -intractable by Theorem 32 and the dichotomy from [12, Theorem 3.2]. Hence, Γ is globally s -intractable by Theorem 26, which gives the result. \square

6 Conclusions

Based on the newly introduced Bounded Generalised Min-Cut problem and its tractability, which might be of independent interest, we have provided a conditional complexity classification of surjective and lower-bounded SEDS VCSPs on non-Boolean domains. Consequently, we obtained a dichotomy theorem with respect to ℓ -tractability for Boolean domains as well as, more interestingly, for SEDS languages on three-element domains.

While our results only pertain to languages that admit multimorphism $\langle c_d \rangle$ for some label d we expect our results and techniques to be useful in identifying new s -tractable and ℓ -tractable languages going beyond those admitting $\langle c_d \rangle$.

As mentioned in Section 1, globally tractable languages that include constant relations are also s -tractable. It is easy to show the same for global ℓ -tractability. For example, this shows that well-studied sources of tractability such as submodularity [26] and its generalisation k -submodularity [16], which are known to be globally tractable [20], are also globally ℓ -tractable.

What other non-Boolean languages are s -tractable and ℓ -tractable? Our results are a first step in this direction. In all cases we encountered global s -(in)tractability coincides with global ℓ -(in)tractability. We do not know whether this is true in general.

The natural next step is to consider languages on three-element domains. As is often the case in the (V)CSP literature, languages on three-element domains are significantly more complex than Boolean languages; for instance, compare two-element CSPs [24] and three-element CSPs [3]. There is an interesting surjective CSP on a three-element domain, known as the *3-No-Rainbow-Colouring* problem [2]. The task is to colour the vertices of a three-regular hypergraph such that every colour is used at least once, while no hyperedge attains all three colours. It has recently been shown that the 3-No-Rainbow-Colouring is NP-hard [30]. Consequently, we expect that it should be possible to classify all three-element surjective CSPs and perhaps even all three-element surjective VCSPs.

Acknowledgements

We would like to thank the anonymous referees of both the conference [22] and this full version of the paper. We also thank Costin-Andrei Onicescu for detailed feedback on a previous version of this paper.

References

- [1] Libor Barto and Marcin Kozik. Constraint Satisfaction Problems Solvable by Local Consistency Methods. *Journal of the ACM*, 61(1), 2014. Article No. 3.

- [2] Manuel Bodirsky, Jan Kára, and Barnaby Martin. The complexity of surjective homomorphism problems – a survey. *Discrete Applied Mathematics*, 160(12):1680–1690, 2012.
- [3] Andrei Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM*, 53(1):66–120, 2006.
- [4] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the Complexity of Constraints using Finite Algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.
- [5] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 319–330, 2017.
- [6] Andrei A. Bulatov and Dániel Marx. The complexity of global cardinality constraints. *Logical Methods in Computer Science*, Volume 6, Issue 4, 2010.
- [7] David A. Cohen, Martin C. Cooper, Peter G. Jeavons, and Andrei A. Krokhin. The Complexity of Soft Constraint Satisfaction. *Artificial Intelligence*, 170(11):983–1016, 2006.
- [8] Nadia Creignou and Jean-Jacques Hébrard. On generating all solutions of generalized satisfiability problems. *Informatique Théorique et Applications*, 31:499–511, 11 1997.
- [9] Elias Dahlhaus, David S. Johnson, Christos H. Papadimitriou, Paul D. Seymour, and Mihalis Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, 1994.
- [10] Víctor Dalmau and Justin Pearson. Set Functions and Width 1 Problems. In *Proceedings of the 5th International Conference on Constraint Programming (CP'99)*, volume 1713 of *Lecture Notes in Computer Science*, pages 159–173. Springer, 1999.
- [11] Tomás Feder and Moshe Y Vardi. The computational structure of monotone monadic snc and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.
- [12] Peter Fulla, Hannes Uppman, and Stanislav Živný. The complexity of Boolean surjective general-valued CSPs. *ACM Transactions on Computation Theory*, 11(1), 2018. Article No. 4.
- [13] Olivier Goldschmidt and Dorit S. Hochbaum. A polynomial algorithm for the k-cut problem for fixed k. *Mathematics of Operations Research*, 19(1):24–37, 1994.
- [14] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990.
- [15] Pavol Hell and Jaroslav Nešetřil. Colouring, constraint satisfaction, and complexity. *Computer Science Review*, 2(3):143 – 163, 2008.
- [16] Anna Huber and Vladimir Kolmogorov. Towards Minimizing k -Submodular Functions. In *Proceedings of the 2nd International Symposium on Combinatorial Optimization (ISCO'12)*, volume 7422 of *Lecture Notes in Computer Science*, pages 451–462. Springer, 2012.

- [17] Peter Jeavons, David Cohen, and Marc Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, July 1997.
- [18] David R. Karger. Global min-cuts in RNC, and other ramifications of a simple min-out algorithm. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'93)*, pages 21–30, 1993.
- [19] Vladimir Kolmogorov, Andrei Krokhin, and Michal Rolínek. The complexity of general-valued CSPs. *SIAM Journal on Computing*, 46(3):1087–1110, 2017.
- [20] Vladimir Kolmogorov, Johan Thapper, and Stanislav Živný. The power of linear programming for general-valued CSPs. *SIAM Journal on Computing*, 44(1):1–36, 2015.
- [21] Marcin Kozik and Joanna Ochremiak. Algebraic properties of valued constraint satisfaction problem. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15)*, volume 9134 of *Lecture Notes in Computer Science*, pages 846–858. Springer Berlin Heidelberg, 2015.
- [22] Gregor Matl and Stanislav Živný. Beyond Boolean surjective VCSPs. In *Proceedings of the 36th Annual Symposium on Theoretical Aspects of Computer Science (STACS'19)*, pages 48:1–48:15, 2019.
- [23] Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *The Handbook of Constraint Programming*. Elsevier, 2006.
- [24] Thomas J. Schaefer. The Complexity of Satisfiability Problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing (STOC'78)*, pages 216–226. ACM, 1978.
- [25] Thomas Schiex, H el ene Fargier, and G erard Verfaillie. Valued Constraint Satisfaction Problems: Hard and Easy Problems. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 631–637, 1995.
- [26] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.
- [27] Johan Thapper and Stanislav Živný. The power of Sherali-Adams relaxations for general-valued CSPs. *SIAM Journal on Computing*, 46(4):1241–1279, 2017.
- [28] Johan Thapper and Stanislav Živný. The complexity of finite-valued CSPs. *Journal of the ACM*, 63(4):37:1–37:33, September 2016.
- [29] D. Zhuk. A proof of CSP dichotomy conjecture. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*, pages 331–342, 2017.
- [30] Dmitriy Zhuk. No-Rainbow Problem is NP-Hard. Technical report, 2020.