A dichotomy theorem for nonuniform CSPs

ANDREI A. BULATOV

In a non-uniform Constraint Satisfaction problem CSP(Γ), where Γ is a set of relations on a finite set A, the goal is to find an assignment of values to variables subject to constraints imposed on specified sets of variables using the relations from Γ. The Dichotomy Conjecture for the non-uniform CSP states that for every constraint language Γ the problem CSP(Γ) is either solvable in polynomial time or is NP-complete. It was proposed by Feder and Vardi in their seminal 1993 paper. In this paper we confirm the Dichotomy Conjecture.

ACM Reference Format:

Andrei A. Bulatov . 2025. A dichotomy theorem for nonuniform CSPs. 1, 1 (December 2025), 43 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

In a Constraint Satisfaction Problem (CSP) the question is to decide whether or not it is possible to satisfy a given set of constraints. One of the standard ways to specify a constraint is to require that a combination of values of a certain set of variables belongs to a given relation. If the constraints allowed in a problem have to come from some set Γ of relations over a fixed finite set, such a restricted problem is referred to as a *nonuniform CSP* and denoted CSP(Γ). The set Γ is then called a *constraint language*. Nonuniform CSPs not only provide a powerful framework ubiquitous across a wide range of disciplines from theoretical computer science to computer vision, but also admit natural and elegant reformulations such as the homomorphism problem, and characterizations, in particular, as the class of problems equivalent to a logic class MMSNP [35]. Many different versions of the CSP have been studied across various fields. These include CSPs over infinite sets, counting CSPs (and related Holant problem and the problem of computing partition functions), several variants of optimization CSPs, valued CSPs, quantified CSPs, and numerous related problems. The reader is referred to the recent book [49] for a survey of the state-of-the art in some of these areas. In this paper we, however, focus on the decision nonuniform CSP and its complexity.

A systematic study of the complexity of nonuniform CSPs was started by Schaefer in 1978 [57] who showed that for every constraint language Γ over a 2-element set the problem CSP(Γ) is either solvable in polynomial time or is NP-complete. Schaefer also asked about the complexity of CSP(Γ) for languages over larger sets. The next step in the study of nonuniform CSPs was made in the seminal paper by Feder and Vardi [34, 35], who apart from considering numerous aspects of the problem, posed the *Dichotomy Conjecture* that states that for every finite constraint language Γ over a finite set the problem CSP(Γ) is either solvable in polynomial time or is NP-complete. This conjecture has become a focal point for research on the complexity of the CSP and most of the effort in this area revolves to some extent around the Dichotomy Conjecture.

The complexity of the CSP in general and the Dichotomy Conjecture in particular has been studied by several research communities using a variety of methods, each contributing an important aspect of the problem. The CSP has been an

Author's address: Andrei A. Bulatov.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Association for Computing Machinery.
Manuscript submitted to ACM

established area in artificial intelligence for decades, and apart from developing efficient general methods of solving CSPs researchers tried to identify tractable fragments of the problem [33]. A very important special case of the CSP, the (Di)Graph Homomorphism problem and the H-Coloring problem have been actively studied in the graph theory community, see, e.g. [39, 40] and subsequent works by Hell, Feder, Bang-Jensen, Rafiey and others. Homomorphism duality introduced in these works has been very useful in understanding the structure of constraint problems. The CSP plays a major role and has been successfully studied in database theory, logic and model theory [38, 46, 47], although the version of the problem mostly used there is not necessarily nonuniform. Logic games and strategies are now a standard tool in many of CSP algorithms. An interesting approach to the Dichotomy Conjecture through long codes was suggested by Kun and Szegedy [50]. Brown-Cohen and Raghavendra proposed to study the conjecture using techniques based on decay of correlations [11]. In this paper we use the algebraic structure of the CSP, which is briefly discussed next.

The most effective approach to the study of the CSP turned out to be the *algebraic approach* that associates every constraint language with its (universal) algebra of polymorphisms. This approach was first developed in a series of papers by Jeavons and coauthors [43–45] and then refined by Bulatov, Krokhin, Barto, Kozik, Maroti, Zhuk and others [5, 6, 8, 16, 27, 29, 53, 54, 61, 62]. While the complexity of $\mathrm{CSP}(\Gamma)$ has been already solved for some interesting classes of structures such as graphs [40], the algebraic approach allowed the researchers to confirm the Dichotomy Conjecture in a number of more general cases: for languages over a set of size up to 7 [12, 17, 52, 62], so called conservative languages [3, 14, 18, 19], and some classes of digraphs [7]. It also helped to design the main classes of CSP algorithms [6, 10, 21, 26, 42], and to refine the exact complexity of the CSP [1, 8, 32, 51].

In this paper we confirm the Dichotomy Conjecture for arbitrary languages over finite sets. More precisely, we prove the following

THEOREM 1.1. *For any finite constraint language* Γ *over a finite set the problem* $\mathrm{CSP}(\Gamma)$ *is either solvable in polynomial time or is NP-complete.*

The same result has been independently obtained by Zhuk [63, 64]. While Zhuk's proof also uses the algebraic approach, the specific algebraic tools he uses significantly differ from what is used in this paper.

The criterion for being solvable in polynomial time that we establish below matches the algebraic form of the Dichotomy Conjecture suggested in [27]. The hardness part of the conjecture has been known for a long time. Therefore the main achievement of this paper is a polynomial time algorithm for problems satisfying the tractability condition from [27].

Using algebraic language we can state the result in a stronger form. Let $\mathbb{A}$ be a finite idempotent algebra and let $\mathrm{CSP}(\mathbb{A})$ denote the union of problems $\mathrm{CSP}(\Gamma)$ such that every term operation of $\mathbb{A}$ is a polymorphism of Γ. Problem $\mathrm{CSP}(\mathbb{A})$ is no longer a nonuniform CSP, and Theorem 1.1 allows for problems $\mathrm{CSP}(\Gamma) \subseteq \mathrm{CSP}(\mathbb{A})$ to have different solution algorithms even when $\mathbb{A}$ meets the tractability condition. We show that the solution algorithm only depends on the algebra $\mathbb{A}$.

THEOREM 1.2. *For a finite idempotent algebra that satisfies the conditions of the Dichotomy Conjecture (see Theorem 2.4) there is a uniform solution algorithm for* $\mathrm{CSP}(\mathbb{A})$.

An interesting question arising from Theorems 1.1,1.2 is known as the *Meta-problem*: Given a constraint language or a finite algebra, decide whether or not it satisfies the conditions of Theorem 2.4. The answer to this question is not straightforward, for a thorough study of the Meta-problem see [31, 37].

We start by introducing the terminology and notation for CSPs that is used throughout the paper and recalling the basics of the algebraic approach. Then in Section 4 we introduce the key ingredients used in the algorithm: separation of congruences and centralizers. Then in Section 5 we apply these concepts to CSPs, first, to demonstrate how centralizers help to decompose an instance into smaller subinstances, and, second, to introduce a new kind of minimality condition for CSPs, *block minimality*. After that we state the main results used by the algorithm and describe the algorithm itself. The last part of the paper, Sections 6–9, is devoted to proving the technical results.

2 CSP, UNIVERSAL ALGEBRA AND THE DICHOTOMY CONJECTURE

For a detailed introduction to the CSP and the algebraic approach to its structure the reader is referred to a recent survey by Barto et al. [9]. Basics of universal algebra can be learned from the textbook [30]. In preliminaries to this paper we therefore focus on what is needed for our result.

2.1 The CSP

The 'AI' formulation of the CSP best fits our purpose. Fix a finite set A and let Γ be a *constraint language* over A, that is, a set — not necessarily finite — of relations over A. The (*nonuniform*) *Constraint Satisfaction Problem* (*CSP*) associated with language Γ is the problem $\text{CSP}(\Gamma)$, in which, an *instance* is a pair (V, C), where V is a set of variables; and C is a set of *constraints*, i.e. pairs $\langle \mathbf{s}, R\rangle$, where $\mathbf{s} = (v_1, \dots, v_k)$ is a tuple of variables from V, the *constraint scope*, and $R \in \Gamma$, the k-ary *constraint relation*. We always assume that relations are given explicitly by a list of tuples[1]. The goal is to find a *solution*, i.e., a mapping $\varphi : V \to A$ such that for every constraint $\langle \mathbf{s}, R\rangle \in C$, $\varphi(\mathbf{s}) \in R$.

2.2 Algebraic methods in the CSP

Jeavons et al. in [43, 44] were the first to observe that higher order symmetries of constraint languages, called polymorphisms, play a significant role in the study of the complexity of the CSP. A *polymorphism* of a relation R over A is an operation $f(x_1, \dots, x_k)$ on A such that for any choice of $\mathbf{a}_1, \dots, \mathbf{a}_k \in R$ we have $f(\mathbf{a}_1, \dots, \mathbf{a}_k) \in R$. If this is the case we also say that f *preserves* R, or that R is *invariant* with respect to f. A polymorphism of a constraint language Γ is an operation that is a polymorphism of every $R \in \Gamma$.

THEOREM 2.1 ([43, 44]). *For constraint languages Γ, Δ, where Γ is finite, if every polymorphism of Δ is also a polymorphism of Γ, then $\text{CSP}(\Gamma)$ is polynomial-time reducible to $\text{CSP}(\Delta)$.*[2]

Listed below are several types of polymorphisms that occur frequently throughout the paper. The presence of each of these polymorphisms imposes strong restrictions on the structure of invariant relations, which can be used in designing a solution algorithm. Some of such results will be mentioned later.

- *Semilattice* operation is a binary operation $f(x, y)$ such that $f(x, x) = x$, $f(x, y) = f(y, x)$, and $f(x, f(y, z)) = f(f(x, y), z)$ for all $x, y, z \in A$;
- k-ary *near-unanimity* operation is a k-ary operation $u(x_1, \dots, x_k)$ such that

$$u(y, x, \dots, x) = u(x, y, x, \dots, x) = \dots = u(x, \dots, x, y) = x$$

  for all $x, y \in A$; a ternary near-unanimity operation m is called a *majority* operation, it satisfies the equations $m(y, x, x) = m(x, y, x) = m(x, x, y) = x$;
- *Mal'tsev* operation is a ternary operation $h(x, y, z)$ satisfying the equations $h(x, y, y) = h(y, y, x) = x$ for all $x, y \in A$; the *affine* operation $x - y + z$ of an Abelian group is a special case of a Mal'tsev operation;

[1]The way constraints are represented does not matter if Γ is finite, but it may change the complexity of the problems for infinite languages.
[2]Using the $s - t$-Connectivity algorithm by Reingold [56] this reduction can be improved to a log-space one.

– k-ary *weak near-unanimity* operation is a k-ary operation w that satisfies the same equations as a near-unanimity operation $w(y, x, \dots, x) = \dots = w(x, \dots, x, y)$, except for the last one $(= x)$, and the equation $w(x, x, \dots, x) = x$.

To illustrate the effect of polymorphisms on the structure of invariant relations we give a few examples that involve polymorphisms of the types introduced above. First, we need some terminology and notation.

By $[n]$ we denote the set $\{1, \dots, n\}$. For sets $A_1, \dots, A_n$ tuples from $A_1 \times \dots \times A_n$ are denoted in boldface, say, $\mathbf{a}$; the ith component of $\mathbf{a}$ is referred to as $\mathbf{a}[i]$. An n-ary relation R over sets $A_1, \dots, A_n$ is any subset of $A_1 \times \dots \times A_n$. For $I = \{i_1, \dots, i_k\} \subseteq [n]$ by $\mathrm{pr}_I \mathbf{a}, \mathrm{pr}_I R$ we denote the *projections* $\mathrm{pr}_I \mathbf{a} = (\mathbf{a}[i_1], \dots, \mathbf{a}[i_k])$, $\mathrm{pr}_I R = \{\mathrm{pr}_I \mathbf{a} \mid \mathbf{a} \in R\}$ of tuple $\mathbf{a}$ and relation R. If $\mathrm{pr}_i R = A_i$ for each $i \in [n]$, relation R is said to be a *subdirect product* of $A_1 \times \dots \times A_n$. Sometimes it is convenient to label the coordinate positions of relations by elements of some set other than $[n]$, e.g. by variables of a CSP.

Example 2.2. (1) Let $\vee$ be the binary operation of disjunction on $\{0, 1\}$, as is easily seen, it is a semilattice operation. The following property of relations invariant under $\vee$ helps solving the corresponding CSP: A relation R contains the tuple $(1, \dots, 1)$ whenever for each coordinate position R contains a tuple with a 1 in that position. Similarly, relations invariant under other semilattice operations on larger sets always contain a sort of a 'maximal' tuple.

(2) By the results of [2] a tuple $\mathbf{a}$ belongs to a (n-ary) relation R invariant under a k-ary near-unanimity operation if and only if for every $(k-1)$-element set $I \subseteq [n]$ we have $\mathrm{pr}_I \mathbf{a} \in \mathrm{pr}_I R$. In particular, if f is the majority operation on $\{0, 1\}$ given by $(x \wedge y) \vee (y \wedge z) \vee (z \wedge x)$, and R is a relation on $\{0, 1\}$, then $\mathbf{a} \in R$ if and only if $(\mathbf{a}[i], \mathbf{a}[j]) \in \mathrm{pr}_{ij} R$ for all $i, j \in [n]$. This property easily gives rise to a reduction of the corresponding CSP to 2-SAT [45].

(3) If $m(x, y, z) = x - y + z$ is the affine operation of, say, $\mathbb{Z}_p$, p prime, then relations invariant with respect to m are exactly those that can be represented as solution sets of systems of linear equations over $\mathbb{Z}_p$, and the corresponding CSP can be solved by Gaussian Elimination. One direction is easy to see. If $R = \{\mathbf{x} \mid \mathbf{x} \cdot M = \mathbf{d}\}$, where M is the matrix of the system of equations, and $\mathbf{a}, \mathbf{b}, \mathbf{c} \in R$, then

$$(\mathbf{a} - \mathbf{b} + \mathbf{c}) \cdot M = \mathbf{a} \cdot M - \mathbf{b} \cdot M + \mathbf{c} \cdot M = \mathbf{d} - \mathbf{d} + \mathbf{d} = \mathbf{d},$$

implying $m(\mathbf{a}, \mathbf{b}, \mathbf{c}) \in R$. The other direction is more involved, see, e.g., [58, Proposition 2.1].

$\diamond$

The next step in discovering more structure behind nonuniform CSPs was made in [27], where universal algebras were brought into the picture. A *(universal) algebra* is a pair $\mathbb{A} = (A, F)$ consisting of a set A, the *universe* of $\mathbb{A}$, and a set F of operations on A. Operations from F (called *basic*) together with operations that can be obtained from them by means of composition are called the *term* operations of $\mathbb{A}$.

Algebras allow for a more general definition of CSPs than the one used above. Let $\mathrm{CSP}(\mathbb{A})$ denote the class of nonuniform CSPs $\{\mathrm{CSP}(\Gamma) \mid \Gamma \subseteq \mathsf{Inv}(F),\ \Gamma \text{ finite}\}$, where $\mathsf{Inv}(F)$ denotes the set of all relations invariant with respect to all operations from F. Note that the tractability of $\mathrm{CSP}(\mathbb{A})$ can be understood in two ways: as the existence of a polynomial-time algorithm for every $\mathrm{CSP}(\Gamma)$ from this class, or as the existence of a uniform polynomial-time algorithm for all such problems. One of the implications of our results is that these two types of tractability are the same. From the formal standpoint we will use the stronger one.

2.3 Structural features of universal algebras

We use some structural elements of algebras, of which the main ones are subalgebras, congruences, and quotient algebras. For $B \subseteq A$ and an operation f on A by $f|_B$ we denote the restriction of f on B. The algebra $\mathbb{B} = (B, \{f|_B \mid f \in F\})$ is a *subalgebra* of $\mathbb{A}$ if $f(b_1, \dots, b_k) \in B$ for any $b_1, \dots, b_k \in B$ and any $f \in F$.

Congruences play a very significant role in our algorithm, so we will now discuss them in more detail. A *congruence* is an equivalence relation $\alpha \in \mathsf{Inv}(F)$. This means that for any operation $f \in F$ and any $(a_1, b_1), \ldots, (a_k, b_k) \in \alpha$ it holds that $(f(a_1, \ldots, a_k), f(b_1, \ldots, b_k)) \in \alpha$. Hence one can define an algebra on $A/_\alpha$, the set of α-blocks, by setting $f/_\alpha(a_1/_\alpha, \ldots, a_k/_\alpha) = (f(a_1, \ldots, a_k))/_\alpha$ for $a_1, \ldots, a_k \in A$, where $a/_\alpha$ denotes the α-block containing a. The algebra $\mathbb{A}/_\alpha$ is called the *quotient algebra modulo* α. Often the fact that a, b are related by a congruence α is denoted $a \stackrel{\alpha}{\equiv} b$.

Example 2.3. The following are examples of congruences and quotient algebras.

(1) Let $\mathbb{A}$ be any algebra. Then the equality relation $\underline{0}_\mathbb{A}$ and the full binary relation $\underline{1}_\mathbb{A}$ on $\mathbb{A}$ are congruences of $\mathbb{A}$. The quotient algebra $\mathbb{A}/_{\underline{0}_\mathbb{A}}$ is $\mathbb{A}$ itself, while $\mathbb{A}/_{\underline{1}_\mathbb{A}}$ is a 1-element algebra.

(2) Let $\mathbb{L}_n$ be an n-dimensional vector space and $\mathbb{L}'$ a k-dimensional subspace of $\mathbb{L}_n$, $k \le n$. The binary relation π given by: $(\overline{a}, \overline{b}) \in \pi$ iff $\overline{a}, \overline{b}$ have the same orthogonal projection on $\mathbb{L}'$, is a congruence of $\mathbb{L}_n$ and $\mathbb{L}_n/_\pi$ is $\mathbb{L}'$.

(3) The next example will be our running example throughout the paper. Let $A = \{0, 1, 2\}$, and let $\mathbb{A}_M$ be the algebra with universe A and two basic operations: a binary operation r such that $r(0,0) = r(0,1) = r(2,0) = r(0,2) = r(2,1) = 0$, $r(1,1) = r(1,0) = r(1,2) = 1$, $r(2,2) = 2$; and a ternary operation t such that $t(x,y,z) = x - y + z$ if $x, y, z \in \{0,1\}$, where $+, -$ are the operations of $\mathbb{Z}_2$, $t(2,2,2) = 2$, and otherwise $t(x,y,z) = t(x',y',z')$, where $x' = x$ if $x \in \{0,1\}$ and $x' = 0$ if $x = 2$; the values y', z' are obtained from y, z by the same rule. It is an easy exercise to verify the following facts: (a) $\mathbb{B} = (\{0,1\}, r|_{\{0,1\}}, t|_{\{0,1\}})$ and $\mathbb{C} = (\{0,2\}, r|_{\{0,2\}}, t|_{\{0,2\}})$ are subalgebras of $\mathbb{A}_M$; (b) the partition $\{0,1\}, \{2\}$ is a congruence of $\mathbb{A}_M$, let us denote it θ; (c) algebra $\mathbb{C}$ is *term equivalent* to a semilattice (it has the same term operations as a semilattice), that is, a set with a semilattice operation, see Fig 1(a). The classes of congruence θ are $0/_\theta = \{0,1\}, 2/_\theta = \{2\}$. Then the quotient algebra $\mathbb{A}_M/_\theta$ is also term equivalent to a semilattice, as $r/_\theta(0/_\theta, 0/_\theta) = r/_\theta(0/_\theta, 2/_\theta) = r/_\theta(2/_\theta, 0/_\theta) = 0/_\theta$ and $r/_\theta(2/_\theta, 2/_\theta) = 2/_\theta$.

$\diamond$

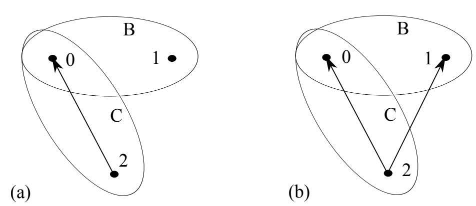

Fig. 1. (a) Algebra $\mathbb{A}_M$. (b) Algebra $\mathbb{A}_N$. Dots represent elements, ovals represent subalgebras, and arrows represent semilattice edges (see Section 3.2).

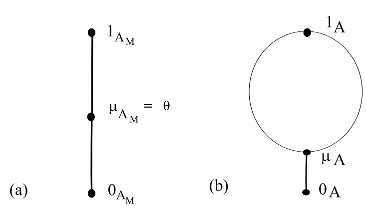

Fig. 2. (a) The congruence lattice of algebra $\mathbb{A}_M$; (b) congruence lattice of a subdirectly irreducible algebra.

The set of all congruences of $\mathbb{A}$ ordered with respect to inclusion is denoted by $\mathsf{Con}(\mathbb{A})$. This set is actually a lattice, that is, the operations of meet $\wedge$ and join $\vee$ can be defined so that $\alpha \wedge \beta$ is the greatest lower bound of $\alpha, \beta \in \mathsf{Con}(\mathbb{A})$ and

$\alpha \vee \beta$ is the least upper bound of α, β. In the case of congruence lattices $\mathrm{Con}(\mathbb{A})$ the meet is given by $\alpha \wedge \beta = \alpha \cap \beta$ and the join $\alpha \vee \beta$ is the transitive closure of $\alpha \cup \beta$. Fig. 2(a) shows $\mathrm{Con}(\mathbb{A}_M)$ for the algebra $\mathbb{A}_M$ from Example 2.3(3). By $\mathrm{HS}(\mathbb{A})$ we denote the set of all quotient algebras of all subalgebras of $\mathbb{A}$.

2.4 The Dichotomy Conjecture

The results of [27] reduce the dichotomy conjecture to idempotent algebras. An algebra $\mathbb{A} = (A, F)$ is said to be *idempotent* if every operation $f \in F$ satisfies the equation $f(x, \dots, x) = x$. If $\mathbb{A}$ is idempotent, then all the *constant* relations $\{(a)\}$ are invariant under F. Therefore studying CSPs over idempotent algebras is the same as studying CSPs that allow all constant relations. Another useful property of idempotent algebras is that every block of each of its congruences is a subalgebra. We now can state the algebraic version of the dichotomy theorem.

THEOREM 2.4. *For a finite idempotent algebra* $\mathbb{A}$ *the following are equivalent (assuming P≠NP):*

(1) $\mathrm{CSP}(\mathbb{A})$ *is solvable in polynomial time;*
(2) $\mathbb{A}$ *has a weak near-unanimity term operation;*
(3) *every algebra from* $\mathrm{HS}(\mathbb{A})$ *has a nontrivial term operation (that is, not a* projection*, an operation of the form* $f(x_1, \dots, x_k) = x_i$*).*

Otherwise $\mathrm{CSP}(\mathbb{A})$ *is NP-complete.*

The hardness part of this theorem is proved in [27]; the equivalence of (2) and (3) was proved in [13] and [55]. That conditions (2) and (3) imply (1) is the main result of this paper. In the rest of the paper we assume all algebras to satisfy conditions (2),(3).

In fact, we will prove a slightly more general result. Let $\mathcal{A}$ be a finite class of finite idempotent *similar* algebras, that is, whose basic operations have the same 'names' and the corresponding arities. If one is only interested in CSPs arising from a single algebra $\mathbb{A}$, one may assume that such a class is produced from $\mathbb{A}$ by taking subalgebras, quotient algebras and also *retractions* introduced in Section 5.5. Then $\mathrm{CSP}(\mathcal{A})$ denotes the class of CSP instances whose variables can have different domains belonging to $\mathcal{A}$, see, e.g. [15]. We will design an algorithm for $\mathrm{CSP}(\mathcal{A})$ whenever there is a weak near-unanimity term for all algebras in $\mathcal{A}$ simultaneously.

3 BOUNDED WIDTH AND THE FEW SUBPOWERS ALGORITHM

Leaving aside occasional combinations thereof, there are only two standard types of algorithms solving the CSP. In this section we give a brief introduction into both of them.

3.1 CSPs of bounded width

Algorithms of the first kind are based on the idea of local propagation, that is formally described below.

Let $\mathcal{P} = (V, \mathcal{C})$ be a CSP instance. For $W \subseteq V$ by $\mathcal{P}_W$ we denote the *restriction* of $\mathcal{P}$ onto W, that is, the instance $(W, \mathcal{C}_W)$, where for each $C = \langle \mathbf{s}, R \rangle \in \mathcal{C}$, the set $\mathcal{C}_W$ includes the constraint $C_W = \langle \mathbf{s} \cap W, \mathrm{pr}_{\mathbf{s} \cap W} R \rangle$, where $\mathbf{s} \cap W$ is the subtuple of $\mathbf{s}$ containing all the elements from W in $\mathbf{s}$, say, $\mathbf{s} \cap W = (i_1, \dots, i_k)$, and $\mathrm{pr}_{\mathbf{s} \cap W} R$ stands for $\mathrm{pr}_{\{i_1, \dots, i_k\}} R$. The set of solutions of $\mathcal{P}_W$ will be denoted by $\mathcal{S}_W$. If W is a singleton, say, $W = \{v\}$, we will abuse the notation and write $\mathcal{P}_v, \mathcal{S}_v$ rather than $\mathcal{P}_{\{v\}}, \mathcal{S}_{\{v\}}$.

Unary solutions, that is, when $|W| = 1$ play a special role. As is easily seen, for $v \in V$ the set $\mathcal{S}_v$ is just the intersections of unary projections $\mathrm{pr}_v R$ of constraints whose scope contains v. An instance $\mathcal{P}$ is said to be *1-minimal* if for every $v \in V$

and every constraint $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ such that $v \in \mathbf{s}$, it holds that $\mathrm{pr}_v R = \mathcal{S}_v$. For a 1-minimal instance one may always assume that the set of allowed values for a variable $v \in V$ is the set $\mathcal{S}_v$. We call this set the *domain* of v and assume that a CSP instance may have different domains for different variables, which nevertheless are always subalgebras or quotient algebras of the original algebra $\mathbb{A}$. It will be convenient to denote the domain of v by $\mathbb{A}_v$. The domain $\mathbb{A}_v$ may change as a result of transformations of the instance.

An instance $\mathcal{P}$ is said to be *(2,3)-consistent* if it has a *(2,3)-strategy*, that is, a collection of relations R^X, for every $X \subseteq V$, $|X| = 2$ satisfying the following conditions (we use R^v, R^{vw} for $R^{\{v\}}, R^{\{v,w\}}$):

- for every $X \subseteq V$ with $|X| \le 2$, $R^X \subseteq \mathcal{S}_X$;
- for every $X = \{u, v\} \subseteq V$, any $w \in V - X$, and any $(a, b) \in R^X$, there is $c \in \mathbb{A}_w$ such that $(a, c) \in R^{uw}$ and $(b, c) \in R^{vw}$.

Let the collection of relations R^X be denoted by $\mathcal{R}$. A tuple $\mathbf{a}$ whose entries are indexed with elements of $W \subseteq V$ and such that $\mathrm{pr}_X \mathbf{a} \in R^X$ for any $X \subseteq W$, $|X| = 2$, will be called *$\mathcal{R}$-compatible*. If a (2,3)-consistent instance $\mathcal{P}$ with a (2,3)-strategy $\mathcal{R}$ satisfies the additional condition

- for every constraint $C = \langle \mathbf{s}, R\rangle$ of $\mathcal{P}$ every tuple $\mathbf{a} \in R$ is $\mathcal{R}$-compatible,

it is called *(2,3)-minimal*. For $k \in \mathbf{N}$, $(k, k+1)$-strategies, $(k, k+1)$-consistency, and $(k, k+1)$-minimality are defined in a similar way replacing 2,3 with $k, k+1$. We will always assume that a (2,3)-consistent or (2,3)-minimal instance contains a constraint with every 1- and 2-element scope.

Instance $\mathcal{P}$ is said to be *minimal* (or *globally minimal*) if for every $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and every $\mathbf{a} \in R$ there is a solution $\varphi \in \mathcal{S}$ such that $\varphi(\mathbf{s}) = \mathbf{a}$. Similarly, $\mathcal{P}$ is said to be *globally 1-minimal* if for every $v \in V$ and $a \in \mathbb{A}_v$ there is a solution φ with $\varphi(v) = a$.

Any instance can be transformed to a 1-minimal, (2,3)-consistent, or (2,3)-minimal instance in polynomial time using standard constraint propagation algorithms (see, e.g. [33]). These algorithms work by changing the constraint relations and the domains of the variables eliminating some tuples and elements from them. We call such a process *tightening* the instance. It is important to notice that if the original instance belongs to $\mathrm{CSP}(\mathbb{A})$ for some algebra $\mathbb{A}$, that is, all its constraint relations are invariant under the basic operations of $\mathbb{A}$, the constraint relations obtained by propagation algorithms are also invariant under the basic operations of $\mathbb{A}$, and so the resulting instance also belongs to $\mathrm{CSP}(\mathbb{A})$. Establishing minimality amounts to solving the problem and therefore is not generally expected to be doable in polynomial time.

If a constraint propagation algorithm solves a CSP, the problem is said to be of bounded width. More precisely, $\mathrm{CSP}(\Gamma)$ (or $\mathrm{CSP}(\mathbb{A})$) is said to have *bounded width* if for some k every $(k, k+1)$-minimal instance from $\mathrm{CSP}(\Gamma)$ (or $\mathrm{CSP}(\mathbb{A})$) has a solution. Problems of bounded width are very well studied, see an older survey [28] and a more recent paper [4].

THEOREM 3.1 ([4, 16, 21, 48]). *For an idempotent algebra $\mathbb{A}$ the following are equivalent:*

(1) $\mathrm{CSP}(\mathbb{A})$ *has bounded width;*
(2) *every (2,3)-minimal instance from* $\mathrm{CSP}(\mathbb{A})$ *has a solution;*
(3) $\mathbb{A}$ *has a weak near-unanimity term of arity k for every $k \ge 3$;*
(4) *every algebra* $\mathsf{HS}(\mathbb{A})$ *has a nontrivial operation, and none of them is equivalent to a module (in a certain precise sense that will be explained in Section 6.1).*

3.2 Omitting semilattice edges and the few subpowers property

The second type of CSP algorithms can be viewed as a generalization of Gaussian elimination, although, it utilizes just one property also used by Gaussian elimination: the set of solutions of a system of linear equations (or a CSP) has a set of generators of size polynomial in the number of variables. The property that for every instance $\mathcal{P}$ of CSP($\mathbb{A}$) its solution space $\mathcal{S}$ has a set of generators of polynomial size is nontrivial, because there are only exponentially many such sets, while, as is easily seen, CSPs may have up to double exponentially many different sets of solutions. Formally, an algebra $\mathbb{A} = (A, F)$ has *few subpowers* if for every n there are only exponentially many n-ary relations in $\mathsf{Inv}(F)$.

Algebras with few subpowers are well studied and the CSP over such an algebra has a polynomial-time solution algorithm, see, [10, 42]. In particular, such algebras admit a characterization in terms of the existence of a term operation with special properties, an *edge* term. We need only a subclass of algebras with few subpowers that appeared in [21, 24] and is defined as follows.

An pair ab (or (a, b)) of elements $a, b \in \mathbb{A}$ is said to be a *semilattice edge* if there is a binary term operation f of $\mathbb{A}$ such that $f(a, a) = a$ and $f(a, b) = f(b, a) = f(b, b) = b$, that is, f is a semilattice operation on $\{a, b\}$. For example, the pair $(2, 0)$ from Example 2.3(3) is a semilattice edge, and the operation r of $\mathbb{A}_M$ witnesses that. An algebra is said to be *semilattice free* if it contains no semilattice edges.

Proposition 3.2 ([21, 24]). *If an idempotent algebra $\mathbb{A}$ semilattice free, it has few subpowers, and therefore* CSP($\mathbb{A}$) *is solvable in polynomial time.*

Semilattice edges have other useful properties including the following one that we use for reducing a CSP to smaller problems.

Lemma 3.3 (Proposition 4.2, [20]). *For any idempotent algebra $\mathbb{A}$ there is a binary term operation xy of $\mathbb{A}$ (think multiplication)*[3] *such that xy is a semilattice operation on any semilattice edge and for any $a, b \in \mathbb{A}$ either $ab = a$ or (a, ab) is a semilattice edge.*

Note that any semilattice operation satisfies the conditions of Lemma 3.3. The operation r of the algebra $\mathbb{A}_M$ from Example 2.3(3) is not a semilattice operation (for instance, it does not satisfy the equation $r(x, y) = r(y, x)$), but it satisfies the conditions of Lemma 3.3.

4 CENTRALIZERS AND DECOMPOSITION OF CSPS

In this section we introduce an alternative definition of the centralizer operator on congruence lattices studied in commutator theory, and study its properties and its connection to decompositions of CSPs. Unlike the vast majority of the literature on the algebraic approach to the CSP we use not only term operations, but also polynomial operations of an algebra. It should be noted however that the first to use polynomials for CSP algorithms was Maróti in [54]. We make use of some ideas from that paper in the next section.

Let $f(x_1, \dots, x_k, y_1, \dots, y_\ell)$ be a $k + \ell$-ary term operation of an algebra $\mathbb{A} = (A, F)$ and $b_1, \dots, b_\ell \in \mathbb{A}$. The operation $g(x_1, \dots, x_k) = f(x_1, \dots, x_k, b_1, \dots, b_\ell)$ is called a *polynomial* of $\mathbb{A}$. The name 'polynomial' is used because of the close connection with polynomials in the standard sense. Indeed, if $\mathbb{A}$ is a ring, its polynomials as just defined are the same as polynomials in the regular sense. A polynomial that depends on only one variable, i.e. $k = 1$, is said to be a *unary* polynomial.

[3] In order to avoid confusion between the edge xy and the binary operation xy we always specify that xy is an edge whenever it is the case.

While polynomials of $\mathbb{A}$ do not have to be polymorphisms of relations from $\mathsf{Inv}(F)$, congruences and unary polynomials are in a special relationship. More precisely, it is a well-known fact (see, e.g., [30]) that an equivalence relation over $\mathbb{A}$ is a congruence if and only if it is preserved by all the unary polynomials of $\mathbb{A}$. If α is a congruence, and f is a unary polynomial, by $f(\alpha)$ we denote the set of pairs $\{(f(a), f(b)) \mid (a,b) \in \alpha\}$.

Example 4.1. The unary polynomials of the algebra $\mathbb{A}_M$ from Example 2.3(3) include the following unary operations (these are the polynomials we will use, there are more unary polynomials of $\mathbb{A}_M$):
$h_1(x) = r(x,0) = r(x,1)$, such that $h_1(0) = h_1(2) = 0, h_1(1) = 1$;
$h_2(x) = r(2,x)$, such that $h_2(0) = h_2(1) = 0$, $h_2(2) = 2$;
$h_3(x) = r(0,x) = 0$.
The lattice $\mathrm{Con}(\mathbb{A}_M)$ has 3 congruences: $\underline{0}, \theta, \underline{1}$ (see Example 2.3(3)). As is easily seen, $h_1(\theta) \not\subseteq \underline{0}$, $h_2(\underline{1}) \not\subseteq \theta$, but $h_1(\underline{1}) \subseteq \theta$, $h_2(\theta) \subseteq \underline{0}$, $h_3(\underline{1}) \subseteq \underline{0}$. ⋄

For an algebra $\mathbb{A}$, a term operation $f(x, y_1, \dots, y_k)$, and $\mathbf{a} \in \mathbb{A}^k$, let $f^{\mathbf{a}}(x) = f(x, \mathbf{a})$. Let $\alpha, \beta \in \mathrm{Con}(\mathbb{A})$, $\alpha \le \beta$, and let $(\alpha : \beta) \subseteq \mathbb{A}^2$ denote the greatest congruence such that for any term operation $f(x, y_1, \dots, y_k)$ and any $\mathbf{a}, \mathbf{b} \in \mathbb{A}^k$ such that $(\mathbf{a}[i], \mathbf{b}[i]) \in (\alpha : \beta)$ for all $i \in [k]$, it holds that $f^{\mathbf{a}}(\beta) \subseteq \alpha$ if and only if $f^{\mathbf{b}}(\beta) \subseteq \alpha$. Polynomials of the form $f^{\mathbf{a}}, f^{\mathbf{b}}$ are often called twin polynomials.

The congruence $(\alpha : \beta)$ will be called the *centralizer*[4] of α, β. The following statement is one of the key ingredients of the algorithm.

LEMMA 4.2 (COROLLARY 50 [25]). *Let $(\alpha : \beta) = \underline{1}_{\mathbb{A}}$, $a, b, c \in \mathbb{A}$ and $b \stackrel{\beta}{\equiv} c$. Then $(ab, ac) \in \alpha$, where multiplication is as in Lemma 3.3.*

Example 4.3. In the algebra $\mathbb{A}_M$, see Example 2.3(3), the centralizer acts as follows: $(\underline{0} : \theta) = \underline{1}$ and $(\theta : \underline{1}) = \theta$. We start with the second centralizer. Since every polynomial preserves congruences, for any term operation $h(x, y_1, \dots, y_k)$ and any $\mathbf{a}, \mathbf{b} \in \mathbb{A}_M^k$ such that $(\mathbf{a}[i], \mathbf{b}[i]) \in \theta$ for $i \in [k]$, we have $(h^{\mathbf{a}}(x), h^{\mathbf{b}}(x)) \in \theta$ for any x. This of course implies $(\theta : \underline{1}) \ge \theta$. On the other hand, let $f(x,y) = r(y,x)$. Then

$$f^0(x) = f(x,0) = r(0,x) = h_3(x),$$
$$f^2(x) = f(x,2) = r(2,x) = h_2(x),$$

and $f^0(\underline{1}) \subseteq \theta$, while $f^2(\underline{1}) \not\subseteq \theta$. This means that $(0,2) \notin (\theta : \underline{1})$ and so $(\theta : \underline{1}) \subset \underline{1}$.

For the first centralizer it suffices to demonstrate that the condition in the definition of centralizer is satisfied for pairs of twin polynomials of the form $(r(a,x), r(b,x))$, $(r(x,a), r(x,b))$, $(t(x, a_1, a_2), t(x, b_1, b_2))$, $(t(a_1, x, a_2), t(b_1, x, b_2))$, $(t(a_1, a_2, x), t(b_1, b_2, x))$ for $a, b, a_1, a_2, b_1, b_2 \in \{0,1,2\}$, which can be verified directly.

Interestingly, Lemma 4.2 implies that if we change the operation r in just one point, it has a profound effect on the centralizer $(\underline{0} : \theta)$. Let $\mathbb{A}_N$ be the same algebra as $\mathbb{A}_M$ with operations r', t' defined in the same way as r, t, except $r'(2,1) = 1$ replacing the value $r(2,1) = 0$. In this case $\{1,2\}$ is also a semilattice edge, see Fig. 1(b). Let again

[4]Traditionally, the centralizer of two congruences is defined in a different way, see, e.g. [36]. Congruence $(\alpha : \beta)$ appeared in [41], but completely inconsequentially, they did not study it at all, and its relation to the standard notion of centralizer remained unknown. We used the current definition in [22] and called it quasi-centralizer, again, not completely aware of its connection to the standard centralizer. Later Willard [60] showed that the two concepts are equivalent, see [25, Proposition 46] for a proof, and we use 'centralizer' here rather than 'quasi-centralizer'.

$f(x,y) = r'(y,x)$ and $a = 0, b = 2$. This time we have

$$f^0(x) = f(x,0) = r'(0,x) = h'_3(x),$$
$$f^2(x) = f(x,2) = r'(2,x) = h'_2(x),$$

where $h'_3(x) = 0$ for all $x \in \{0,1,2\}$ and $h'_2(0) = 0, h'_2(1) = 1$ showing that $f^0(\theta) \subseteq \underline{0}$, while $f^2(\theta) \not\subseteq \underline{0}$. ⋄

Fig. 3(a),(b) shows the effect of large centralizers $(\alpha : \beta)$ on the structure of algebra $\mathbb{A}$, which is a generalization of the phenomena observed in Example 4.3. Dots there represent α-blocks (assume α is the equality relation), ovals represent β-blocks, let they be B and C, and such that there is at least one semilattice edge between B and C. If $(\alpha : \beta)$ is the full relation, Lemmas 3.3 and 4.2 imply that for any $a \in B$ and any $b, c \in C$ we have $ab = ac$, and so ab is the only element of C such that (a, ab) is a semilattice edge (represented by arrows). In other words, we have a mapping from B to C that can also be shown to be injective. We will use this mapping to lift any solution with a value from B to a solution with a value from C.

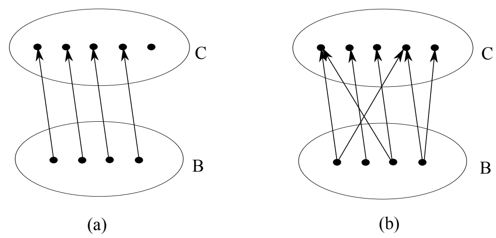

Fig. 3. (a) $(\alpha : \beta)$ is the full relation; (b) $(\alpha : \beta)$ is not the full relation

Finally, we prove an easy corollary from Lemma 4.2.

Corollary 4.4. *Let $\alpha, \beta \in \mathsf{Con}(\mathbb{A})$, $\alpha \le \beta$, be such that $(\alpha : \beta) \ge \beta$. Then for every β-block B if (a,b) is a semilattice edge and $a, b \in B$, then $a \overset{\alpha}{\equiv} b$.*

Proof. Let $a, b \in \mathbb{B}$, $a \overset{\alpha}{\not\equiv} b$, form a semilattice edge, that is, $ab = ba = b$. However, since $a \overset{(\alpha:\beta)}{\equiv} b$, by Lemma 4.2 applied to the subalgebra of $\mathbb{A}$ with the universe B it must hold that $aa \overset{\alpha}{\equiv} ab$, a contradiction. □

5 THE ALGORITHM

In this section we introduce the reductions used in the algorithm, and then explain the algorithm itself. The reductions heavily use the algebraic structure of the domains of an instance, and the structure of the instance itself.

5.1 Decomposition of CSPs

We have seen in the previous section that large centralizers impose strong restrictions on the structure of an algebra. We start this section by introducing some concepts that will be used later to show the effect of 'small' centralizers on the structure of CSP instances.

Let R be a binary relation, a subdirect product of $\mathbb{A} \times \mathbb{B}$, and $\alpha \in \mathsf{Con}(\mathbb{A})$, $\gamma \in \mathsf{Con}(\mathbb{B})$. Relation R is said to be *$\alpha\gamma$-aligned* if, for any $(a,c), (b,d) \in R$, $(a,b) \in \alpha$ if and only if $(c,d) \in \gamma$. This means that if $A_1, \dots, A_k$ are the α-blocks of $\mathbb{A}$, then there are also k γ-blocks of $\mathbb{B}$ and they can be labeled $B_1, \dots, B_k$ in such a way that

$$R = (R \cap (A_1 \times B_1)) \cup \dots \cup (R \cap (A_k \times B_k)).$$

This definition provides a way to decompose CSP instances. Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal instance from CSP($\mathbb{A}$). Recall that we assume that a (2,3)-consistent or (2,3)-minimal instance has a constraint $C^X = \langle X, R^X = \mathcal{S}_X\rangle$ for every $X \subseteq V$, $|X| \le 2$. So, C contains a constraint $C^{vw} = \langle (v, w), R^{vw}\rangle$ for every $v, w \in V$, and these relations form a (2,3)-strategy for $\mathcal{P}$. Recall that $\mathbb{A}_v$ denotes the domain of $v \in V$. Let $W \subseteq V$ and $\alpha_v \in \mathsf{Con}(\mathbb{A}_v) - \{\underline{1}_v\}$, $v \in W$, be such that for any $v, w \in W$ the relation R^{vw} is $\alpha_v\alpha_w$-aligned. The set W is then called a *strand* of $\mathcal{P}$. We will also say that $\mathcal{P}_W$ is *$\overline{\alpha}$-aligned*.

For a strand W and congruences α_v as above, there is a one-to-one correspondence between α_v- and α_w-blocks of $\mathbb{A}_v$ and $\mathbb{A}_w$, $v, w \in W$. Moreover, by (2,3)-minimality these correspondences are consistent, that is, if $u, v, w \in W$ and B_u, B_v, B_w are α_u-, α_v- and α_w-blocks, respectively, such that $R^{uv} \cap (B_u \times B_v) \neq \varnothing$ and $R^{vw} \cap (B_v \times B_w) \neq \varnothing$, then $R^{uw} \cap (B_u \times B_w) \neq \varnothing$. This means that $\mathcal{P}_W$ can be split into several instances, whose domains are α_v-blocks.

LEMMA 5.1. *Let $\mathcal{P}, W, \alpha_v$ for each $v \in W$, be as above. Then $\mathcal{P}_W$ can be decomposed into a collection of instances $\mathcal{P}_1, \dots, \mathcal{P}_k$, k constant, $\mathcal{P}_i = (W, C_i)$ such that every solution of $\mathcal{P}_W$ is a solution of one of the $\mathcal{P}_i$ and for every $v \in W$ its domain in $\mathcal{P}_i$ is an α_v-block.*

All the strands of a CSP instance $\mathcal{P}$ can be efficiently found. Indeed, as is easily seen, for any variables $v, w \in V$ and any $\alpha \in \mathsf{Con}(\mathbb{A}_v)$, there is at most one congruence $\beta \in \mathsf{Con}(\mathbb{A}_w) - \{\underline{1}_w\}$ such that R^{vw} is $\alpha\beta$-aligned. This means, there is only one maximal (with respect to inclusion) strand $W \subseteq V$ witnessed by congruences $\alpha_w \in \mathsf{Con}(\mathbb{A}_w) - \{\underline{1}_w\}$ and such that $v \in W$ and $\alpha_v = \alpha$. In particular, a CSP instance can only have linearly many maximal strands, as every domain has only a constant number of congruences. Moreover, this observation suggests a way to find all the strands: for each v and each $\alpha \in \mathsf{Con}(\mathbb{A}_v) - \{\underline{1}_v\}$ construct a set $W \subseteq V$ that contains all $w \in V$ such that R^{vw} is $\alpha\beta$-aligned for some $\beta \in \mathsf{Con}(\mathbb{A}_w) - \{\underline{1}_w\}$.

Example 5.2. Let $\mathbb{A}_M$ be the algebra introduced in Example 2.3(3), and let R be the following ternary relation over $\mathbb{A}_M$ invariant under r, t, given by

$$R = \begin{pmatrix} 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 2 & 2 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 2 & 2 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 2 \end{pmatrix},$$

where triples, the elements of the relation, are written vertically. Consider the following simple CSP instance from CSP($\mathbb{A}_M$): $\mathcal{P} = (V = \{v_1, v_2, v_3, v_4, v_5\}, \{C^1 = \langle \mathbf{s}_1 = (v_1, v_2, v_3), R_1\rangle, C^2 = \langle \mathbf{s}_2 = (v_2, v_4, v_5), R_2\rangle\}$, where $R_1 = R_2 = R$. To make the instance (2,3)-minimal we run the appropriate local propagation algorithm on it. First, such an algorithm adds new binary constraints $C^{v_iv_j} = \langle (v_i, v_j), R^{v_iv_j}\rangle$ for $i, j \in [5]$ starting with $R^{v_iv_j} = \mathbb{A}_M \times \mathbb{A}_M$. It then iteratively removes pairs from these relations that do not satisfy the (2,3)-minimality condition. Similarly, it tightens the original constraint relations if they violate the conditions of (2,3)-minimality. It is not hard to see that this algorithm does not change constraints C^1, C^2, and that the new binary relations are as follows: $R^{v_1v_2} = R^{v_2v_4} = R^{v_1v_4} = \theta$, $R^{v_1v_3} = R^{v_2v_3} = R^{v_2v_5} = R^{v_4v_5} = Q$, and $R^{v_1v_5} = R^{v_3v_4} = R^{v_3v_5} = S$, where

$$Q = \mathrm{pr}_{13}R = \begin{pmatrix} 0 & 0 & 1 & 1 & 2 & 2 \\ 0 & 1 & 0 & 1 & 0 & 2 \end{pmatrix}, \qquad S = \begin{pmatrix} 0 & 0 & 1 & 1 & 0 & 2 & 2 \\ 0 & 1 & 0 & 1 & 2 & 0 & 2 \end{pmatrix}.$$

In order to distinguish elements and congruences of domains belonging to different variables let the domain of v_i be denoted by $\mathbb{A}_i$, its elements by $0_i, 1_i, 2_i$, and the congruences of $\mathbb{A}_i$ by $\underline{0}_i, \theta_i, \underline{1}_i$.

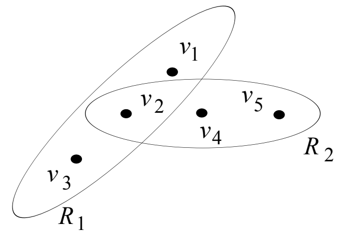

Fig. 4. Instance $\mathcal{P}$ from Example 5.2

Let $W = \{v_1, v_2, v_4\}$, $\alpha_i = \theta_i$ for $v_i \in W$. Then, since $R^{v_1v_2} = R^{v_2v_4} = R^{v_1v_4} = \theta$ and therefore are $\alpha_i\alpha_j$-aligned, $i, j \in \{1, 2, 4\}$, W is a strand of $\mathcal{P}$. Therefore the instance $\mathcal{P}_W = (\{v_1, v_2, v_4\}, \{C_W^1 = \langle (v_1, v_2), \mathrm{pr}_{v_1v_2} R_1\rangle, C_W^2 = \langle (v_2, v_4), \mathrm{pr}_{v_2v_4} R_2\rangle\})$ can be decomposed into a disjoint union of two instances

$$\mathcal{P}_1 = (\{v_1, v_2, v_4\}, \{\langle (v_1, v_2), Q_1\rangle, \langle (v_2, v_4), Q_2\rangle),$$
$$\mathcal{P}_2 = (\{v_1, v_2, v_4\}, \{\langle (v_1, v_2), S_1\rangle, \langle (v_2, v_4), S_2\rangle),$$

where $Q_1 = \{0_1, 1_1\} \times \{0_2, 1_2\}, Q_2 = \{0_2, 1_2\} \times \{0_4, 1_4\}, S_1 = \{(2_1, 2_2)\}, S_2 = \{(2_2, 2_4)\}$. $\diamond$

5.2 Irreducibility

In order to formulate the algorithm properly we need one more transformation of algebras. An algebra $\mathbb{A}$ is said to be *subdirectly irreducible* if the intersection of all its nontrivial (different from the equality relation) congruences is nontrivial. This smallest nontrivial congruence $\mu_{\mathbb{A}}$ is called the *monolith* of $\mathbb{A}$, see Fig. 2(b). For instance, the algebra $\mathbb{A}_M$ from Example 2.3(3) is subdirectly irreducible, because all its nontrivial congruences contain the nontrivial congruence θ. It is a folklore observation that any CSP instance can be transformed in polynomial time to an instance, in which the domain of every variable is a subdirectly irreducible algebra. One way to see this is to observe that if $\mathbb{A}$ is a finite algebra and $\theta_1, \ldots, \theta_k$ its minimal nontrivial congruences, then any constraint language Γ invariant under $\mathbb{A}$ can be pp-interpreted by a constraint language over $\mathbb{A}/_{\theta_1}, \ldots, \mathbb{A}/_{\theta_k}$, see, e.g., [9]. We will assume this property of all the instances we consider.

5.3 Block-minimality

Using Lemma 5.1 we introduce a new type of consistency of a CSP instance, block-minimality, which will be crucial for our algorithm. In a certain sense it is similar to the standard local consistency notions, as it is also defined through a family of relations that have to be consistent in a certain way. However, block-minimality is not quite local, and is more difficult to establish, as it involves solving smaller CSP instances recursively. The definitions below are designed to allow for an efficient procedure to establish block-minimality. This is achieved either by allowing for decomposing a subinstance into instances over smaller domains as in Lemma 5.1, or by replacing large domains with their quotient algebras.

Let α_v be a congruence of $\mathbb{A}_v$ for $v \in V$. By $\mathcal{P}/_{\overline{\alpha}}$ we denote the instance $(V, C_{\overline{\alpha}})$ constructed as follows: the domain of $v \in V$ is $\mathbb{A}_v/_{\alpha_v}$; for every constraint $C = \langle \mathbf{s}, R\rangle \in C$, $\mathbf{s} = (v_1, \ldots, v_k)$, the set $C_{\overline{\alpha}}$ includes the constraint $\langle \mathbf{s}, R/_{\overline{\alpha}}\rangle$, where $R/_{\overline{\alpha}} = \{(\mathbf{a}[v_1]/_{\alpha_{v_1}}, \ldots, \mathbf{a}[v_k]/_{\alpha_{v_k}}) \mid \mathbf{a} \in R\}$.

Example 5.3. Consider the instance $\mathcal{P}$ from Example 5.2, and let $\alpha_{v_i} = \theta_i$ (recall the θ_i is the congruence θ introduced in Example 2.3(3)) for each $i \in [5]$. Then $\mathcal{P}/_{\overline{\alpha}}$ is the instance over $\mathbb{A}_M/_\theta$ given by $\mathcal{P}/_{\overline{\alpha}} = (V, \{\langle \mathbf{s}_1, R_1/_{\overline{\alpha}}\rangle, \langle \mathbf{s}_2, R_2/_{\overline{\alpha}}\rangle\})$, where

$$R_1/_{\overline{\alpha}} = R_2/_{\overline{\alpha}} = \begin{pmatrix} 0/_\theta & 2/_\theta & 2/_\theta \\ 0/_\theta & 2/_\theta & 2/_\theta \\ 0/_\theta & 0/_\theta & 2/_\theta \end{pmatrix}.$$

◇

Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal instance with subdirectly irreducible domains. As $\mathcal{P}$ is (2,3)-minimal, for $X \subseteq V$, $|X| \leq 2$, it contains a constraint $C^X = \langle X, R^X \rangle$, where R^X is the set of partial solutions on X.

Recall that an algebra $\mathbb{A}_v$ is said to be semilattice free if it does not contain semilattice edges. Let $\mathsf{size}(\mathcal{P})$ denote the maximal size of domains of $\mathcal{P}$ that are not semilattice free and $\mathsf{MAX}(\mathcal{P})$ be the set of variables $v \in V$ such that $|\mathbb{A}_v| = \mathsf{size}(\mathcal{P})$ and $\mathbb{A}_v$ is not semilattice free. Finally, for $Y \subseteq V$ let $\mu_v^Y = \mu_v$ if $v \in Y$ and $\mu_v^Y = \underline{0}_v$ otherwise.

An instance $\mathcal{P}$ is said to be *block-minimal* if

(BM) for every strand $U \subseteq V$ the problem $\mathcal{P}_{/U}$, where $\mathcal{P}_{/U} = \mathcal{P}/_{\overline{\mu}^Y}$ and $Y = \mathsf{MAX}(\mathcal{P}) - U$, is minimal.

The definition of block-minimality is designed in such a way that block-minimality can be efficiently established. Observe that a strand can be large, even equal to V. However, following the approach of Lemma 5.1 we will show that $\mathcal{P}_{/U}$ splits into a union of disjoint problems over smaller domains.

Example 5.4. Let us consider again the instance $\mathcal{P}$ from Example 5.2. In that example we found all its binary solutions, and now we use them to find strands and to verify that this instance is block-minimal. As we saw in Example 5.2, unless $i, j \in \{1, 2, 4\}$ the relation $R^{v_i v_j}$ equals Q or S, which are not $\alpha\beta$-aligned for any congruences α, β except the full ones. This means that the only strands of $\mathcal{P}$ are $W = \{v_1, v_2, v_4\}$ and all the 1-element sets of variables.

Now we check the condition (BM) for $\mathcal{P}$. Consider W. For this strand we have $Y = \{3, 5\}$, and so $\mu_1^Y = \mu_2^Y = \mu_4^Y = \underline{0}$ and $\mu_3^Y = \mu_5^Y = \theta$. The problem $\mathcal{P}_{/W}$ is the following problem: $(V, \{C_1', C_2'\})$, where $C_1' = \langle \mathbf{s}_1, R^\theta \rangle$, $C_2' = \langle \mathbf{s}_2, R^\theta \rangle$, and

$$R^\theta = \begin{pmatrix} 0 & 0 & 1 & 1 & 2 & 2 \\ 0 & 1 & 0 & 1 & 2 & 2 \\ 0/_\theta & 0/_\theta & 0/_\theta & 0/_\theta & 0/_\theta & 2/_\theta \end{pmatrix}.$$

Now, consider first C_1'. For any tuple $(a_1, a_2, a_3) \in R^\theta$, that is, assignment $v_1 = a_1 \in \mathbb{A}_M, v_2 = a_2 \in \mathbb{A}_M, v_3 = a_3 \in \mathbb{A}_M/_\theta$, we can extend this assignment to $v_4 = v_2$ and $v_5 = 0/_\theta$ to obtain a satisfying assignment of $\mathcal{P}_{/W}$. For C_2' the argument is the same.

For 1-element strands consider $\{v_2\}$. Then $Y = \{v_1, v_3, v_4, v_5\}$, and $\mu_1^Y = \mu_3^Y = \mu_4^Y = \mu_5^Y = \theta$. We have $\mathcal{P}_{/\{v_2\}} = (V, \{C_1'', C_2''\})$, where $C_1'' = \langle \mathbf{s}_1, R_1^{\theta\theta} \rangle$, $C_2'' = \langle \mathbf{s}_2, R_2^{\theta\theta} \rangle$, and

$$R_1^{\theta\theta} = \begin{pmatrix} 0/_\theta & 0/_\theta & 2/_\theta & 2/_\theta \\ 0 & 1 & 2 & 2 \\ 0/_\theta & 0/_\theta & 0/_\theta & 2/_\theta \end{pmatrix}, \qquad R_2^{\theta\theta} = \begin{pmatrix} 0 & 1 & 2 & 2 \\ 0/_\theta & 0/_\theta & 2/_\theta & 2/_\theta \\ 0/_\theta & 0/_\theta & 0/_\theta & 2/_\theta \end{pmatrix}.$$

As is easily seen, any assignment to v_1, v_2, v_3 or to v_2, v_4, v_5 can be extended to a solution of $\mathcal{P}_{/\{v_2\}}$. ◇

For an instance $\mathcal{P}$ we say that an instance $\mathcal{P}'$ is *strictly smaller* than instance $\mathcal{P}$ if $\mathsf{size}(\mathcal{P}') < \mathsf{size}(\mathcal{P})$.

LEMMA 5.5. *Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal instance. Then $\mathcal{P}$ can be transformed to an equivalent block-minimal instance $\mathcal{P}'$ by solving a quadratic number of strictly smaller CSPs.*

Proof. To establish the block-minimality of $\mathcal{P}$, for every strand $U \subseteq V$, we need to check if the problem given in condition (BM) is minimal. If they are then $\mathcal{P}$ is block-minimal, otherwise some tuples can be removed from some constraint relation R (the set of tuples that remain in R is always a subalgebra, as is easily seen), and the instance $\mathcal{P}$ tightened, in which case we need to repeat the procedure with the tightened instance. Therefore we just need to show how to reduce solving those subproblems to solving strictly smaller CSPs.

By the definition of a strand, there is a partition $B_{w1}, \dots, B_{w\ell}$ of $\mathbb{A}_w$ for $w \in U$ such that for every constraint $\langle \mathbf{s}, R\rangle \in \mathcal{C}$, for any $w_1, w_2 \in \mathbf{s} \cap U$, any $\mathbf{b} \in R$, and any $i \in [\ell]$ it holds that $\mathbf{b}[w_1] \in B_{w_1 i}$ if and only if $\mathbf{b}[w_2] \in B_{w_2 i}$. Then the problem $\mathcal{P}_{/U}$ is a disjoint union of instances $\mathcal{P}_1, \dots, \mathcal{P}_\ell$ given by: $\mathcal{P}_i = (V, \mathcal{C}_i)$, where for every constraint $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ there is $C_i = \langle \mathbf{s}, R_i\rangle \in \mathcal{C}_i$ such that

$$R_i = \{\mathbf{a}' \mid \mathbf{a} \in R, \mathbf{a}[w] \in B_{wi} \text{ for each } w \in \mathbf{s} \cap U\},$$

with $\mathbf{a}'[u] = \mathbf{a}[u]/_{\mu_u^Y}$, $Y = \mathsf{MAX}(\mathcal{P}) - U$, for each $u \in \mathbf{s}$. Clearly, $\mathsf{size}(\mathcal{P}_i) < \mathsf{size}(\mathcal{P})$ for each $i \in [\ell]$.

In order to establish the minimality of $\mathcal{P}_{/U}$ it suffices to do the following. Take $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and $\mathbf{a} \in R$. We need to check that $\mathbf{a}' = \mathbf{a}/_{\overline{\mu}^Y}$, $Y = \mathsf{MAX}(\mathcal{P}) - U$, extends to a solution of at least one of the problems $\mathcal{P}_1, \dots, \mathcal{P}_\ell$. For $i \in [\ell]$ let $\mathcal{P}'_i$ be the problem obtained from $\mathcal{P}_i$ as follows: fix the values of variables from $\mathbf{s}$ to those of $\mathbf{a}'$, or in other words, add the constraint $\langle (w), \{\mathbf{a}[w]/_{\mu_w^Y}\}\rangle$ for each $w \in \mathbf{s}$. Then $\mathbf{a}'$ can be extended to a solution of $\mathcal{P}_i$ if and only if $\mathcal{P}'_i$ has a solution. □

5.4 The algorithm

We are now in a position to describe our solution algorithm. In the algorithm we distinguish three cases depending on the presence of semilattice edges and the centralizers of the domains of variables. In each case we employ different methods of solving or reducing the instance to a strictly smaller one. Algorithm 1, SolveCSP, gives a more formal description of the solution algorithm.

Let $\mathcal{P} = (V, \mathcal{C})$ be a subdirectly irreducible (2,3)-minimal instance. Let $\mathsf{Center}(\mathcal{P})$ denote the set of variables $v \in V$ such that $(\underline{0}_v : \mu_v) = \underline{1}_v$. Let $\mu_v^* = \mu_v$ if $v \in \mathsf{MAX}(\mathcal{P}) \cap \mathsf{Center}(\mathcal{P})$ and $\mu_v^* = \underline{0}_v$ otherwise.

Semilattice free domains. If all domains of $\mathcal{P}$ are semilattice free then $\mathcal{P}$ can be solved in polynomial time, using the few subpowers algorithm, as shown in [21, 42].

Small centralizers. If $\mu_v^* = \underline{0}_v$ for all $v \in V$, by Theorem 5.6 block-minimality guarantees that a solution exists, and we can use Lemma 5.5 to solve the instance.

Theorem 5.6. *If $\mathcal{P}$ is subdirectly irreducible, (2,3)-minimal, block-minimal, and $\mathsf{MAX}(\mathcal{P}) \cap \mathsf{Center}(\mathcal{P}) = \varnothing$, then $\mathcal{P}$ has a solution.*

Large centralizers. Suppose that $\mathsf{MAX}(\mathcal{P}) \cap \mathsf{Center}(\mathcal{P}) \neq \varnothing$. In this case the algorithm proceeds in three steps.

Stage 1. Consider the problem $\mathcal{P}/_{\overline{\mu}^*}$. We establish the global 1-minimality of this problem. If it is tightened in the process, we start solving the new problem from scratch. To check global 1-minimality, for each $v \in V$ and every $a \in \mathbb{A}_v/_{\mu_v^*}$, we need to find a solution of the instance, or show it does not exist. To this end, add the constraint $\langle (v), \{a\}\rangle$ to $\mathcal{P}/_{\overline{\mu}^*}$. The resulting problem belongs to $\mathrm{CSP}(\mathbb{A})$, since $\mathbb{A}_v$ is idempotent, and hence $\{a\}$ is a subalgebra of $\mathbb{A}_v/_{\mu_v^*}$. Then we establish (2,3)-minimality and block minimality of the resulting problem. Let us denote it $\mathcal{P}'$. There are two possibilities. First, if $\mathsf{size}(\mathcal{P}') < \mathsf{size}(\mathcal{P})$ then $\mathcal{P}'$ is a problem strictly smaller than $\mathcal{P}$ and can be solved by recursively calling Algorithm 1 on $\mathcal{P}'$. If $\mathsf{size}(\mathcal{P}') = \mathsf{size}(\mathcal{P})$ then, as all the domains $\mathbb{A}_v$ of maximal size for $v \in \mathsf{Center}(\mathcal{P})$ are replaced with their quotient algebras, there is $w \notin \mathsf{Center}(\mathcal{P})$ such that $|\mathbb{A}_w| = \mathsf{size}(\mathcal{P})$ and $\mathbb{A}_w$ is not semilattice free. Therefore for every $u \in \mathsf{Center}(\mathcal{P}')$, for the corresponding domain $\mathbb{A}'_u$ we have $|\mathbb{A}'_u| < \mathsf{size}(\mathcal{P}) = \mathsf{size}(\mathcal{P}')$. Thus, $\mathsf{MAX}(\mathcal{P}') \cap \mathsf{Center}(\mathcal{P}') = \varnothing$, and $\mathcal{P}'$ has a solution by Theorem 5.6.

Stage 2. For every $v \in \mathsf{MAX}(\mathcal{P})$ we find a solution φ of $\mathcal{P}/_{\overline{\mu}^*}$ such that there is $a \in \mathbb{A}_v$ such that $(a, \varphi(v))$ is a semilattice edge if $\mu^*_v = \underline{0}_v$, or, if $\mu^*_v = \mu_v$, there is $b \in \varphi(v)$ such that (a, b) is a semilattice edge. Take $v \in \mathsf{MAX}(\mathcal{P})$ and $b \in \mathbb{A}_v/_{\mu^*_v}$ such that (a, b) is a semilattice edge in $\mathbb{A}_v/_{\mu^*_v}$ for some $a \in \mathbb{A}_v/_{\mu^*_v}$. Such a semilattice edge exists, because $\mathbb{A}_v$ is not semilattice free. Also, if $\mu^*_v \neq \underline{0}_v$, then $v \in \mathsf{Center}(\mathcal{P})$ and $(\underline{0}_v : \mu_v) = \underline{1}_v$ and by Corollary 4.4 its semilattice edges are all between μ_v-blocks. Since $\mathcal{P}/_{\overline{\mu}^*}$ is globally 1-minimal, there is a solution $\varphi_{v,b}$ such that $\varphi_{v,b}(v) = b$, and therefore $\varphi_{v,b}$ satisfies the condition. Let $\mathsf{MAX}(\mathcal{P}) = \{v_1, \dots, v_\ell\}$ and $b_1, \dots, b_\ell$ the values satisfying the requirements above.

Stage 3. We apply the transformation of $\mathcal{P}$ suggested by Maróti in [54]. For a solution φ of $\mathcal{P}/_{\overline{\mu}^*}$, by $\mathcal{P} \cdot \varphi$ we denote the instance (V, C_φ) given by the rule: for every $C = \langle \mathbf{s}, R \rangle \in C$ the set C_φ contains a constraint $\langle \mathbf{s}, R \cdot \varphi \rangle$. To construct $R \cdot \varphi$ choose a tuple $\mathbf{b} \in R$ such that $\mathbf{b}[v]/_{\mu^*_v} = \varphi(v)$ for all $v \in \mathbf{s}$; this is possible because φ is a solution of $\mathcal{P}/_{\overline{\mu}^*}$. Then set $R \cdot \varphi = \{\mathbf{ab} \mid \mathbf{a} \in R\}$, where $\mathbf{ab}$ denotes coordinate-wise multiplication as introduced in Lemma 3.3. Note that by Lemma 4.2 even if $\mu^*_v \neq \underline{0}_v$ the tuple $\mathbf{ab}$ does not depend on the choice of $\mathbf{b}$ as long as $\mathbf{b}[v]/_{\mu^*_v} = \varphi(v)$ for all $v \in \mathbf{s}$. By the results of [54] and Lemma 4.2 the instance $\mathcal{P} \cdot \varphi$ has a solution if and only if $\mathcal{P}$ does. We now use the solutions $\varphi_{v_1,b_1}, \dots, \varphi_{v_\ell,b_\ell}$ to construct a new problem

$$\mathcal{P}^1 = (\dots((\mathcal{P} \cdot \varphi_{v_1,b_1}) \cdot \varphi_{v_2,b_2}) \cdot \dots) \cdot \varphi_{v_\ell,b_\ell}.$$

Note that the transformation of $\mathcal{P}$ above boils down to a collection of mappings $p_v : \mathbb{A}_v \to \mathbb{A}_v$, $v \in V$, so called *consistent mappings*, see Section 5.5, that also satisfy some additional properties. If we now repeat the procedure above starting from $\mathcal{P}^1$ and using the same solutions φ_{v_i,b_i}, we obtain an instance $\mathcal{P}^2$, for which the corresponding collection of consistent mappings is $p_v \circ p_v$, $v \in V$. More generally,

$$\mathcal{P}^{i+1} = (\dots((\mathcal{P}^i \cdot \varphi_{v_1,b_1}) \cdot \varphi_{v_2,b_2}) \cdot \dots) \cdot \varphi_{v_\ell,b_\ell}.$$

There is k such that p_v^k is idempotent for every $v \in V$, that is, $(p_v^k \circ p_v^k)(x) = p_v^k(x)$ for all x (see, e.g., [59, Chapter 5]). Set $\mathcal{P}^\dagger = \mathcal{P}^k$. We will show later that $\mathsf{size}(\mathcal{P}^\dagger) < \mathsf{size}(\mathcal{P})$.

This last case can be summarized as the following

THEOREM 5.7. *If $\mathcal{P}/_{\overline{\mu}^*}$ is globally 1-minimal, then $\mathcal{P}$ can be reduced in polynomial time to a strictly smaller instance over a class of algebras satisfying the conditions of the Dichotomy Conjecture.*

We now illustrate the algorithm on our running example.

Example 5.8. We illustrate the algorithm SolveCSP on the instance from Example 5.2. Recall that the domain of each variable is $\mathbb{A}_M$, its monolith is θ, and $(\underline{0} : \theta)$ is the full relation. This means that $\mathsf{size}(\mathcal{P}) = 3$, $\mathsf{MAX}(\mathcal{P}) = V$ and $\mathsf{Center}(\mathcal{P}) = V$, as well. Therefore we are in the case of large centralizers. Set $\mu^*_{v_i} = \theta_i$ for each $i \in [5]$ and consider the problem $\mathcal{P}/_{\overline{\mu}*} = (V, \{C^*_1 = \langle \mathbf{s}_1, R^*_1 \rangle, C^*_2 = \langle \mathbf{s}_2, R^*_2 \rangle)$, where

$$R^*_1 = R^*_2 = \begin{pmatrix} 0/_\theta & 2/_\theta & 2/_\theta \\ 0/_\theta & 2/_\theta & 2/_\theta \\ 0/_\theta & 0/_\theta & 2/_\theta \end{pmatrix}.$$

It is an easy exercise to show that this instance is globally 1-minimal (every value $0/_\theta$ can be extended to the all-$0/_\theta$ solution, and every value $2/_\theta$ can be extended to the all-$2/_\theta$ solution). This completes *Stage 1*. For every variable v_i we

Algorithm 1 Procedure SolveCSP

Require: A CSP instance $\mathcal{P} = (V, C)$ over $\mathcal{A}$
Ensure: A solution of $\mathcal{P}$ if one exists, 'NO' otherwise
1: **if** all the domains are semilattice free **then**
2: Solve $\mathcal{P}$ using the few subpowers algorithm and RETURN the answer
3: **end if**
4: Transform $\mathcal{P}$ to a subdirectly irreducible, block-minimal and (2,3)-minimal instance
5: $\mu^*_v = \mu_v$ for $v \in \mathsf{MAX}(\mathcal{P}) \cap \mathsf{Center}(\mathcal{P})$ and $\mu^*_v = \underline{0}_v$ otherwise
6: $\mathcal{P}^* = \mathcal{P}/_{\overline{\mu}^*}$
7: /* establish the global 1-minimality of $\mathcal{P}^*$
8: **for** every $v \in V$ and $a \in \mathbb{A}_v/_{\mu^*_v}$ **do**
9: $\mathcal{P}' = \mathcal{P}^* \cup \{\langle v, a\rangle\}$ /* Add constraint $\langle (v), \{a\}\rangle$ fixing the value of v to a
10: Transform $\mathcal{P}'$ to a subdirectly irreducible, (2,3)-minimal instance $\mathcal{P}''$
11: If $\mathsf{size}(\mathcal{P}'') < \mathsf{size}(\mathcal{P})$ call SolveCSP on $\mathcal{P}''$ and flag a if $\mathcal{P}''$ has no solution
12: Establish the block-minimality of $\mathcal{P}''$; if the problem changes, return to Step 10
13: If the resulting instance is empty, flag the element a
14: **end for**
15: If there are flagged values, tighten the instance by removing the flagged elements and return to Step 1
16: Use Theorem 5.7 to reduce $\mathcal{P}$ to an instance $\mathcal{P}^\dagger$ with $\mathsf{size}(\mathcal{P}^\dagger) < \mathsf{size}(\mathcal{P})$
17: Call SolveCSP on $\mathcal{P}^\dagger$ and RETURN the answer

choose $b \in \mathbb{A}_M/_\theta$ such that for some $a \in \mathbb{A}_M/_\theta$ the pair (a, b) is a semilattice edge. Since $\mathbb{A}_M/_\theta$ is a 2-element semilattice, setting $b = 0/_\theta$ and $a = 2/_\theta$ is the only choice. Therefore φ_{v_i,b_i} in our case can be chosen to be the same solution φ given by $\varphi(v_i) = 0/_\theta$; and *Stage 2* is completed. For *Stage 3* first note that in $\mathbb{A}_M$ the operation r plays the role of multiplication $\cdot$. Then for each of the constraints C^1, C^2 choose a representative $\mathbf{a}_1 \in R_1 \cap (\varphi(v_1) \times \varphi(v_2) \times \varphi(v_3)) = R_1 \cap \{0,1\}^3$, $\mathbf{a}_2 \in R_2 \cap (\varphi(v_2) \times \varphi(v_4) \times \varphi(v_5)) = R_2 \cap \{0,1\}^3$, and set $\mathcal{P}' = (\{v_1, \dots, v_5\}, \{C'_1 = \langle (v_1, v_2, v_3), R'_1\rangle, C'_2 = \langle (v_2, v_4, v_5), R'_2\rangle\})$, where $R'_1 = r(R_1, \mathbf{a}_1)$, $R'_2 = r(R_2, \mathbf{a}_2)$. Since $r(2,0) = r(2,1) = 0$, regardless of the choice of $\mathbf{a}_1, \mathbf{a}_2$ in our case $R'_1 \subseteq R_1, R'_2 \subseteq R_2$, and are invariant with respect to the affine operation of $\mathbb{Z}_2$. Therefore the instance $\mathcal{P}'$ can be viewed as a system of linear equations over $\mathbb{Z}_2$ (this system is actually empty in our case), and can be easily solved. ⋄

Using Lemma 5.5 and Theorems 5.6,5.7 it is not difficult to see that the algorithm runs in polynomial time.

THEOREM 5.9. *Algorithm SolveCSP (Algorithm 1) correctly solves every instance from* $\mathrm{CSP}(\mathcal{A})$ *and runs in polynomial time.*

PROOF. By the results of [21, 24] the algorithm correctly solves the given instance $\mathcal{P}$ in polynomial time if the conditions of Step 1 are true. Lemma 5.5 implies that Steps 4 and 12 can be completed by recursing to strictly smaller instances.

Next we show that the for-loop in Steps 8-14 checks if $\mathcal{P}^* = \mathcal{P}/_{\overline{\mu}^*}$ is globally 1-minimal. For this we need to verify that a value a is flagged if and only if $\mathcal{P}^*$ has no solution φ with $\varphi(v) = a$, and therefore if no values are flagged then $\mathcal{P}^*$ is globally 1-minimal. If $\varphi(v) = a$ for some solution φ of $\mathcal{P}^*$, then φ is a solution to the instance $\mathcal{P}'$ constructed in Step 9. In this case Steps 11,12 cannot result in an empty instance. Suppose $a \in \mathbb{A}_v/_{\mu^*_v}$ is not flagged. If $\mathsf{size}(\mathcal{P}'') < \mathsf{size}(\mathcal{P})$ this means that $\mathcal{P}''$ and therefore $\mathcal{P}'$ has a solution. Otherwise this means that establishing block-minimality of $\mathcal{P}''$ is successful. In this case $\mathcal{P}''$ has a solution by Theorem 5.6, because $\mathsf{MAX}(\mathcal{P}'') \cap \mathsf{Center}(\mathcal{P}'') = \varnothing$. This in turn implies that $\mathcal{P}'$ has a solution. Observe also that the set of unflagged values for each variable $v \in V$ is a subalgebra of $\mathbb{A}/_{\mu^*}$.

Indeed, the set of solutions of $\mathcal{P}^*$ is a subalgebra $\mathcal{S}^*$ of $\prod_{v\in V} \mathbb{A}/_{\mu^*}$, and the set of unflagged values is the projection of $\mathcal{S}^*$ on the coordinate position v.

Finally, if Steps 8–15 are completed without restarts, Steps 16,17 can be completed by Theorem 5.7, and recursing on $\mathcal{P}'$ such that either $\mathsf{size}(\mathcal{P}') < \mathsf{size}(\mathcal{P})$ or $\mathsf{MAX}(\mathcal{P}') \cap \mathsf{Center}(\mathcal{P}') = \varnothing$.

To see that the algorithm runs in polynomial time it suffices to observe that

(1) The number of restarts in Steps 4 and 15 is at most linear, as the instance becomes smaller after every restart; therefore the number of times Steps 4–15 are executed together is at most linear.
(2) The number of iterations of the for-loop in Steps 8–14 is linear.
(3) The number of restarts in Steps 10 and 12 is at most linear, as the instance becomes smaller after every iteration.
(4) Every call of SolveCSP when establishing block-minimality in Steps 4, and 12 is made on an instance strictly smaller than $\mathcal{P}$, and therefore the depth of recursion is bounded by $\mathsf{size}(\mathcal{P})$ in Step 4,11,12 and 17.

Thus a more thorough estimation gives a bound on the running time of $O(n^{3k})$, where k is the maximal size of an algebra in $\mathcal{A}$. □

5.5 Proof of Theorem 5.7

Following [54] let $\mathcal{P} = (V, \mathcal{C})$ be an instance and $p_v\colon \mathbb{A}_v \to \mathbb{A}_v$, $v \in V$. Mappings p_v, $v \in V$, are said to be *consistent* if for any $\langle \mathbf{s}, R\rangle \in \mathcal{C}$, $\mathbf{s} = (v_1, \dots, v_k)$, and any tuple $\mathbf{a} \in R$ the tuple $(p_{v_1}(\mathbf{a}[1]), \dots, p_{v_k}(\mathbf{a}[k]))$ belongs to R. It is easy to see that the composition of two families of consistent mappings is also a consistent mapping. For consistent idempotent mappings p_v by $p(\mathcal{P})$ we denote the *retraction* of $\mathcal{P}$, that is, $\mathcal{P}$ restricted to the images of p_v. It was proved in [54, Lemma 3] that in this case $\mathcal{P}$ has a solution if and only if $p(\mathcal{P})$ has. For the sake of completeness we reproduce here this short proof. Firstly, for any constraint $\langle \mathbf{a}, R\rangle$, $\mathbf{s} = (v_1, \dots, v_k)$, of $\mathcal{P}$ the corresponding constraint relation $p(R) = \{(p_{v_1}(\mathbf{a}[1]), \dots, p_{v_k}(\mathbf{a}[k])) \mid \mathbf{a} \in R\}$ is a subset of R, which means that every solution of $p(\mathcal{P})$ is also a solution of $\mathcal{P}$. On the other hand, if φ is a solution of $\mathcal{P}$ then for every constraint $\langle \mathbf{a}, R\rangle$, $\mathbf{s} = (v_1, \dots, v_k)$, of $\mathcal{P}$ it hold that $(\varphi(v_1), \dots, \varphi(v_k)) \in R$. Let $p(\varphi)$ denote the mapping given by $p(\varphi)(v) = p_v(\varphi(v))$ for every $v \in V$. Then $(p_{v_1}(\varphi(v_1)), \dots, p_{v_k}(\varphi(v_k))) \in p(R)$ implying that $p(\varphi)$ is a solution of $p(\mathcal{P})$.

Let φ be a solution of $\mathcal{P}/_{\overline{\mu}^*}$. We define $p_v^\varphi : \mathbb{A}_v \to \mathbb{A}_v$ as follows: $p_v^\varphi = q_v^k$, where $q_v(a) = a \cdot b_v$, element b_v is any element of $\varphi(v)$, and k is such that q_v^k is idempotent for all $v \in V$. Note that by Lemma 4.2 this mapping is properly defined even if $\mu_v^* \neq \underline{0}_v$.

LEMMA 5.10. *Mappings p_v^φ, $v \in V$, are consistent.*

PROOF. Take any $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$. Since φ is a solution of $\mathcal{P}/_{\overline{\mu}^*}$, there is $\mathbf{b} \in R$ such that $\mathbf{b}[v] \in \varphi(v)$ for $v \in \mathbf{s}$. Then for any $\mathbf{a} \in R$, $q(\mathbf{a}) = \mathbf{a} \cdot \mathbf{b} \in R$, and this product does not depend on the choice of $\mathbf{b}$, as follows from Lemma 4.2. Iterating this operation also produces a tuple from R. □

We are now in a position to prove Theorem 5.7.

PROOF OF THEOREM 5.7. We need to show 3 properties of the problem $\mathcal{P}^\dagger$ constructed in Stage 3: (a) $\mathcal{P}$ has a solution if and only if $\mathcal{P}^\dagger$ does; (b) for every $v \in \mathsf{MAX}(\mathcal{P})$, $|\mathbb{A}_v^\dagger| < |\mathbb{A}_v|$, where $\mathbb{A}_v^\dagger$ is the domain of v in $\mathcal{P}^\dagger$; and (c) every algebra $\mathbb{A}_v^\dagger$ has a weak near-unanimity term operation. We use the inductive definition of $\mathcal{P}^\dagger$ given in Stage 3.

Recall that $\mathsf{MAX}(\mathcal{P}) = \{v_1, \dots, v_\ell\}$, $a_i, b_i \in \mathbb{A}_{v_i}$ are such that (a_i, b_i) is a semilattice edge and $b_i \in \varphi_{v_i,b_i}(v_i)$, where φ_{v_i,b_i} is a solution of $\mathcal{P}/_{\overline{\mu}^*}$. For $v \in V$ let mapping $p_{vi} : \mathbb{A}_v \to \mathbb{A}_v$ be given by

$$p_{vi}(x) = (\dots (x \cdot \varphi_{v_1,b_1}(v)) \cdot \dots) \cdot \varphi_{v_i,b_i}(v),$$

where if $\mu^*_v = \mu_v$ by Lemma 4.2 the multiplication by $\varphi_{v_j,b_j}(v)$ does not depend on the choice of a representative from $\varphi_{v_j,b_j}(v)$. By Lemma 5.10 $\{p_{vi}\}$ for every i, and so $\{p_v\}$ and $\{p^k_v\}$ are collections of consistent mappings. Now (a) follows from [54].

Next we show that for every $j \le i \le \ell$ it holds that $|p_{vi}(\mathbb{A}_{v_j})| < |\mathbb{A}_{v_j}|$. Since applying mappings to a set does not increase its cardinality, this implies (b). If $|p_{vj-1}(\mathbb{A}_{v_j})| < |\mathbb{A}_{v_j}|$, we have the desired inequality applying the observation in the previous sentence. Otherwise $a_j \in \mathbb{A}_{v_j} = p_{vj-1}(\mathbb{A}_{v_j})$, and it suffices to notice that $a_j \cdot \varphi_{v_j,b_j}(v_j) = b_j \cdot \varphi_{v_j,b_j}(v_j) = b_j$.

To prove (c) observe that if $\mathbb{A}_v$ is semilattice free then p^φ_v is the identity mapping for any φ by Lemma 3.3, and so $\mathbb{A}^\dagger_v = \mathbb{A}_v$. For the remaining domains let f be a weak near-unanimity term of the class $\mathcal{A}$. Then for any idempotent mapping p the operation $p \circ f$ given by $(p \circ f)(x_1, \dots, x_n) = p(f(x_1, \dots, x_n))$ is a weak near-unanimity term of $p(\mathcal{A}) = \{p(\mathbb{A}) \mid \mathbb{A} \in \mathcal{A}\}$. The result follows. □

6 ALGEBRA TECHNICALITIES

The rest of the paper is dedicated to proving Theorem 5.6. This part assumes some familiarity with algebraic terminology. A brief review of the necessary facts from universal algebra can be found in [25]. In this section we recall some results from [25] necessary for our proof.

6.1 Coloured graphs

In [16, 29] we introduced a local approach to the structure of finite algebras. All algebras in this and subsequent sections are supposed to be finite and idempotent. As we use this approach in the proof of Theorem 5.6, we present the necessary elements of it here, see also [20, 23]. For the sake of the definitions below we slightly abuse terminology and by a module mean the full idempotent reduct of a module.

For an algebra $\mathbb{A}$ the graph $\mathcal{G}(\mathbb{A})$ is defined as follows. The vertex set is the universe A of $\mathbb{A}$. A pair ab of vertices is an *edge*[5] if and only if there exists a congruence θ of $\mathrm{Sg}(a, b)$, and a term operation f of $\mathbb{A}$ such that either $\mathrm{Sg}(a, b)/_\theta$ is a module and f is an affine operation on it, or f is a semilattice operation on $\{a/_\theta, b/_\theta\}$, or f is a majority operation on $\{a/_\theta, b/_\theta\}$. (Note that we use the same operation symbol in this case.) If there are a congruence θ and a term operation f of $\mathbb{A}$ such that f is a semilattice operation on $\{a/_\theta, b/_\theta\}$ then ab is said to have the *semilattice type*. An edge ab is of *majority type* if there are a congruence θ and a term operation f such that f is a majority operation on $\{a/_\theta, b/_\theta\}$ and there is no semilattice term operation on $\{a/_\theta, b/_\theta\}$. Finally, ab has the *affine type* if there are θ and f such that f is an affine operation on $\mathrm{Sg}(a, b)/_\theta$ and $\mathrm{Sg}(a, b)/_\theta$ is a module. Pairs of the form $\{a/_\theta, b/_\theta\}$ will be referred to as *thick edges*.

Properties of $\mathcal{G}(\mathbb{A})$ are related to the properties of the algebra $\mathbb{A}$.

THEOREM 6.1 (THEOREM 3.5 OF [20]). *Let $\mathbb{A}$ be an algebra $\mathbb{A}$ such that* $\mathsf{var}(\mathbb{A})$ *omits type* **1**. *Then*

(1) *any two elements of $\mathbb{A}$ are connected by a sequence of edges of the semilattice, majority, and affine types;*
(2) $\mathsf{var}(\mathbb{A})$ *omits types* **1** *and* **2** *if and only if $\mathcal{G}(\mathbb{A})$ satisfies the conditions of item (1) and contains no edges of the affine type.*

[5]The notion of edge we introduce here is more general than semilattice edges used before. We will show how semilattice edges fit into the picture shortly.

Algebra $\mathbb{A}$ is said to be *smooth* if for every edge ab of the semilattice or majority type, $a/_\theta \cup b/_\theta$, where θ is a congruence witnessing that ab is an edge, is a subalgebra of $\mathbb{A}$.

THEOREM 6.2 (THEOREM 3.12 OF [20]). *For any algebra $\mathbb{A}$ such that $\mathsf{var}(\mathbb{A})$ omits type* **1** *there is a reduct $\mathbb{A}'$ of $\mathbb{A}$ that is smooth and such that $\mathsf{var}(\mathbb{A}')$ omits type* **1**.

Moreover, if $\mathbb{A}$ does not contain edges of the affine types, $\mathbb{A}'$ can be chosen such that it does not contain edges of the affine type.

For the rest of the paper we fix a finite class $\mathcal{K}$ of smooth idempotent algebras closed under taking subalgebras and homomorphic images and such that $\mathsf{var}(\mathcal{K})$ omits type **1**, and let $\mathcal{V}$ be the class of finite algebras from the variety generated by $\mathcal{K}$, that is, the pseudovariety generated by $\mathcal{K}$. We will slightly abuse the terminology and call $\mathcal{V}$ the variety generated by $\mathcal{K}$. If we are interested in a particular algebra $\mathbb{A}$, set $\mathcal{K} = \mathsf{HS}(\mathbb{A})$.

We use the following refinement of the construction of edges. A ternary term operation g' of $\mathcal{K}$ is said to satisfy the *majority condition* for $\mathcal{K}$ if g' is a majority operation on every thick majority edge of every algebra from $\mathcal{K}$ and satisfies the equation $g(x, g(x, y, y), g(x, y, y)) = g(x, y, y)$. A ternary term operation h' is said to satisfy the *minority condition* for $\mathcal{K}$ if h' is a Mal'tsev operation on every thick affine edge of every algebra of $\mathcal{K}$ and satisfies the equation $h(h(x, y, y), y, y) = h(x, y, y)$. Operations satisfying the majority and minority conditions always exist, as is proved in [20, Theorem 3.21, Corollary 3.22]. Fix an operation h satisfying the minority condition. A pair of elements $a, b \in \mathbb{A} \in \mathcal{V}$ is said to be

(1) a *semilattice edge* if there is a term operation f such that $f(a, b) = f(b, a) = b$;
(2) a *thin majority edge* if for any term operation g' satisfying the majority condition the subalgebras $\mathsf{Sg}(a, g'(a, b, b))$, $\mathsf{Sg}(a, g'(b, a, b))$, $\mathsf{Sg}(a, g'(b, b, a))$ contain b.
(3) a *thin affine edge* if $h(b, a, a) = b$, where h if the fixed operation chosen above and $b \in \mathsf{Sg}(a, h'(a, a, b))$ for any term operation h' satisfying the minority condition.

Note that thin edges are directed, as a and b appear asymmetrically. By $\mathcal{G}'(\mathbb{A})$ we denote the digraph whose vertices are the elements of $\mathbb{A}$, and the edges are the thin edges defined above. Theorem 3.21 and Corollary 3.22 from [20] also imply that there exists a binary term operation $\cdot$ of $\mathbb{A}$ that is a semilattice operation on every thin semilattice edge.

LEMMA 6.3 (LEMMA 1.11 OF [23], LEMMA 12 OF [25]). (1) *Let $\mathbb{A} \in \mathcal{V}$, $\overline{a}\overline{b}$ be a thin edge in $\mathbb{A}/_\theta$, $\theta \in \mathsf{Con}(\mathbb{A})$, and $a \in \overline{a}$, then there is $b \in \overline{b}$ such that ab is a thin edge in $\mathbb{A}$ of the same type.*
(2) *Let $\mathbb{A} \in \mathcal{V}$ and ab be a thin edge. Then ab is a thin edge of the same type in any subalgebra of $\mathbb{A}$ containing a, b, and $a/_\theta b/_\theta$ is a thin edge in $\mathbb{A}/_\theta$ of the same type for any congruence θ.*
(3) *Let $\mathbb{A} \in \mathcal{V}$ and $\mathbb{B}$ its subalgebra. Then every thin edge of $\mathbb{B}$ is a thin edge of $\mathbb{A}$ of the same type.*

We distinguish several types of paths in $\mathcal{G}'(\mathbb{A})$ depending on the types of edges involved. A directed path in $\mathcal{G}'(\mathbb{A})$ is called an *asm-path*, if there is an asm-path from a to b we write $a \sqsubseteq^{asm} b$. If all edges of this path are semilattice or affine [only semilattice], it is called an *affine-semilattice path* [*semilattice path*, respectively] or an *as-path* [*s-path*, respectively]. If there is an as-path [s-path] from a to b we write $a \sqsubseteq^{as} b$ [respectively, $a \sqsubseteq b$]. We consider strongly connected components of $\mathcal{G}'(\mathbb{A})$ with majority edges removed, and the natural partial order on such components. The maximal components will be called *as-components*, and the elements from as-components are called *as-maximal*; the set of all as-maximal elements of $\mathbb{A}$ is denoted by $\mathsf{amax}(\mathbb{A})$. For $a \in \mathsf{amax}(\mathbb{A})$ by $\mathsf{as}(a)$ we denote the as-component containing a. Then *maximal components* and *maximal* elements are defined in the same way by removing from $\mathcal{G}'(\mathbb{A})$ all majority and affine edges. The set of all maximal elements of $\mathbb{A}$ is denoted by $\max(\mathbb{A})$. An alternative way to define as-maximal [maximal] elements is as follows: a is as-maximal [respectively, maximal] if for every $b \in \mathbb{A}$ such that $a \sqsubseteq^{as} b$ [respectively, $a \sqsubseteq b$] it also holds that $b \sqsubseteq^{as} a$ [respectively, $b \sqsubseteq a$]. Finally, element $a \in \mathbb{A}$ is said to be *universally maximal* (or *u-maximal* for short) if for every $b \in \mathbb{A}$ such that $a \sqsubseteq^{asm} b$ it also holds that $b \sqsubseteq^{asm} a$. The set of all u-maximal elements of $\mathbb{A}$ is denoted $\mathsf{umax}(\mathbb{A})$.

U-maximality has additional useful properties.

LEMMA 6.4 (THEOREM 2.3, COROLLARY 2.4 OF [23]; LEMMA 17, [25]). (1) *Any two u-maximal elements are connected with an asm-path.*
(2) *Every maximal and as-maximal element is also u-maximal.*
(3) *Let $\mathbb{B}$ be a subalgebra of $\mathbb{A}$ containing a u-maximal element of $\mathbb{A}$. Then every element u-maximal in $\mathbb{B}$ is also u-maximal in $\mathbb{A}$. In particular, if α is a congruence of $\mathbb{A}$ and $\mathbb{B}$ is a u-maximal α-block, that is, $\mathbb{B}$ is a u-maximal element in $\mathbb{A}/_\alpha$, then* $\mathsf{umax}(\mathbb{B}) \subseteq \mathsf{umax}(\mathbb{A})$.

The following statement shows how thin edges and the connectivity notions are related to subalgebras of direct products.

LEMMA 6.5 (THE MAXIMALITY LEMMA, COROLLARIES 1.13,1.14 OF [23], LEMMA 18 OF [25]). *Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$, $I \subseteq [n]$.*

(1) *For any $\mathbf{a} \in R$, and an as-path (asm-path, s-path) $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathrm{pr}_I R$ with $\mathrm{pr}_I \mathbf{a} = \mathbf{b}_1$, there is an as-path (asm-path, s-path) $\mathbf{b}'_1, \dots, \mathbf{b}'_k \in R$ such that $\mathbf{b}'_1 = \mathbf{a}$ and $\mathrm{pr}_I \mathbf{b}'_i = \mathbf{b}_i$, $i \in [k]$.*
(2) *For any $\mathbf{b} \in \mathsf{amax}(\mathrm{pr}_I R)$ ($\mathbf{b} \in \mathsf{umax}(\mathrm{pr}_I R)$, $\mathbf{b} \in \max(\mathrm{pr}_I R)$) there is $\mathbf{b}' \in \mathsf{amax}(R)$ ($\mathbf{b}' \in \mathsf{umax}(R)$, $\mathbf{b}' \in \max(\mathrm{pr}_I R)$), such that $\mathrm{pr}_I \mathbf{b}' = \mathbf{b}$.*
(3) *If $\mathbf{a} \in R$ is an as-maximal, u-maximal, or maximal element then so is $\mathrm{pr}_I \mathbf{a}$.*

We complete this section with an auxiliary statement that will be needed later. Let $\mathbb{A}$ be a finite algebra and $\alpha, \beta \in \mathrm{Con}(\mathbb{A})$. The pair α, β is said to be a *prime interval*, denoted $\alpha \prec \beta$ if $\alpha < \beta$ and for any $\gamma \in \mathrm{Con}(\mathbb{A})$ with $\alpha \le \gamma \le \beta$ either $\alpha = \gamma$ or $\beta = \gamma$. By $\mathsf{typ}(\alpha, \beta)$ we denote the type of (α, β) in the sense of Tame Congruence Theorem, [41].

LEMMA 6.6 (LEMMA 22 OF [25], LEMMA 4.14 OF [41]). (1) *Let $\mathbb{A} \in \mathcal{V}$, $\alpha \prec \beta$, $\alpha, \beta \in \mathrm{Con}(\mathbb{A})$, let B be a β-block and $\mathsf{typ}(\alpha, \beta) = \mathbf{2}$. Then $B/_\alpha$ is term equivalent to a module. In particular, every pair of elements of $B/_\alpha$ is a thin affine edge in $\mathbb{A}/_\alpha$.*
(2) *If $(\alpha : \beta) \ge \beta$, then $\mathsf{typ}(\alpha, \beta) = \mathbf{2}$.*

6.2 Quasi-decomposition and rectangularity

We make use of the property of quasi-2-decomposability proved in [23].

THEOREM 6.7 (THEOREM 4.1, COROLLARY 4.3 OF [23]). *Let $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$. If R is a subdirect product of $\mathbb{A}_1 \times \dots \times \mathbb{A}_n$ and tuple $\mathbf{a}$ is such that $\mathrm{pr}_J \mathbf{a} \in \mathrm{pr}_J R$ for any $J \subseteq [n]$, $|J| = 2$, then there is a tuple $\mathbf{b} \in R$ with $\mathrm{pr}_J \mathbf{a} \sqsubseteq^{as} \mathrm{pr}_J \mathbf{b}$ in $\mathrm{pr}_J R$ for any $J \subseteq [n]$, $|J| = 2$. Moreover, if for some $I \subseteq [n]$ it holds that $\mathrm{pr}_I \mathbf{a} \in \mathsf{amax}(\mathrm{pr}_I R)$ then $\mathbf{b}$ can be chosen such that $\mathrm{pr}_I \mathbf{b} = \mathrm{pr}_I \mathbf{a}$.*

Another property of relations was also introduced in [23] and is similar to the rectangularity property of relations with a Mal'tsev polymorphism. Let R be a subdirect product of $\mathbb{A}_1, \mathbb{A}_2$. By $\mathsf{lk}_1, \mathsf{lk}_2$ we denote the congruences of $\mathbb{A}_1, \mathbb{A}_2$, respectively, generated by the sets of pairs $\{(a, b) \in \mathbb{A}_1^2 \mid \text{there is } c \in \mathbb{A}_2 \text{ such that } (a, c), (b, c) \in R\}$ and

$\{(a,b) \in \mathbb{A}_2^2 \mid$ there is $c \in \mathbb{A}_1$ such that $(c,a),(c,b) \in R\}$, respectively. Congruences $\mathsf{lk}_1, \mathsf{lk}_2$ are called *link congruences*. Relation R is said to be *linked* if the link congruences are full congruences.

PROPOSITION 6.8 (PROPOSITIONS 3.3,3.5 COROLLARY 3.4 OF [23]). *Let R be a subdirect product of $\mathbb{A}_1$ and $\mathbb{A}_2$ from $\mathcal{V}$ and $\mathsf{lk}_1, \mathsf{lk}_2$ the link congruences.*

(1) *Let B_1, B_2 be as-components of a lk_1-block and a lk_2-block, respectively, such that $R \cap (B_1 \times B_2) \neq \varnothing$. Then $B_1 \times B_2 \subseteq R$.*
In particular, if R is linked and B_1, B_2 are as-components of $\mathbb{A}_1, \mathbb{A}_2$, respectively, such that $R \cap (B_1 \times B_2) \neq \varnothing$, then $B_1 \times B_2 \subseteq R$.
(2) *Let B_1 be an as-component of an lk_1-block and $B_2' = R[B_1]$; let $B_2 = \mathsf{umax}(B_2')$. Then $B_1 \times B_2 \subseteq R$.*

6.3 Separating congruences

Let $\mathbb{A}$ be a finite algebra and $\alpha, \beta \in \mathsf{Con}(\mathbb{A})$, $\alpha \prec \beta$. An (α,β)-*minimal set* is a set minimal with respect to inclusion among the sets of the form $f(\mathbb{A})$, where f is a unary polynomial of $\mathbb{A}$ such that $f(\beta) \not\subseteq \alpha$. For an (α,β)-minimal set U and a β-block B such that $\beta|_{U\cap B} \neq \alpha|_{U\cap B}$, the set $U \cap B$ is said to be an (α,β)-*trace*. A 2-element set $\{a,b\} \subseteq U \cap B$ such that $(a,b) \in \beta - \alpha$, is called an (α,β)-*subtrace*.

Let $\alpha \prec \beta$ and $\gamma \prec \delta$ be prime intervals in $\mathsf{Con}(\mathbb{A})$. We say that (α,β) can be *separated* from (γ,δ) if there is a unary polynomial f of $\mathbb{A}$ such that $f(\beta) \not\subseteq \alpha$, but $f(\delta) \subseteq \gamma$. The polynomial f in this case is said to *separate* (α,β) from (γ,δ).

In a similar way separation can be defined for prime intervals in different coordinate positions of a relation. Let R be a subdirect product of $\mathbb{A}_1 \times \cdots \times \mathbb{A}_n$. Then R is also an algebra and its polynomials can be defined in the same way as for a single algebra. Let $I, J \subseteq [n]$ and let $\alpha \prec \beta$, $\gamma \prec \delta$ be prime intervals in $\mathsf{Con}(\mathrm{pr}_I R)$ and $\mathsf{Con}(\mathrm{pr}_J R)$, respectively. Interval (α,β) can be separated from (γ,δ) if there is a unary polynomial f of R such that $f(\beta) \not\subseteq \alpha$ but $f(\delta) \subseteq \gamma$ (note that the actions of f on $\mathrm{pr}_I R, \mathrm{pr}_J R$ are polynomials of those algebras).

LEMMA 6.9 (LEMMA 35(2), LEMMA 45 OF [25]). *Let $\mathbb{A}_1, \ldots, \mathbb{A}_n$ be arbitrary algebras. Let R, Q be subdirect products of $\mathbb{A}_1, \ldots, \mathbb{A}_n$, $Q \subseteq R$, $I, J \subseteq [n]$, $\mathrm{pr}_I Q = \mathrm{pr}_I R$, $\mathrm{pr}_J Q = \mathrm{pr}_J R$, and $\alpha_I \prec \beta_I$, $\alpha_J \prec \beta_J$ for $\alpha_I, \beta_I \in \mathsf{Con}(\mathrm{pr}_I R)$, $\alpha_J, \beta_J \in \mathsf{Con}(\mathrm{pr}_J R)$.*

(1) *If (α_I, β_I) can be separated from (α_J, β_J) in Q, (α_I, β_I) can also be separated from (α_J, β_J) in R.*
(2) *If (α_I, β_I) and (α_J, β_J) cannot be separated from each other in R, then $\mathsf{typ}(\alpha_I, \beta_I) = \mathsf{typ}(\alpha_J, \beta_J)$.*

If $\mathbb{A}_1, \ldots, \mathbb{A}_n$ are algebras and $B_1, \ldots, B_n$ are their subsets $B_i \subseteq \mathbb{A}_i$, $i \in [n]$, and $\beta_1, \ldots, \beta_n$ are congruences of the $\mathbb{A}_i$'s, it will be convenient to denote $B_1 \times \cdots \times B_n$ by $\overline{B}$ and $\beta_1 \times \cdots \times \beta_n = \{(\mathbf{a},\mathbf{b}) \in (\mathbb{A}_1 \times \cdots \times \mathbb{A}_n)^2 \mid \mathbf{a}[i] \stackrel{\beta_i}{\equiv} \mathbf{b}[i], i \in [n]\}$ by $\overline{\beta}$. For $I \subseteq [n]$ we will use $\overline{B}_I$, for $\prod_{i\in I} B_i$, and $\overline{\beta}_I$ for $\prod_{i \in I} \beta_i$, or if I is clear from the context just $\overline{B}, \overline{\beta}$. By $\mathrm{Cg}_{\mathbb{A}}(D)$, or just $\mathrm{Cg}(D)$ if $\mathbb{A}$ is clear from the context, we denote the congruence of $\mathbb{A}$ generated by a set D of pairs from $\mathbb{A}^2$.

For an algebra $\mathbb{A} \in \mathcal{V}$, $\alpha \in \mathsf{Con}(\mathbb{A})$, a set $\mathcal{U}$ of unary polynomials, and $B \subseteq \mathbb{A}^2$, we denote by $\mathrm{Cg}_{\mathbb{A},\alpha,\mathcal{U}}(B)$ the transitive-symmetric closure of the set $T(B,\mathcal{U}) = \{(f(a), f(b)) \mid (a,b) \in B, f \in \mathcal{U}\} \cup \alpha$. Let $\alpha, \beta \in \mathsf{Con}(\mathbb{A})$, $\alpha \leq \beta$ and D a subuniverse of $\mathbb{A}$. We say that α and β are $\mathcal{U}$-*chained* with respect to D if for any $a, b \in D$, $a \stackrel{\beta}{\equiv} b$, such that $\beta = \mathrm{Cg}_{\mathbb{A}}(\{(a,b)\}) \vee \alpha$ and any β-block B such that $B' = B \cap \mathsf{umax}(D) \neq \varnothing$, we have $(\mathsf{umax}(B'))^2 \subseteq \mathrm{Cg}_{\mathbb{A},\alpha,\mathcal{U}}(\{(a,b)\})$. Note that if $\beta = \mathrm{Cg}_{\mathbb{A}}(\{(a,b)\}) \vee \alpha$ for no $a, b \in D$, then α, β are trivially $\mathcal{U}$-chained.

Let R be a subdirect product of $\mathbb{A}_1, \ldots, \mathbb{A}_n \in \mathcal{V}$, $\beta_i \in \mathsf{Con}(\mathbb{A}_i)$, let B_i be a β_i-block for $i \in [n]$, and let $R' = R \cap \overline{B}$, $B_i' = \mathrm{pr}_i R'$. A unary polynomial f of R is said to be $\overline{B}$-*preserving* if $f(R') \subseteq R'$. We call an n-ary relation R *chained* with respect to $\overline{\beta}, \overline{B}$ if

(Q1) for any $I \subseteq [n]$ and $\alpha, \beta \in \mathsf{Con}(\mathrm{pr}_I R)$ such that $\alpha \le \beta \le \overline{\beta}_I$, α, β are $\mathcal{U}_{\overline{B}}$-chained with respect to $\mathrm{pr}_I R'$, where $\mathcal{U}_{\overline{B}}$ is the set of all $\overline{B}$-preserving polynomials of R;

(Q2) for any $I, J \subseteq [n]$ (note that it may happen that $I \cap J \neq \varnothing$) and $\alpha, \beta \in \mathsf{Con}(\mathrm{pr}_I R)$, $\gamma, \delta \in \mathsf{Con}(\mathrm{pr}_J R)$ such that $\alpha \prec \beta \le \overline{\beta}_I$, $\gamma \prec \delta \le \overline{\beta}_J$, and (α, β) can be separated from (γ, δ), the congruences α and β are $\mathcal{U}(\gamma, \delta, \overline{B})$-chained with respect to $\mathrm{pr}_I R'$, where $\mathcal{U}(\gamma, \delta, \overline{B})$ is the set of all $\overline{B}$-preserving polynomials g of R such that $g(\delta) \subseteq \gamma$.

The following lemma claims that the property to be chained is preserved under certain transformations of $\overline{\beta}$ and $\overline{B}$.

Lemma 6.10 (Lemmas 36,37(3),59,60 of [25]). *Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$. Let $\beta_i \in \mathsf{Con}(\mathbb{A}_i)$ and B_i a β_i-block, $i \in [n]$, and $R' = R \cap \overline{B}$, $B'_i = \mathrm{pr}_i R'$.*

(1) *Let $K \subseteq [n]$ and consider R as a subdirect product R^* of $\mathrm{pr}_K R$ and $\mathbb{A}_i$, $i \in [n] - K$. Let $\beta^*_i = \beta_i$ and $B^*_i = B_i$ for $i \in [n] - K$ and $\beta^*_K = \overline{\beta}_K$ and $B^*_K = \overline{B}_K$. If R is chained with respect to $\overline{\beta}, \overline{B}$ then the relation R^* is chained with respect $\overline{\beta}^*, \overline{B}^*$.*

(2) *Let R be chained with respect to $\overline{\beta}, \overline{B}$. Let $\alpha_i \in \mathsf{Con}(\mathbb{A}_i)$ with $\alpha_i \le \beta_i$, $i \in [n]$. The relation $R/_{\overline{\alpha}}$ is chained with respect to $\overline{\beta}/_{\overline{\alpha}}, \overline{B}/_{\overline{\alpha}}$.*

(3) *For any $I \subseteq [n]$, any $\alpha \le \beta \le \overline{\beta}_I$, any $\mathbf{a}, \mathbf{b} \in \overline{B}'_I$ with $(\mathbf{a}, \mathbf{b}) \in \beta - \alpha$ and $\beta = \mathrm{Cg}_{\mathrm{pr}_I R}(\{(\mathbf{a}, \mathbf{b})\}) \vee \alpha$, and any $\mathbf{c}, \mathbf{d} \in \overline{B}'_I$ such that $(\mathbf{c}, \mathbf{d}) \in \beta$ and $\mathbf{c}/_\alpha, \mathbf{d}/_\alpha$ belong to the same as-component of an $\beta/_\alpha$-block, there is a $\overline{B}$-preserving polynomial f such that $f(\mathbf{a}) \stackrel{\alpha}{\equiv} \mathbf{c}$ and $f(\mathbf{b}) \stackrel{\alpha}{\equiv} \mathbf{d}$.*

(4) *If $\beta_i = \underline{1}_{\mathbb{A}_i}$ and $B_i = \mathbb{A}_i$ for $i \in [n]$, then R is chained with respect to $\overline{\beta}, \overline{B}$.*

(5) *Fix $i \in [n]$, $\beta'_i \prec \beta_i$, and let D_i be a β'_i-block that is as-maximal in $B'_i/_{\beta'_i}$. Let also $\beta'_j = \beta_j$ and $D_j = B_j$ for $j \neq i$. If R is chained with respect to $\overline{\beta}, \overline{B}$ then R is chained with respect to $\overline{\beta}', \overline{D}$.*

Let again R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$ and let $\mathcal{W}^R$ denote the set of triples (i, α, β), where $i \in [n]$ and $\alpha, \beta \in \mathsf{Con}(\mathbb{A}_i)$, $\alpha \prec \beta$. We say that (i, α, β) cannot be separated from (j, γ, δ) if (α, β) cannot be separated from (γ, δ) in R. Then the relation 'cannot be separated' on $\mathcal{W}^R$ is clearly reflexive and transitive. The next lemma shows that it is to some extent symmetric.

Lemma 6.11 (Theorem 38 of [25]). *Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$, for each $i \in [n]$, $\beta_i \in \mathsf{Con}(\mathbb{A}_i)$, B_i a β_i-block such that R is chained with respect to $\overline{\beta}, \overline{B}$; $R' = R \cap \overline{B}$, $B'_i = \mathrm{pr}_i R'$. Let also $I, J \subseteq [n]$, $\alpha \prec \beta \le \overline{\beta}_I$, $\gamma \prec \delta = \overline{\beta}_J$, where $\alpha, \beta \in \mathsf{Con}(\mathrm{pr}_I R)$, $\gamma, \delta \in \mathsf{Con}(\mathrm{pr}_J R)$. If $\overline{B}'_J/_\gamma$ has a nontrivial as-component C_J and (α, β) can be separated from (γ, δ), then there is a $\overline{B}$-preserving polynomial g of R such that $g(\beta|_{\overline{B}'_I}) \subseteq \alpha$ and $g(\delta|_{\overline{B}'_J}) \not\subseteq \gamma$. Moreover, for any $\mathbf{c}, \mathbf{d} \in C_J$, $\mathbf{c} \neq \mathbf{d}$, the polynomial g can be chosen such that $g(\mathbf{c}) = \mathbf{c}, g(\mathbf{d}) = \mathbf{d}$.*

We also introduce polynomials that collapse all prime intervals in congruence lattices of factors of a subproduct, except for a set of intervals that cannot be separated from each other.

Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$, $\beta_j \in \mathsf{Con}(\mathbb{A}_j)$, and B_j a β_j-block, $j \in [n]$. Let also $i \in [n]$, and $\alpha, \beta \in \mathsf{Con}(\mathbb{A}_i)$ be such that $\alpha \prec \beta \le \beta_i$. We call an idempotent unary polynomial f of R *$\alpha\beta$-collapsing for $\overline{\beta}, \overline{B}$* if

(a) f is $\overline{B}$-preserving;

(b) $f(\mathbb{A}_i)$ contains an (α, β)-minimal set, in particular $f(\beta) \not\subseteq \alpha$;

(c) $f(\delta|_{B_j}) \subseteq \gamma|_{B_j}$ for every $(j, \gamma, \delta) \in \mathcal{W}^R$ with $\gamma \prec \delta \le \beta_j$ and such that (α, β) can be separated from (γ, δ) or (γ, δ) can be separated from (α, β).

Lemma 6.12 (Theorem 53 of [25]). *Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$, and choose $\beta_j \in \mathsf{Con}(\mathbb{A}_j)$ and a β_j-block B_j for each $j \in [n]$; let R be chained with respect to $\overline{\beta}, \overline{B}$ and $R' = R \cap \overline{B}$, $B'_j = \mathrm{pr}_j R'$, $j \in [n]$. Let* *also $I \subseteq [n]$, and $\alpha \in \mathsf{Con}(\mathrm{pr}_I R)$ be such that $\alpha \prec \overline{\beta}_I$. If $\overline{B}'_I/_\alpha$ contains a nontrivial as-component, then there exists an $\alpha\overline{\beta}_I$-collapsing polynomial f for $\overline{\beta}, \overline{B}$. Moreover, f can be chosen to satisfy any one of the following conditions:*

(d) *for any $(\alpha, \overline{\beta}_I)$-subtrace $\{\mathbf{a}, \mathbf{b}\} \subseteq \mathsf{amax}(\mathrm{pr}_I R')$ with $\mathbf{b} \in \mathsf{as}(\mathbf{a})$ and any $J \subseteq [n]$, the polynomial f can be chosen such that $\mathbf{a}/_\alpha, \mathbf{b}/_\alpha \in f(\mathrm{pr}_I R)/_\alpha$ and $f(\overline{B}_J) \cap \mathsf{umax}(\overline{B}'_J) \neq \varnothing$;*

(e) *if $\mathsf{typ}(\alpha, \overline{\beta}_I) = 2$, for any $\mathbf{a} \in \mathsf{umax}(R')$ the polynomial f can be chosen such that $f(\mathbf{a}) = \mathbf{a}$;*

(f) *if $\mathsf{typ}(\alpha, \overline{\beta}_I) = 2$, a tuple $\mathbf{a} \in R'$ is such that $\mathbf{a} \in \mathsf{umax}(R'')$, where $R'' = \{\mathbf{b} \in R \mid \mathrm{pr}_I \mathbf{b} \stackrel{\alpha}{\equiv} \mathrm{pr}_I \mathbf{a}\}$ and $\{\mathbf{a}_0, \mathbf{b}_0\} \subseteq \mathsf{amax}(\mathrm{pr}_I R')$ is an $(\alpha, \overline{\beta}_I)$-subtrace such that $\mathrm{pr}_I \mathbf{a} = \mathbf{a}_0$ and $\mathbf{b}_0 \sqsubseteq^{as} \mathbf{b}_0$, then the polynomial f can be chosen such that $f(\mathbf{a}) = \mathbf{a}$ and $\mathbf{b}' \in f(\mathrm{pr}_I R)$ for some $\mathbf{b}' \stackrel{\alpha}{\equiv} \mathbf{b}_0$.*

We will also need the following auxiliary lemma.

LEMMA 6.13 (LEMMA 52 OF [25]). *Let $\mathbb{A} \in \mathcal{V}$ and $\alpha, \beta \in \mathsf{Con}(\mathbb{A})$ be such that $\alpha \prec \beta$ and $\mathsf{typ}(\alpha, \beta) = 2$; let B be a β-block containing more than one α-block, and $a, b \in B$ with $a \sqsubseteq^{asm} b$ in B.*

(1) *There exists a polynomial f such that $f(a) = b$ and $f(\beta|_B) \not\subseteq \alpha$.*

(2) *If a belongs to an (α, β)-trace, so does b. In particular, every element from $\mathsf{umax}(B)$ belongs to an (α, β)-trace.*

(3) *Let $a \stackrel{\alpha}{\equiv} b$, $a \sqsubseteq^{asm} b$ in $a/_\alpha$, and N an (α, β)-trace with $a \in N$. Then there is a polynomial f such that $f(a) = b$, $N' = f(N)$ is an (α, β)-trace containing b, and $N'/_\alpha = N/_\alpha$.*

6.4 The Congruence Lemma

This section contains a technical result, the Congruence Lemma 6.15, that will be used when proving Theorem 5.6. We start with introducing two closure properties of algebras and their subdirect products.

Let $\mathbb{C}$ be a subalgebra of $\mathbb{A} \in \mathcal{V}$. A subset $B \subseteq \mathbb{C}$ is *as-closed* in $\mathbb{C}$ if for any $a, b \in \mathbb{C}$ with $a \in \mathsf{umax}(B)$ and $a \sqsubseteq^{as} b$ in $\mathbb{C}$, it holds that $b \in B$. Similarly, the set B is *s-closed* in $\mathbb{C}$ if for any $a, b \in \mathbb{C}$ with $a \in \mathsf{umax}(B)$ and $a \leq b$ in $\mathbb{C}$, it holds that $b \in B$. Thus, an as-closed (s-closed) set is just a set of elements closed under thin semilattice and affine edges (respectively, thin semilattice edges). Note that the subalgebra $\mathbb{C}$ is very important here, as we normally want to 'contain' as-closed (s-closed) sets within some subalgebra, and thin edges do not respect subalgebras.

Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$ and Q a subalgebra of R. By $\mathsf{Cg}(Q)$ we denote the congruence of R generated by pairs of elements from Q. That is, $\mathsf{Cg}(Q)$ is the smallest congruence such that Q is a subset of a $\mathsf{Cg}(Q)$-block, denote it $\mathsf{Block}(Q)$. The subalgebra Q is said to be *polynomially closed* in R if Q is as-closed in $\mathsf{Block}(Q)$. A subset $S \subseteq Q$ is said to be *weakly as-closed* in Q if for any $i \in [n]$, $\mathrm{pr}_i S$ is as-closed in $\mathrm{pr}_i Q$.

Polynomially closed subalgebras, as- and s-closed subsets are well behaved with respect to some standard algebraic transformations.

LEMMA 6.14 (LEMMA 55 OF [25]). *Let $\mathbb{A} \in \mathcal{V}$ and let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$.*

(1) *R is polynomially closed in R and $\mathbb{A}$ is as-closed and s-closed in $\mathbb{A}$.*

(2) *For any congruence $\beta \in \mathsf{Con}(R)$ and a β-block Q, the subalgebra Q is polynomially closed in R.*

(3) *Let Q_1, Q_2 be subalgebras of R, Q_1, Q_2 polynomially closed in R, and $\mathsf{umax}(Q_1) \cap \mathsf{umax}(Q_2) \neq \varnothing$. Then $Q_1 \cap Q_2$ is polynomially closed in R.*

In particular, let $\beta \in \mathsf{Con}(R)$ and T a β-block such that $\mathsf{umax}(Q_1) \cap \mathsf{umax}(T) \neq \varnothing$. Then $Q_1 \cap T$ is polynomially closed in R.

Let $\mathbb{B}$ be a subalgebra of $\mathbb{A}$, and let $\mathbb{C}_1, \mathbb{C}_2$ be subalgebras of $\mathbb{B}$ as-closed (s-closed) in $\mathbb{B}$ and such that $\mathsf{umax}(\mathbb{C}_1) \cap \mathsf{umax}(\mathbb{C}_2) \neq \varnothing$. Then $\mathbb{C}_1 \cap \mathbb{C}_2$ is as-closed (respectively, s-closed) in $\mathbb{B}$.

(4) *Let $\mathbb{B}$ be a subalgebra of $\mathbb{A}$ and a subalgebra $\mathbb{C} \subseteq \mathbb{B}$ as-closed (s-closed) in $\mathbb{B}$. Then for any $\beta \in \mathsf{Con}(\mathbb{A})$, the algebra $\mathbb{C}/_\beta$ is as-closed (respectively, s-closed) in $\mathbb{B}/_\beta$.*

(5) *Let R_i, $i \in [k]$ be a subdirect product of some algebras from $\mathcal{V}$, and Q_i polynomially closed in R_i, $i \in [k]$. Let R, Q be conjunctive-defined through $R_1, \dots, R_k$ and $Q_1, \dots, Q_k$, respectively, by the same conjunctive formula Φ; that is, $R = \Phi(R_1, \dots, R_k)$ and $Q = \Phi(Q_1, \dots, Q_k)$. In other words R consists of all tuples that are satisfying assignments of Φ which uses the R_i's as atoms, while Q is obtained in the same way only replacing the R_i's with the Q_i's. Suppose that for every atom $Q_i(x_{j_1}, \dots, x_{j_\ell})$ it holds that $\mathsf{umax}(\mathrm{pr}_{\{x_{j_1},\dots,x_{j_\ell}\}} Q) \subseteq \mathsf{umax}(Q_i)$. Then Q is polynomially closed in R.*

(6) *Let R_i be a subdirect product of some algebras from $\mathcal{V}$, Q_i subalgebra of R_i, and T_i a subalgebra of Q_i as-closed (s-closed) in Q_i, $i \in [k]$. Let R, Q, and T be pp-defined through $R_1, \dots, R_k$, $Q_1, \dots, Q_k$, and $T_1, \dots, T_k$, respectively, by the same pp-formula $\exists\overline{x}\Phi$; that is, $R = \exists\overline{x}\Phi(R_1, \dots, R_k)$, $Q = \exists\overline{x}\Phi(Q_1, \dots, Q_k)$, and $T = \exists\overline{x}\Phi(T_1, \dots, T_k)$. Let also $T' = \Phi(T_1, \dots, T_k)$ and, and suppose that for every atom $T_i(x_{j_1}, \dots, x_{j_\ell})$ it holds that $\mathsf{umax}(\mathrm{pr}_{\{x_{j_1},\dots,x_{j_\ell}\}} T') \subseteq \mathsf{umax}(T_i)$. Then T is as-closed (respectively, s-closed) in Q.*

In most cases we use the condition of polynomial closedness as follows. Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$ and Q its subalgebra. Let $\mathbf{a}, \mathbf{b} \in Q$, $\mathbf{a} \in \mathsf{umax}(Q)$, and f a polynomial of R such that $f(\mathbf{a}) = \mathbf{a}$. Then $(f(\mathbf{a}), f(\mathbf{b})) \in \mathsf{Cg}(Q)$ and therefore $f(\mathbf{b}) \in \mathsf{Block}(Q)$. From this we conclude that any $\mathbf{c} \in R$ such that $\mathbf{a} \sqsubseteq^{as} \mathbf{c}$ in the subalgebra generated by $\mathbf{a}, f(\mathbf{b})$ belongs to Q.

We are now in a position to state the Congruence Lemma.

LEMMA 6.15 (THE CONGRUENCE LEMMA, LEMMA 56 OF [25]). *Let R be a subdirect product of $\mathbb{A}_1, \dots, \mathbb{A}_n \in \mathcal{V}$, β_i a congruence of $\mathbb{A}_i$ and let B_i be a β_i-block, $i \in [n]$. Also, let R be chained with respect to $\overline{\beta}, \overline{B}$ and $R' = R \cap \overline{B}$, $B'_i = \mathrm{pr}_i R'$. Let $I_0 \subseteq [n]$ (without loss of generality assume say, $I_0 = [\ell]$) and $\alpha \in \mathsf{Con}(\mathrm{pr}_{I_0} R)$ be such that $\alpha \prec \overline{\beta}_{I_0} = \prod_{i \in I_0} \beta_i$, let $I \subseteq [n] - I_0$, $I' = I \cup I_0$, $J = [n] - I'$, and let Q be a subalgebra of R' polynomially closed in R and such that $E_1 = \mathrm{pr}_{I_0} Q$ contains an as-component C of $\overline{B}'_{I_0}$ and $Q \cap \mathsf{umax}(R') \neq \varnothing$. We consider R, R', and Q as subdirect products of $\mathrm{pr}_{I_0} R \times \mathrm{pr}_I R \times \mathrm{pr}_J R$, $\mathrm{pr}_{I_0} R' \times \mathrm{pr}_I R' \times \mathrm{pr}_J R'$, and $\mathrm{pr}_{I_0} Q \times \mathrm{pr}_I Q \times \mathrm{pr}_J Q$, respectively. Let $R'/_\alpha = R'/_{\alpha \times \underline{0}_{\ell+1} \times \dots \times \underline{0}_n}$, $Q/_\alpha = Q/_{\alpha \times \underline{0}_{\ell+1} \times \dots \times \underline{0}_n}$, and let $\mathsf{lk}'_1, \mathsf{lk}'_2$ and $\mathsf{lk}^Q_1, \mathsf{lk}^Q_2$ be the link congruences of $\mathrm{pr}_{I'} R'/_\alpha$ and $\mathrm{pr}_{I'} Q/_\alpha$, respectively, and $E_2 = \mathrm{pr}_I Q$. Let also $E_2^C = Q[C/_\alpha] = \{\mathrm{pr}_I \mathbf{a} \mid \mathbf{a} \in Q, \mathrm{pr}_{I_0} \mathbf{a}/_\alpha \in C/_\alpha\}$. Then either*

(1) *$C/_\alpha \times \mathsf{umax}(E_2^C) \subseteq \mathrm{pr}_{I'} Q/_\alpha$, or*

(2) *there is $\eta \le \overline{\beta}_I$, maximal (under inclusion) congruence of $\mathrm{pr}_I R$ with $\eta|_{\mathsf{umax}(E_2^C)} \subseteq \mathsf{lk}^Q_2$, and $\gamma = \mathsf{Cg}_{\mathrm{pr}_I R}(\mathsf{umax}(E_2^C)) \vee \eta$, such that $\eta \prec \gamma \le \beta_I$ and the intervals $(\alpha, \overline{\beta}_{I_0})$ and (η, γ) cannot be separated.*

Moreover, in case (2) $\mathrm{pr}_{I'} Q/_\alpha$ is the graph of a mapping $\varphi : E_2 \to E_1/_\alpha$ such that the kernel lk^Q_2 of φ is the restriction of η on E_2, and for any $\overline{B}$-preserving polynomial f such that $f(\overline{\beta}_{I_0}|_{\overline{B}_{I_0}}) \not\subseteq \alpha$ it holds that $f(\gamma|_{\overline{B}_I}) \not\subseteq \eta$.[6]

7 DECOMPOSITIONS AND COMPRESSED PROBLEMS

In this section we apply the machinery developed in the previous section to constraint satisfaction problems in order to prove Theorem 5.6.

[6]This property is somewhat stronger than non-separability. The non-separability of $(\alpha, \overline{\beta}_{I_0})$ and (η, γ) only implies that $f(\gamma) \not\subseteq \eta$ without any restrictions on $f(\gamma|_{\overline{B}_I})$.

7.1 Decomposition of CSPs

We begin with showing how separating congruence intervals and centralizers can be combined to obtain strands and therefore useful decompositions of CSPs. The case of binary relations is settled in [25].

LEMMA 7.1 (LEMMA 47 OF [25]). *Let R be a subdirect product of $\mathbb{A}_1 \times \mathbb{A}_2$, $\alpha_i, \beta_i \in \mathsf{Con}(\mathbb{A}_i)$, $\alpha_i \prec \beta_i$, for $i = 1, 2$. If (α_1, β_1) and (α_2, β_2) cannot be separated from each other, then the coordinate positions 1,2 are $\zeta_1\zeta_2$-aligned in R, where $\zeta_1 = (\alpha_1 : \beta_1), \zeta_2 = (\alpha_2 : \beta_2)$.*

Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal instance with subdirectly irreducible domains and let $\overline{\beta}$, $\beta_v \in \mathsf{Con}(\mathbb{A}_v)$, $v \in V$, be a collection of congruences. Let $\mathcal{W}^{\mathcal{P}}(\overline{\beta})$ denote the set of triples (v, α, β) such that $v \in V$, $\alpha, \beta \in \mathsf{Con}(\mathbb{A}_v)$, and $\alpha \prec \beta \leq \beta_v$. Also, $\mathcal{W}^{\mathcal{P}}$ denotes $\mathcal{W}^{\mathcal{P}}(\overline{\beta})$ when $\beta_v = \underline{1}_v$ for all $v \in V$. We will omit the superscript $\mathcal{P}$ whenever it is clear from the context. Let also $\mathcal{W}'^{\mathcal{P}}(\overline{\beta})$, $\mathcal{W}'^{\mathcal{P}}$, $\mathcal{W}'$ denote the set of triples (v, α, β) from $\mathcal{W}^{\mathcal{P}}(\overline{\beta})$, $\mathcal{W}^{\mathcal{P}}$, $\mathcal{W}$, respectively, for which $(\alpha : \beta) = \underline{1}_v$. For every $(v, \alpha, \beta) \in \mathcal{W}(\overline{\beta})$, let $Z(v, \alpha, \beta, \overline{\beta})$ denote the set of triples $(w, \gamma, \delta) \in \mathcal{W}(\overline{\beta})$ such that (α, β) and (γ, δ) cannot be separated in R^{vw} (recall that we always assume that a (2,3)-minimal instance contains all unary and binary constraints). Slightly abusing the terminology we will also say that (α, β) and (γ, δ) cannot be separated in $\mathcal{P}$. Then let $W(v, \alpha, \beta, \overline{\beta}) = \{w \in V \mid (w, \gamma, \delta) \in Z(v, \alpha, \beta, \overline{\beta})$ for some $\gamma, \delta \in \mathsf{Con}(\mathbb{A}_w)\}$. We will omit mentioning of $\overline{\beta}$ whenever possible. Sets of the form $W(v, \alpha, \beta, \overline{\beta})$ will be called *$\overline{\beta}$-coherent sets*, or just coherent sets if $\overline{\beta}$ is clear from the context. Also, if $(\alpha : \beta) \neq \underline{1}_v$ then the corresponding coherent set is called *non-central*. The following statement is an easy corollary of Lemma 7.1.

THEOREM 7.2. *Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal instance and $(v, \alpha, \beta) \in \mathcal{W}$, let $\beta_v = \underline{1}_v$ for $v \in V$. For $w \in W(v, \alpha, \beta, \overline{\beta})$ let $(w, \gamma, \delta) \in \mathcal{W}$ be such that (α, β) and (γ, δ) cannot be separated and $\zeta_w = (\gamma : \delta)$. Then $\mathcal{P}_{W(v,\alpha,\beta,\overline{\beta})}$ is $\overline{\zeta}$-aligned.*

Corollary 7.3 relates domains with congruence intervals that cannot be separated with strands.

COROLLARY 7.3. *Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal instance and W a non-central coherent set. Then W is a subset of a strand.*

For technical reasons we will also count the empty set as a non-central coherent set.

7.2 Compound relations

In our proof we will need the property of chaining for several auxiliary constructions. Due to the way we prove chaining — through Lemma 6.10(5), — it is convenient to include them in the conditions (S1)–(S6) in Section 7.3. The following five types of relations associated with an instance $\mathcal{P}$ will be called *compound relations*.

Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal and block-minimal instance over $\mathcal{V}$ with subdirectly irreducible domains. Let $C = \{C_1, \ldots, C_\ell\}$, $C_j = \langle \mathbf{s}_j, R_j \rangle$. Recall that for a strand $W \subseteq V$ by $\mathcal{P}_{/W}$ we denote the problem $\mathcal{P}/_{\overline{\mu}_{/W}}$, where $\overline{\mu}_{/W} = \overline{\mu}^Y$ and $Y = \mathsf{MAX}(\mathcal{P}) - W$. We will refer to a component in the coordinate position x of the product congruence $\overline{\mu}_{/W}$ as $\mu_{x/W}$. Let also $\mathcal{S}_{/W}$ denote the set of solutions of $\mathcal{P}_{/W}$. If W is a non-central coherent set, the problem $\mathcal{P}_{/W}$ is defined in the same way.

- Let $W \subseteq V$ be a non-central coherent set, and let $C = \langle \mathbf{s}, R \rangle \in C$. We define $P[C, W]$ to be the subalgebra of the product $R \times \mathcal{S}_{/W}$ that consists of all tuples $(\mathbf{a}, \varphi)$ such that φ is a solution of $\mathcal{S}_{/W}$ with $\mathbf{a} \in \varphi(\mathbf{s})$. For a non-central coherent set U we also define

$$P^U[C, W] = \{(\mathbf{a}/_{\overline{\mu}_{/U}}, \varphi) \mid (\mathbf{a}, \varphi) \in P[C, W]\}.$$

- Let $W \subseteq V$ be a non-central coherent set and $x, y, w \in V$, where x, y, w are all distinct. Set $V^i = \{u^i \mid u \in V\}$, $i = 1, 2$, and

$$S[x,y,w,W](x,y,w,V^1,V^2) = R^{xy}(x,y) \wedge P[C^{xw},W](x,w,V^1) \wedge P[C^{yw},W](y,w,V^2).$$

- Let $W, U \subseteq V$ be non-central coherent sets. Set $V^i = \{u^i \mid u \in V\}$, $i \in [\ell]$, and

$$S[W,U](V,V^1,\dots,V^\ell) = \mathcal{S}_{/U}(V) \wedge \bigwedge_{j=1}^{\ell} P^U[C_j,W](\mathbf{s}_j,V^j).$$

- Let $x, y, w, v \in V$. Then set

$$S[x,y,w,v](x,y,w,v) = R^{xy}(x,y) \wedge R^{xw}(x,w) \wedge R^{yw}(y,w) \wedge R^{wv}(w,v).$$

- Let U be a non-central coherent set and $v, w \in V$. Then set

$$S[U,w,v](V,v') = \mathcal{S}_{/U}(V) \wedge R^{wv}/_{\mu_{w/U}\times \underline{0}_v}(w,v').$$

The set of all compound relations associated with the instance $\mathcal{P}$ will be denoted by $\mathsf{Comp}(\mathcal{P})$.

We will use the notation $\overline{\beta}, \overline{B}$ for compound relations choosing a congruence β_x and a β_x-block for each variable x involved. However, since in such a relation (a copy of) a variable $x \in V$ may occur more than once, $\overline{\beta}, \overline{B}$ in this case use the same β_x, B_x for all copies of x.

The following property of compound relations will be important.

LEMMA 7.4. *Let $\mathcal{P} = (V, \mathcal{C})$ be a $(2,3)$-minimal and block-minimal instance over $\mathcal{V}$ with subdirectly irreducible domains, $\beta_v \in \mathsf{Con}(\mathbb{A}_v)$, let B_v be a β_v-block for $v \in V$, and let $\overline{\beta} = \prod_{v\in V}\beta_v$, $\overline{B} = \prod_{v\in V} B_v$. Let a problem instance $\mathcal{P}^\dagger = (V, \mathcal{C}^\dagger)$ be such that for each $\langle \mathbf{s}, R\rangle \in \mathcal{C}$, $\mathcal{C}^\dagger$ contains $\langle \mathbf{s}, R^\dagger\rangle$, $R^\dagger \subseteq R \cap \overline{B}$, and such that $R^{xy\dagger}$, $x, y \in V$, constitute a $(2,3)$-strategy for $\mathcal{P}^\dagger$, and for every non-central coherent set W the problem $\mathcal{P}^\dagger_{/W}$ is minimal (that is, $\mathcal{P}^\dagger$ satisfies the conditions (S1)–(S3) from Section 7.3). Then for any $Q \in \mathsf{Comp}(\mathcal{P})$ it holds that $R^{t\dagger}/_\gamma \subseteq \mathrm{pr}_t Q \cap \overline{B}$ for every t from the scope of the relation Q, where $\gamma \in \{\underline{0}_t, \mu_t\}$ depending on the domain of t in Q.*

PROOF. We consider all five types of compound relations one by one. Let $W, U \subseteq V$ be non-central coherent sets, $\mathcal{S}^\dagger_{/W}, \mathcal{S}^\dagger_{/U}$ denote the set of solutions of $\mathcal{P}^\dagger_{/W}, \mathcal{P}^\dagger_{/U}$, respectively, $C \in \mathcal{C}$, and $x, y, z, w, v \in V$.

$\mathbf{P[C,W]}$. Let $C = \langle \mathbf{s}, R\rangle$, $x \in \mathbf{s}$, and $a \in \mathsf{umax}(R^{x\dagger})$. By the Maximality Lemma 6.5(2) there is $\mathbf{a} \in \mathsf{umax}(R^\dagger)$ such that $\mathbf{a}[x] = a$. By (S3) there is $\varphi \in \mathcal{S}^\dagger_{/W}$ such that $\mathbf{a}/_{\overline{\mu}_{/W}} = \varphi(\mathbf{s})$. Thus, by definition $(\mathbf{a}, \varphi) \in P[C,W]$, and the result holds for $x \in \mathbf{s}$. Note that in this case an even stronger property holds: $R^\dagger \subseteq \mathrm{pr}_{\mathbf{s}} P[C,W] \cap \overline{B}$.

Now, let $a \in \mathsf{umax}(R^{v\dagger}/_{\mu_{v/U}})$, $v \in V$. Then again by (S3) there is a solution $\varphi \in \mathcal{S}^\dagger_{/W}$ such that $\varphi(v) = a$. Let $\overline{\mathbf{a}} = \varphi(\mathbf{s})$. Since $\overline{\mathbf{a}} \in R/_{\overline{\mu}_{/W}}$, there is $\mathbf{a} \in R$ such that $\mathbf{a}/_{\overline{\mu}_{/W}} = \overline{\mathbf{a}}$. By definition $(\mathbf{a}, \varphi) \in P[C,W]$, and the result holds for the remaining coordinates as well.

$\mathbf{S[x,y,w,W]}$. By (S2) for any $a \in R^{x\dagger}$ there are $b \in R^{y\dagger}, c \in R^{w\dagger}$ such that $(a,b) \in R^{xy\dagger}, (b,c) \in R^{yw\dagger}, (a,c) \in R^{xw\dagger}$. By what is proved for $P[C,W]$ the pairs $(a,c), (b,c)$ can be extended to triples $(a,c,\varphi_1) \in P[C^{xw},W]$ and $(b,c,\varphi_2) \in P[C^{yw},W]$ such that $\varphi_1, \varphi_2 \in \mathcal{S}^\dagger_{/W}$.

By the result for $P[C,W]$, for any $\varphi_1 \in \mathcal{S}^\dagger_{/W}$ there is $(a,c) \in R^{xw\dagger}$ such that $(a,c,\varphi_1) \in P[C^{xw},W]$. By (S2) there is $b \in R^{y\dagger}$ such that $(a,b) \in R^{xy\dagger}$ and $(b,c) \in R^{yw\dagger}$. Again by the result for $P[C,W]$ there is $\varphi_2 \in \mathcal{S}^\dagger_{/W}$ such that $(b,c,\varphi_2) \in P[C^{yw},W]$. For the remaining variable the proof is basically the same.

$S[\mathbf{W}, \mathbf{U}]$. We follow the same lines as in the previous item. By (S3) for any $x \in V$ and $a \in R^{x\dagger}/_{\overline{\mu}_{x/U}}$ there is a solution $\varphi \in \mathcal{S}^{\dagger}_{/U}$ such that $a = \varphi(x)$. Then for every $C_j = \langle \mathbf{s}_j, R_j \rangle \in C$ we have $\varphi(\mathbf{s}_j) = R^{\dagger}_j/_{\overline{\mu}_{/U}}$, hence, there is $\mathbf{a}_j \in R^{\dagger}_j$ with $\mathbf{a}_j/_{\overline{\mu}_{/U}} = \varphi(\mathbf{s}_j)$. By what is proved for $P[C_j, W]$ this tuple can be extended to $(\mathbf{a}_j/_{\overline{\mu}_{/U}}, \varphi_j) \in P^U[C_j, W]$ such that $\varphi_j \in \mathcal{S}^{\dagger}_{/W}$.

For $j \in [\ell]$ and $\varphi_j \in \mathcal{S}^{\dagger}_{/W}$ there is $\mathbf{a}_j \in R^{\dagger}_j$ such that $(\mathbf{a}_j, \varphi_j) \in P[C_j, W]$. By (S3) there exists a solution $\varphi \in \mathcal{S}^{\dagger}_{/U}$ such that $\varphi(\mathbf{s}_j) = \mathbf{a}_j/_{\overline{\mu}_{/U}}$. This solution can then be extended to V^i, $i \in [\ell] - \{j\}$, as before.

$S[\mathbf{x}, \mathbf{y}, \mathbf{w}, \mathbf{v}], S[\mathbf{U}, \mathbf{w}, \mathbf{v}]$. In these cases the proof is similar to the previous cases. □

7.3 Compressed problems

In this section we define a way to tighten a block-minimal problem instance in such a way that it remains (similar to) block-minimal. More precisely, we introduce several properties of a subproblem of a CSP instance $\mathcal{P}$ that are preserved when the problem is restricted in a certain way.

LEMMA 7.5. *Let $\mathcal{P}$ be a (2,3)-minimal and block minimal problem with subdirectly irreducible domains. Then for every non-central coherent set W the problem $\mathcal{P}_{/W}$ is minimal.*

PROOF. By Corollary 7.3 there is a strand $U \subseteq V$ such that $W \subseteq U$. It now suffices to observe that for every solution $\varphi \in \mathcal{S}_{/U}$ of $\mathcal{P}_{/U}$ the mapping $\varphi/_{\overline{\mu}_{/W}}$ is a solution of $\mathcal{P}_{/W}$. □

Let $\mathcal{P} = (V, C)$ over $\mathcal{V}$ be a (2,3)-minimal and block minimal problem with subdirectly irreducible domains. Let $\beta_v \in \mathrm{Con}(\mathbb{A}_v)$ and let B_v be a β_v-block, $v \in V$, $\overline{\beta} = \prod_{v \in V} \beta_v$, $\overline{B} = \prod_{v \in V} B_v$. A problem instance $\mathcal{P}^{\dagger} = (V, C^{\dagger})$, where $\langle \mathbf{s}, R^{\dagger} \rangle \in C^{\dagger}$ if and only if $\langle \mathbf{s}, R \rangle \in C$, is said to be $(\overline{\beta}, \overline{B})$-*compressed* from $\mathcal{P}$ if the following conditions hold:

(S1) for every $\langle \mathbf{s}, R \rangle \in C$ the relation $R^{\dagger}$ is a nonempty subalgebra of $R \cap \overline{B}$ such that $\mathsf{umax}(\mathrm{pr}_v R^{\dagger}) = \mathsf{umax}(R^{v\dagger})$ for $v \in \mathbf{s}$;
(S2) the relations $R^{X\dagger}$, where $R^{X\dagger}$ is obtained from R^X for $X \subseteq V$, $|X| \le 2$, form a nonempty $(2,3)$-strategy for $\mathcal{P}^{\dagger}$;
(S3) for every non-central coherent set W the problem $\mathcal{P}^{\dagger}_{/W} = \mathcal{P}^{\dagger}/_{\overline{\mu}_{/W}}$ is minimal;
(S4) the relation R for every $\langle \mathbf{s}, R \rangle \in C$, the relations $\mathcal{S}_{/W}$ for every non-central coherent set $W \subseteq V$ (this includes $W = \varnothing$), and every compound relation from $\mathrm{Comp}(\mathcal{P})$ are chained with respect to $\overline{\beta}, \overline{B}$;
(S5) for every $\langle \mathbf{s}, R \rangle \in C$ and any non-central coherent set W the subalgebras $R^{\dagger}, R^{\dagger}/_{\overline{\mu}_{/W}}$ are polynomially closed in $R, R/_{\overline{\mu}_{/W}}$, respectively;
(S6) for every $\langle \mathbf{s}, R \rangle \in C$ the subalgebra $R^{\dagger}$ is weakly as-closed and s-closed in $R \cap \overline{B}$.

Conditions (S1)–(S3) are the conditions we actually want to maintain when constructing a compressed instance, and these are the ones that provide the desired results. However, to prove that (S1)–(S3) are preserved under transformations of compressed instances we also need more technical conditions (S4)–(S6).

We now show how we plan to use compressed instances. Let $\mathcal{P}$ be a subdirectly irreducible, (2,3)-minimal, and block-minimal instance, $\beta_v = \underline{1}_v$ and $B_v = \mathbb{A}_v$ for $v \in V$. Then as is easily seen the instance $\mathcal{P}$ itself is $(\overline{\beta}, \overline{B})$-compressed from $\mathcal{P}$. Also, by (S1) a $(\overline{\gamma}, \overline{D})$-compressed instance with $\gamma_v = \underline{0}_v$ for all $v \in V$ gives a solution of $\mathcal{P}$. Our goal is therefore to show that a $(\overline{\beta}, \overline{B})$-compressed instance for any $\overline{\beta}$ and an appropriate $\overline{B}$ can be 'reduced', that is, transformed to a $(\overline{\beta}', \overline{B}')$-compressed instance for some $\overline{\beta}' < \overline{\beta}$. Note that this reduction of instances is where the condition $\mathsf{MAX}(\mathcal{P}) \cap \mathsf{Center}(\mathcal{P}) = \varnothing$ is used. Indeed, suppose that $\beta_v = \mu^*_v$ (see Section 5.4). Then by conditions (S1)–(S6) we only have information about solutions to problems of the form $\mathcal{P}/_{\overline{\mu}^*}$ or something very close to that. Therefore this barrier cannot be penetrated. We consider two cases.

CASE 1. There are $v \in V$ and $\alpha \prec \beta_v$ nontrivial on B_v, $\mathsf{typ}(\alpha, \beta_v) = \mathbf{2}$. This case is considered in Section 8.

CASE 2. For all $v \in V$ and $\alpha \prec \beta_v$ nontrivial on B_v, $\mathsf{typ}(\alpha, \beta_v) \in \{\mathbf{3}, \mathbf{4}, \mathbf{5}\}$. This case is considered in Section 9.

There is also the possibility that $\alpha|_{R^{v\dagger}} = \beta_v|_{R^{v\dagger}}$ for all $\alpha \prec \beta_v$. In this case we can replace β_v with a smaller congruence without violating any of the conditions (S1)–(S6).

8 PROOF OF THEOREM 5.6: AFFINE FACTORS

In this section we consider Case 1 of tightening instances: there is $\alpha \in \mathsf{Con}(\mathbb{A}_v)$ for some $v \in V$ such that $\alpha \prec \beta_v$ and $\mathsf{typ}(\alpha, \beta_v) = \mathbf{2}$.

Let $\mathcal{P} = (V, C)$ be a $(2, 3)$-minimal and block-minimal instance with subdirectly irreducible domains, $\overline{\beta} = \prod_{w \in V} \beta_w$, $\beta_w \in \mathsf{Con}(\mathbb{A}_w)$, and $\overline{B} = \prod_{w \in V} B_w$, B_w is a β_w-block, $w \in V$. Let $\mathcal{W}, \mathcal{W}'$ denote $\mathcal{W}^{\mathcal{P}}(\overline{\beta})$, $\mathcal{W}'^{\mathcal{P}}(\overline{\beta})$, respectively. Let also $\mathcal{P}^\dagger = (V, C^\dagger)$ be a $(\overline{\beta}, \overline{B})$-compressed instance, and for $C = \langle \mathbf{s}, R\rangle \in C$ there is $C^\dagger = \langle \mathbf{s}, R^\dagger\rangle \in C^\dagger$. We select $v \in V$ and $\alpha \in \mathsf{Con}(\mathbb{A}_v)$ with $\alpha \prec \beta_v$, $\alpha|_{R^{v\dagger}} \neq \beta_v|_{R^{v\dagger}}$, $\mathsf{typ}(\alpha, \beta_v) = \mathbf{2}$, and an α-block $B \in R^{v\dagger}/_\alpha$. Note that since $\mathsf{typ}(\alpha, \beta_v) = \mathbf{2}$, $B_v/_\alpha$ is a module, and therefore B is as-maximal in this set. In this section we show how $\mathcal{P}^\dagger$ can be transformed to a $(\overline{\beta}', \overline{B}')$-compressed instance such that $\beta'_w \le \beta_w$, $B'_w \subseteq B_w$ for $w \in V$, and $\beta'_v = \alpha$, $B'_v = B$. Let also $W = W(v, \alpha, \beta_v, \overline{\beta})$, and let $\mathcal{S}^\dagger_{/U}$ denote the set of solutions of $\mathcal{P}^\dagger_{/U}$ for a non-central coherent set U.

8.1 Properties of compound relations

We start with proving several useful properties of compound relations. Let $W' = W$ if $(v, \alpha, \beta_v) \notin \mathcal{W}'$ and $W' = \varnothing$ otherwise. To save on notation we will use $\overline{\mu}^\circ$ for $\overline{\mu}_{/W'}$. Finally, let $P^\dagger[C, W']$ be defined the same way as $P[C, W']$ only using $R^\dagger, \mathcal{S}^\dagger_{/W'}$ instead of $R, \mathcal{S}_{/W'}$, and for a non-central coherent set U, let $P^{U\dagger}[C, W'] = P^\dagger[C, W']/_{\overline{\mu}_{/U}}$. We use compound relations $P[C, W']$, $P^\dagger[C, W']$, $S[x, y, w, W']$, and $S[W', U]$.

LEMMA 8.1. *Let $C = \{C_1, \dots, C_\ell\}$.*

(1) *For any non-central coherent set U the relation $\mathcal{S}^\dagger_{/U}$ is polynomially closed in $\mathcal{S}_{/U}$.*

(2) *$P[C, W'], P^U[C, W']$, $C = \langle \mathbf{s}, R\rangle \in C$, U a non-central coherent set, can be represented by the pp-definitions*

$$
\begin{aligned}
P[C, W'](\mathbf{s}, V^0) &= R(\mathbf{s}) \wedge \mathcal{S}_{/W'}(V^0) \wedge \bigwedge_{w \in \mathbf{s}} T_w(w, w^0),\\
P^U[C, W'](\mathbf{s}, V^0) &= R/_{\overline{\mu}_{/U}}(\mathbf{s}) \wedge \mathcal{S}_{/W'}(V^0) \wedge \bigwedge_{w \in \mathbf{s}} T_w/_{\mu_{w/U} \times \mu^\circ_w}(w, w^0),
\end{aligned}
$$

where $V^0 = \{u^0 \mid u \in V\}$ and $T_w = \{(a, a/_{\mu^\circ_w}) \mid a \in \mathbb{A}_w\}$. The relations $P^\dagger[C, W'], P^{U\dagger}[C, W']$ can be represented by the pp-definitions

$$
\begin{aligned}
P^\dagger[C, W'](\mathbf{s}, V^0) &= R^\dagger(\mathbf{s}) \wedge \mathcal{S}^\dagger_{/W'}(V^0) \wedge \bigwedge_{w \in \mathbf{s}} T'_w(w, w^0),\\
P^{U\dagger}[C, W'](\mathbf{s}, V^0) &= R^\dagger/_{\overline{\mu}_{/U}}(\mathbf{s}) \wedge \mathcal{S}^\dagger_{/W'}(V^0) \wedge \bigwedge_{w \in \mathbf{s}} T'_w/_{\mu_{w/U} \times \mu^\circ_w}(w, w^0),
\end{aligned}
$$

where $T'_w = T_w \cap (B_w \times B_w/_{\mu^\circ_w})$.

(3) *$P[C, W']$, $C = \langle \mathbf{s}, R\rangle \in C$, is a subdirect product of R and $\mathcal{S}_{/W'}$; and $P^\dagger[C, W']$ is a subdirect product of $R^\dagger$ and $\mathcal{S}^\dagger_{/W'}$. For any $t \in \mathbf{s}, u \in V^0$, $\mathrm{pr}_{tu^0} P[C, W'](\mathbf{s}, V^0) = R^{tu}/_{\underline{0}_t \times \mu^\circ_u}$, $\mathrm{pr}_{tu^0} P^\dagger[C, W'](\mathbf{s}, V^0) = R^{tu\dagger}/_{\underline{0}_t \times \mu^\circ_u}$.*

(4) *$P^\dagger[C, W']$, $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$, is polynomially closed in $P[C, W']$, and $P^{U\dagger}[C, W']$ is polynomially closed in $P^U[C, W']$ for any non-central coherent set U.*

(5) *Let $S^\dagger[x, y, w, W']$ be given by*

$$S^\dagger[x, y, w, W'](x, y, w, V^1, V^2) = R^{xy\dagger}(x, y) \wedge P^\dagger[C^{xw}, W'](x, w, V^1) \wedge P^\dagger[C^{yw}, W'](x, y, V^2),$$

where $V^i = \{u^i \mid u \in V\}$, $i = 1, 2$, and $S'[x, y, w, W'] = S[x, y, w, W'] \cap \overline{B}$. Then $\mathrm{pr}_{tu} S[x, y, w, W'] = R^{tu}$, $\mathrm{pr}_{tu} S^\dagger[x, y, w, W'] = R^{tu\dagger}$ for any $t, u \in \{x, y, w\}$, $\mathrm{pr}_{V^i} S[x, y, w, W'] = \mathcal{S}_{/W'}$, $\mathrm{pr}_{V^i} S^\dagger[x, y, w, W'] = \mathcal{S}^\dagger_{/W'}$, $i = 1, 2$, $R^{tu}/_{\underline{0}_t \times \mu^\circ_u} \subseteq \mathrm{pr}_{tu^i} S[x, y, w, W']$, $R^{tu\dagger}/_{\underline{0}_t \times \mu^\circ_u} \subseteq \mathrm{pr}_{tu^i} S^\dagger[x, y, w, W']$ for any $t \in \{x, y, w\}$, $u \in V$, $i = 1, 2$; and $\mathsf{umax}(S^\dagger[x, y, w, W']) \subseteq \mathsf{umax}(S'[x, y, w, W'])$.

(6) *In the notation of item (5) $S^\dagger[x, y, w, W']$ is polynomially closed in $S[x, y, w, W']$.*

(7) *Let $S^\dagger[W', U]$ be given by*

$$S^\dagger[W', U](V, V^1, \ldots, V^\ell) = \mathcal{S}^\dagger_{/U}(V) \wedge \bigwedge_{j=1}^{\ell} P^{U\dagger}[C_j, W'](\mathbf{s}_j, V^j),$$

where $C_j = \langle \mathbf{s}_j, R_j\rangle$, and $V^i = \{u^i \mid u \in V\}$, $i \in [\ell]$, are distinct variables not in V, and let $S'[W', U] = S[W', U] \cap \overline{B}$. Then $S[W', U]$, $S^\dagger[W', U]$ are subdirect products of $\mathcal{S}_{/U}$ and ℓ copies of $\mathcal{S}_{/W'}$ (respectively, of $\mathcal{S}^\dagger_{/U}$ and ℓ copies of $\mathcal{S}^\dagger_{/W'}$), $R^{tu}/_{\mu_{t/U} \times \mu^\circ_u} \subseteq \mathrm{pr}_{tu^i} S[W', U]$, $R^{tu\dagger}/_{\mu_{t/U} \times \mu^\circ_u} \subseteq \mathrm{pr}_{tu^i} S^\dagger[W', U]$ for any $t, u \in V$, $i \in [\ell]$; and $\mathsf{umax}(S^\dagger[W', U]) \subseteq \mathsf{umax}(S^U[W', U])$.

(8) *In the notation of item (7) $S^\dagger[W', U]$ is polynomially closed in $S[W', U]$.*

PROOF. (1) follows from the block-minimality of $\mathcal{P}$ and (S3), (S5) for $\mathcal{P}^\dagger$ by Lemma 6.14(5), because $\mathcal{S}_{/U}, \mathcal{S}^\dagger_{/U}$ are obtained through the same conjunctive definition represented by the instance $\mathcal{P}$.

Item (2) follows from the definition of $P[C, W'], P^U[C, W'], P^\dagger[C, W'], P^{U\dagger}[C, W']$. For $P[C, W']$ item (3) follows from the block-minimality of $\mathcal{P}$, as $\mathcal{P}_{/W'}$ is minimal. For $P^\dagger[C, W']$ it follows immediately from (S3) for $\mathcal{P}^\dagger$ (recall that the empty set of coordinates is also non-central).

To prove item (4) observe that $T'_w = T_w \cap (B_w \times B_w/_{\mu^\circ_w})$, is polynomially closed in T_w, by Lemma 6.14(2,3), as it is a block of $\beta_w \times \beta_w/_{\mu^\circ_w}$ of T_w. Similarly, $T'_w/_{\mu_{w/U} \times \mu^\circ_w}$ is polynomially closed in $T_w/_{\mu_{w/U} \times \mu^\circ_w}$. Therefore $P^\dagger[C, W'], P^{U\dagger}[C, W']$ are polynomially closed in $P[C, W'], P^U[C, W']$, respectively, by Lemma 6.14(5).

(5) The first part of item (5) follows from (S3) for $\mathcal{P}^\dagger$ and item (3). For the second part, by Lemma 6.4(3) it suffices to show that $S^\dagger[x, y, w, W']$ contains an element u-maximal in $S'[x, y, w, W']$. Take any $\mathbf{a} \in S^\dagger[x, y, w, W']$, and let $\mathbf{b} \in S'[x, y, w, W']$ be an element maximal in $S'[x, y, w, W']$ and such that $\mathbf{a} \sqsubseteq \mathbf{b}$ in $S'[x, y, w, W']$. Then Lemma 6.14(4,6) implies that $S^\dagger[x, y, w, W']$ is s-closed in $S'[x, y, w, W']$, because $R^\dagger$ is s-closed in R for every constraint relation R by (S6) for $\mathcal{P}^\dagger$, and since T'_w is trivially s-closed in $T_w \cap (B_w \times B_w/_{\mu^\circ_w})$ for $w \in \mathbf{s}$. We obtain $\mathbf{b} \in S^\dagger[x, y, w, W']$, and, as $\mathbf{b}$ is maximal in $S'[x, y, w, W']$, it is also u-maximal in $S'[x, y, w, W']$.

(6) follows from (S5) for $\mathcal{P}^\dagger$, item (4), and Lemma 6.14(5).

(7) The first part of item (7) follows from (S3) for $\mathcal{P}^\dagger$ and item (3). For the second part, by Lemma 6.4(3) it suffices to show that $S^\dagger[W', U]$ contains an element u-maximal in $S'[W', U]$. Take any $\mathbf{a} \in S^\dagger[W', U]$, and let $\mathbf{b} \in S'[W', U]$ be an element maximal in $S'[W', U]$ and such that $\mathbf{a} \sqsubseteq \mathbf{b}$ in $S'[W', U]$. Then, as similar to item (5) $S^\dagger[W', U]$ is s-closed in $S'[W', U]$, we obtain $\mathbf{b} \in S^\dagger[W', U]$. Also, as $\mathbf{b}$ is maximal in $S'[W', U]$, it is also u-maximal in $S'[W', U]$.

(8) follows from (S5) for $\mathcal{P}^\dagger$ and Lemma 6.14(5). □

8.2 Tightening the instance and induced congruences

Let $\mathcal{P}^{\ddagger} = (V, C^{\ddagger})$ be the following instance.

(R1) For every $C^{\dagger} = \langle \mathbf{s}, R^{\dagger}\rangle \in C^{\dagger}$, the set $R'^{\ddagger}$ includes every $\mathbf{a} \in \mathsf{umax}(R^{\dagger})$ such that $\mathbf{a}/_{\overline{\mu}^{\circ}}$ extends to a solution $\varphi \in \mathsf{umax}(\mathcal{S}^{\dagger}_{/W'})$ with $\varphi(v) \in B/_{\overline{\mu}^{\circ}}$;

(R2) for every $C^{\dagger} = \langle \mathbf{s}, R^{\dagger}\rangle \in C^{\dagger}$, there is $C^{\ddagger} = \langle \mathbf{s}, R^{\ddagger}\rangle$, where $R^{\ddagger} = \mathrm{Sg}_R(R'^{\ddagger})$.

The following two statements show how relations $R^{\ddagger}$ are related to $R^{\dagger}$. They amount to saying that either $R^{\ddagger}$ is (almost) the intersection of $R^{\dagger}$ with a block of a congruence of R, or $\mathsf{umax}(R^{\ddagger}) = \mathsf{umax}(R^{\dagger})$. Recall that for congruences β_w, $w \in V$, and $U \subseteq V$ by $\overline{\beta}_U$ we denote the product $\prod_{w\in U} \beta_w$.

Lemma 8.2. *Let $C = \langle \mathbf{s}, R\rangle \in C$, and let $\mathcal{S}_{/W'}, \mathcal{S}^{\dagger}_{/W'}$ be the set of solutions of $\mathcal{P}_{/W'}$ (respectively, $\mathcal{P}^{\dagger}_{/W'}$). Then there is a congruence τ_C of R satisfying the following conditions.*

(a) *Either* $\mathsf{umax}(R^{\ddagger}) = \mathsf{umax}(R^{\dagger})$*, or for a τ_C-block T it holds that $R^{\ddagger} \subseteq R^{\dagger} \cap T$ and* $\mathsf{umax}(R^{\ddagger}) = \mathsf{umax}(R^{\dagger} \cap T) \subseteq \mathsf{umax}(R^{\dagger}) \cap T$.

(b) *Either $\tau_C|_{R^{\dagger}} = \overline{\beta}_{\mathbf{s}}|_{R^{\dagger}}$, or $R^{\dagger}/_{\tau_C}$ is isomorphic to $R^{v\dagger}/_{\alpha}$. Moreover, in the latter case $\tau_C \prec \gamma_C \le \overline{\beta}_{\mathbf{s}}$, where $\gamma_C = \mathrm{Cg}(R^{\dagger}) \vee \tau_C$, the smallest congruence of R greater than τ_C for which $R^{\dagger}$ is a subset of a γ_C-block.*

If, according to item (b) of the lemma, $\tau_C|_{R^{\dagger}} = \overline{\beta}_{\mathbf{s}}|_{R^{\dagger}}$, we say that τ_C is the *full congruence*; if the latter option of item (b) holds we say that τ_C is a *maximal congruence*.

Proof. We apply the Congruence Lemma 6.15 to the relations $P = P[C, W'](\mathbf{s}, V^0)$ and $P^{\dagger} = P^{\dagger}[C, W'](\mathbf{s}, V^0)$. The relation $P[C, W']$ is chained with respect to $\overline{\beta}, \overline{B}$ by (S4) and $P^{\dagger}[C, W']$ is polynomially closed in $P[C, W']$. The parameters of the Congruence Lemma are then set as follows: $I_0 = \{v^0\}$, $\mathbb{A}_w = \mathbb{A}_w$, $\beta_w = \beta_w$ for $w \in \mathbf{s}$ and $\mathbb{A}_{w^0} = \mathbb{A}_w/_{\mu^{\circ}_w}$, $\beta_{w^0} = \beta_w$, $w \in V$, $\alpha = \alpha$, $R' = P \cap \overline{B}$, $Q = P^{\dagger}$, $I = \mathbf{s}$, $E_1 = R^{v\dagger}/_{\mu^{\circ}_v}$, $E^C_2 = R^{\dagger}$. As $R^{v\dagger}$ is weakly as-closed in B_v by (S6), if $a \in R^{v\dagger}$ then the as-component C of B_v such that $a \in C$ belongs to $R^{v\dagger}$. Moreover, note that since $\mathsf{typ}(\alpha, \beta_v) = \mathbf{2}$, the algebra $B_v/_{\alpha}$ is a module. Therefore $E_1/_{\alpha}$ is a submodule of $B_v/_{\alpha}$, and so $C/_{\alpha} = B_v/_{\alpha}$ for any as-component C of B_v. Thus, $E_1/_{\alpha} = B_v/_{\alpha}$.

Let $Q' = \mathrm{pr}_{\mathbf{s}\cup\{v^0\}} P^{\dagger}$. If the first option of the Congruence Lemma 6.15 holds, set $\tau_C = \overline{\beta}_{\mathbf{s}}$. If the second option is the case, set τ_C to be the congruence η of $\mathrm{pr}_{\mathbf{s}} Q'$ identified in the Congruence Lemma 6.15. Note that in the latter case the restriction of τ_C on $R^{\dagger}$ is nontrivial, because tuples from a τ_C-block are related in Q' only to elements from one α-block, while the domain of v in Q' spans all the α-blocks in B_v, as we observed above.

(a) If the first option of the Congruence Lemma 6.15 holds for $P^{\dagger}$ then $\mathsf{umax}(R^{\ddagger}) = \mathsf{umax}(R^{\dagger})$ by that lemma. If the second option holds then by the Congruence Lemma 6.15 the relation $Q'/_{\alpha}$ is the graph of a mapping $\varphi : R^{\dagger} \to B_v/_{\alpha}$ whose kernel is the restriction of τ_C to $R^{\dagger}$. This means there is a τ_C-block T such that $R^{\ddagger} \subseteq R^{\dagger} \cap T$. Moreover, as $R^{\ddagger}$ contains all tuples u-maximal in $R^{\dagger}$ belonging to T, the result follows.

(b) If $\tau_C \ne \overline{\beta}_{\mathbf{s}}$, by construction $R^{\dagger}/_{\tau_C}$ is isomorphic to $\mathrm{pr}_{v^0} P^{\dagger}/_{\alpha}$, which is isomorphic to $R^{v\dagger}/_{\alpha}$.

It remains to show that $\tau_C \prec \gamma_C$. Let γ be the congruence of R constructed in the Congruence Lemma 6.15, that is, $\gamma = \mathrm{Cg}_R(\mathsf{umax}(R^{\dagger})) \vee \tau_C$ in our case. Lemma 6.15 states that $\tau_C \prec \gamma$. It therefore suffices to observe that $\gamma_C = \gamma$. It follows from the fact that $\mathsf{typ}(\alpha, \beta_v) = \mathbf{2}$, and therefore $R^{\dagger}/_{\tau_C} = (\mathsf{umax}(R^{\dagger}))/_{\tau_C}$. □

Next we identify variables $w \in V$ for which β'_w has to be different from β_w. Since $\mathcal{P}$ is (2,3)-minimal, for every $w \in V$ there is $C^w = \langle (w), R^w\rangle \in C$. For $w \in V$ there are two cases. In the first case, when τ_{C^w} is the full congruence, we set $\beta'_w = \beta_w$. Otherwise τ_{C^w} is a congruence of $\mathbb{A}_w$ with $\tau_{C^w} \prec \beta_w$ in $\mathrm{Con}(\mathbb{A}_w)$. Set $\beta'_w = \tau_{C^w}$. If $\beta'_w \ne \beta_w$ then there is a

β'_w-block B'_w such that $b \in B'_w$ whenever $(a, b) \in R^{vw\dagger}$ and $a \in B$. Let V' be the set of such variables. For the remaining variables w we set $B'_w = B_w$.

LEMMA 8.3. *In the notation above*

(1) *Let $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and let $\gamma, \delta \in \mathsf{Con}(\mathbb{A}_u)$, $u \in U = \mathbf{s} \cap W$ be such that $(u, \gamma, \delta) \in \mathcal{W}$ and $(\alpha, \beta_v), (\gamma, \delta)$ cannot be separated from each other. Then if τ_C is a maximal congruence, for any polynomial f of R, $f(\gamma_C) \subseteq \tau_C$ if and only if $f(\delta) \subseteq \gamma$. If γ, δ are considered as congruences of R, this condition means that (τ_C, γ_C) and (γ, δ) cannot be separated.*

(2) *Assuming $\mathsf{MAX}(\mathcal{P}) \cap \mathsf{Center}(\mathcal{P}) = \varnothing$, for every $z \in V$ either $\mu^\circ_z = \underline{0}_z$, or $(\underline{0}_z, \mu^\circ_z)$ can be separated from (α, β_v) in $\mathcal{P}$ or (α, β_v) can be separated from $(\underline{0}_z, \mu^\circ_z)$ in $\mathcal{P}$.*

(3) *Let $V^i = \{u^i \mid u \in V\}$, $i = 1, 2$, be the coordinates of the relation $S[x, y, w, W']$. For each of $t \in \{x, y, w\}$ and each $i = 1, 2$, either $\mu^\circ_t = \underline{0}_t$, or the interval (α, β_v) (in the v^i-coordinate) can be separated in $S[x, y, w, W']$ from $(\underline{0}_t, \mu^\circ_t)$ or the other way round.*

(4) *Let $\mathcal{C} = \{C_1, \dots, C_\ell\}$ and let $V^i = \{u^i \mid u \in V\}$, $i \in [\ell]$, be the coordinates of the relation $S[W', U]$. For each $x \in V$ with $\mu_{x/U} = \underline{0}_x$ and each $i \in [\ell]$, either $\mu^\circ_x = \underline{0}_x$, or the interval (α, β_v) (in the v^i-coordinate) can be separated in $S[W', U]$ from $(\underline{0}_x, \mu^\circ_x)$ or the other way round.*

PROOF. (1) Let $\mathcal{S}_{/W'}$ be defined as in Lemma 8.2 and τ_C a maximal congruence. Take a polynomial f of R. Since $\mathcal{P}$ is a block-minimal instance, the polynomial f can be extended from a polynomial of R to a polynomial of $\mathcal{S}_{/W'}$; we keep notation f for this polynomial. Since τ_C is maximal, by the Congruence Lemma 6.15 the intervals (α, β_v) and (τ_C, γ_C) in the congruence lattices of $\mathbb{A}_v/_{\mu^\circ_v}$ and $R/_{\overline{\mu}^\circ}$, respectively, cannot be separated in $\mathcal{S}_{/W'}$. Therefore $f(\beta_v) \subseteq \alpha$ if and only if $f(\gamma_C) \subseteq \tau_C$. Since (α, β_v) and (γ, δ) cannot be separated in $\mathcal{P}$, the first inclusion holds if and only if $f(\delta) \subseteq \gamma$, and we infer the result.

(2) If $(v, \alpha, \beta_v) \in \mathcal{W}'$, the centralizer $(\alpha : \beta_v) = \underline{1}_v$. On the other hand, if $w \in \mathsf{MAX}(\mathcal{P})$, then $w \notin \mathsf{Center}(\mathcal{P})$ and $(\underline{0}_w : \mu_w) \neq \underline{1}_w$. Therefore (α, β_v) can be separated from $(\underline{0}_w, \mu_w)$ or the other way round, as it follows from Lemma 7.1. If $(v, \alpha, \beta_v) \notin \mathcal{W}'$ then if $(\underline{0}_z, \mu_z)$ and (α, β_v) cannot be separated then $z \in W$ and $\mu^\circ_z = \underline{0}_z$ by construction.

(3) Let $S = S[x, y, w, W']$. By Lemma 8.1(5) we have $R^{xv}/_{\underline{0}_x \times \mu^\circ_v} \subseteq \mathrm{pr}_{xv^i} S$, $R^{yv}/_{\underline{0}_y \times \mu^\circ_v} \subseteq \mathrm{pr}_{yv^i} S$, $R^{xv}/_{\underline{0}_z \times \mu^\circ_v} \subseteq \mathrm{pr}_{wv^i} S$, $i = 1, 2$. The result follows from item (2) and Lemma 6.9(1).

(4) Let $S = S[W', U]$. Then by Lemma 8.1(7) we have $R^{xv}/_{\overline{\mu}_{x/U} \times \mu^\circ_v} \subseteq \mathrm{pr}_{xv^i} S$. The result follows from item (2) and Lemma 6.9(1). □

Now we are in a position to prove that $\mathcal{P}^\ddagger$ is a $(\overline{\beta}', \overline{B}')$-compressed instance.

THEOREM 8.4. *In the notation above, $\mathcal{P}^\ddagger$ is a $(\overline{\beta}', \overline{B}')$-compressed instance.*

8.3 Conditions (S1), and (S4)–(S6)

We start with conditions (S1), and (S4)–(S6).

We start with condition (S1).

LEMMA 8.5. *Condition (S1) holds for $\mathcal{P}^\ddagger$.*

PROOF. Let $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$. That $R^\ddagger \neq \varnothing$ is straightforward by construction, item (R2). If $w \in \mathbf{s} \cap V'$ then a tuple $\mathbf{a} \in \mathsf{umax}(R^\dagger)$ belongs to $R^\ddagger$ if and only if $\mathbf{a} \in \varphi(\mathbf{s})$ for some $\varphi \in \mathcal{S}^\dagger_{/W'}$ with $\varphi(v) \in B/_{\mu^\circ_v}$. The latter implies that $\mathbf{a}[w] \in R^{w\ddagger}$, and so $\mathsf{umax}(\mathrm{pr}_w R^\ddagger) \subseteq \mathsf{umax}(R^{w\ddagger})$. On the other hand, if $a \in \mathsf{umax}(R^{w\ddagger}) \subseteq \mathsf{umax}(R^{w\dagger})$ and

$\mathbf{a} \in \mathsf{umax}(R^\dagger)$ with $\mathbf{a}[w] = a$ (such $\mathbf{a}$ exists by (S1) for $\mathcal{P}^\dagger$), then there is $\varphi \in \mathcal{S}^\dagger_{W'}$ such that $\mathbf{a} \in \varphi(\mathbf{s})$. As $a \in R^{w\ddagger}$, it holds that $\varphi(v) \in B/_{\mu^\circ_v}$, and so $\mathbf{a} \in R^\ddagger$ implying $\mathsf{umax}(R^{w\ddagger}) \subseteq \mathsf{umax}(\mathrm{pr}_w R^\ddagger)$.

Now let $w \in \mathbf{s} - V'$. We prove that $\mathsf{umax}(\mathrm{pr}_w R^\ddagger) = \mathsf{umax}(\mathrm{pr}_w R^\dagger)$. Suppose first that $\mu^\circ_w = \underline{0}_w$. In this case for any $a \in \mathsf{umax}(R^{w\dagger})$ there is $b \in B$ with $(a, b) \in R^{wv\dagger}$, and by (S3) a solution $\varphi \in \mathcal{S}^\dagger_{/W}$ such that $\varphi(w) = a, b \in \varphi(v)$. Let $\mathbf{a} \in \varphi(\mathbf{s}) \cap \mathsf{umax}(R^\dagger)$. We have $\mathbf{a} \in R^\ddagger$ and $\mathbf{a}[w] = a$.

Suppose that $\mu^\circ_w = \mu_w$, and let $D = \mathrm{pr}_w R^\dagger$. By Lemma 8.3(1,2) (τ_C, γ_C) can be separated from $(\underline{0}_w, \mu_w)$ or the other way round. As R is chained, by Lemma 6.12(d), for any $\mathbf{a}, \mathbf{b} \in R^\dagger$ there is a polynomial f of R such that $f(\mathbf{a}) \in \mathbf{a}/_{\tau_C}, f(\mathbf{b}) \in \mathbf{b}/_{\tau_C}$, and $f(\mu_w|_D) \subseteq \underline{0}_w$. Let $Q = \{(\mathbf{a}/_{\tau_C}, \mathbf{a}[w]) \mid \mathbf{a} \in R^\dagger\} \subseteq R^\dagger/_{\tau_C} \times D$. By the observation above Q is linked, and by Proposition 6.8(2) $R^\dagger/_{\tau_C} \times \mathsf{umax}(D) \subseteq Q$. Therefore for every τ_C-block T and any $a \in \mathsf{umax}(D)$ there exists $\mathbf{a} \in T$ with $\mathbf{a}[w] = a$. □

Since $B_v/_\alpha$ is a module, and therefore is a nontrivial as-component, Lemma 6.10(5) applied to every variable $w \in V'$ immediately implies that condition (S4) for $\mathcal{P}^\ddagger$ holds for constraint relations and relations of the form $\mathcal{S}_{/U}$. For compound relations the result follows from Lemma 6.10(5), Lemma 7.4, and properties (S2),(S3) for $\mathcal{P}^\ddagger$ that will be proved in Sections 8.4, 8.5. Indeed, by Lemma 7.4 for $\mathcal{P}^\dagger$ and $\mathcal{P}^\ddagger$ the u-maximal elements of the domains of the compound relations are the same as those for $R^\dagger, R^\ddagger$, which means that we can apply the same restrictions of the domains required in Lemma 6.10(5) as for $R^\dagger, R^\ddagger$. Condition (S5) is not too difficult as well.

LEMMA 8.6. *Condition (S5) for $\mathcal{P}^\ddagger$ holds. That is, for every $\langle \mathbf{s}, R\rangle \in \mathcal{C}$ and every non-central coherent set U the relations $R^\ddagger$ and $R^\ddagger/_{\overline{\mu}_{/U}}$ are polynomially closed in R and $R/_{\overline{\mu}_{/U}}$, respectively.*

PROOF. Let $\mathsf{Cg}(R^\dagger), \mathsf{Cg}(R^\ddagger)$ and $\mathsf{Block}(R^\dagger), \mathsf{Block}(R^\ddagger)$ be the congruences and blocks of those congruences identified in the condition of polynomial closedness. Let also $\mathbf{a} \in \mathsf{umax}(R^\ddagger)$ and $\mathbf{b} \in \mathsf{Block}(R^\ddagger)$ satisfy the conditions of polynomial closedness. Since $\mathsf{umax}(R^\ddagger) \subseteq \mathsf{umax}(R^\dagger)$, and since $\mathsf{Cg}(R^\ddagger) \le \mathsf{Cg}(R^\dagger)$, by (S5) for $\mathcal{P}^\dagger$, $\mathbf{b} \in R^\dagger$. Let T be the τ_C-block containing $R^\ddagger$. As $\mathbf{a} \in \mathsf{umax}(R^\ddagger)$ and $\mathbf{a} \sqsubseteq^{as} \mathbf{b}$, by Lemma 8.2(a) it suffices to show that $\mathbf{b} \in T$. However, this is straightforward, because $R^\ddagger \subseteq T$ by construction, and so $\mathsf{Block}(R^\ddagger) \subseteq T$.

Now, let $\mathsf{Cg}(R^\ddagger/_{\overline{\mu}_{/U}})$ and $T^\mu = \mathsf{Block}(R^\ddagger/_{\overline{\mu}_{/U}})$ be the congruence and its block identified in the condition of polynomial closedness for $R^\ddagger/_{\overline{\mu}_{/U}}$. Let also $\overline{\mathbf{a}} \in \mathsf{umax}(R^\ddagger/_{\overline{\mu}_{/U}})$, $\overline{\mathbf{b}} \in T^\mu$ satisfy the conditions of polynomial closedness. Similar to the argument above, since $\mathsf{umax}(R^\ddagger/_{\overline{\mu}_{/U}}) \subseteq \mathsf{umax}(R^\dagger/_{\overline{\mu}_{/U}})$, and, since $\mathsf{Cg}(R^\ddagger/_{\overline{\mu}_{/U}}) \le \mathsf{Cg}(R^\dagger/_{\overline{\mu}_{/U}})$, by (S5) for $\mathcal{P}^\dagger$, $\overline{\mathbf{b}} \in R^\dagger/_{\overline{\mu}_{/U}}$. This means, in particular, that there is $\mathbf{b} \in \overline{\mathbf{b}} \cap \mathsf{umax}(R^\dagger)$.

Let $\eta = \tau_C \vee \overline{\mu}_{/U}$. Then, since $\tau_C \prec \gamma_C$, $\eta \wedge \gamma_C = \tau_C$ or $\eta \wedge \gamma_C = \gamma_C$. In the former case by the choice of $\mathbf{b}$, it belongs to $R^\dagger \cap S$, where $S = \{\mathbf{c} \in R \mid \mathbf{c}/_{\overline{\mu}_{/U}} \in \mathsf{Block}(R^\ddagger/_{\overline{\mu}_{/U}})\}$, i.e., S is a subset of an η-block. Therefore for any $\mathbf{c} \in \mathsf{umax}(R^\ddagger)$ we have $(\mathbf{b}, \mathbf{c}) \in \gamma_C$ and $(\mathbf{b}, \mathbf{c}) \in \eta$, implying $\mathbf{b} \in T$. Thus, by (R1) $\mathbf{b} \in R^\ddagger$, hence, $\overline{\mathbf{b}} \in R^\ddagger/_{\overline{\mu}_{/U}}$. If $\overline{\mu}_{/U} \le \gamma_C$, consider $\theta = \tau_C \wedge \overline{\mu}_{/U}$. If $\theta = \tau_C$, that is, $\overline{\mu}_{/U} \le \tau_C$, we complete the proof the same way as for $R^\ddagger$. Otherwise consider $R^\dagger/_\theta$ as a subdirect product P of $R^\dagger/_{\tau_C} \times R^\dagger/_{\overline{\mu}_{/U}}$. Since $R^\dagger/_{\tau_C}$ is a module, it is an as-maximal component. Hence, as $\tau_C \vee \overline{\mu}_{/U} = \gamma_C$ the relation P is linked. By Proposition 6.8(2) $R^\dagger/_{\tau_C} \times \mathsf{umax}(R^\dagger/_{\overline{\mu}_{/U}}) \subseteq P$. Let T be the τ_C-block containing $R^\ddagger$. Then as $\overline{\mathbf{b}} \in \mathsf{umax}(R^\dagger/_{\overline{\mu}_{/U}})$, there is $\mathbf{c} \in \overline{\mathbf{b}} \cap T \cap R^\dagger$. By (R1) $\mathbf{c} \in R^\ddagger$ and therefore $\overline{\mathbf{b}} \in R^\ddagger/_{\overline{\mu}_{/U}}$. □

Finally, condition (S6) also holds.

LEMMA 8.7. *Condition (S6) for $\mathcal{P}^\ddagger$ holds.*

PROOF. Let $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$. By Lemma 8.2(a) either $\mathsf{umax}(R^\ddagger) = \mathsf{umax}(R^\dagger)$, in which case we are done, or $\mathsf{umax}(R^\ddagger) \subseteq \mathsf{umax}(R^\dagger) \cap T$, where T is a τ_C-block. For $w \in \mathbf{s} - V'$ we have $\mathsf{umax}(\mathrm{pr}_w R^\ddagger) = \mathsf{umax}(\mathrm{pr}_w R^\dagger)$ and the property of weak as-closedness holds for such variables. Otherwise if $\mathbf{s}\cap V' \neq \emptyset$, $\mathsf{umax}(R^\ddagger) \subseteq \mathsf{umax}(R^\dagger)\cap\overline{B}'$. Moreover by (S2) for any $\mathbf{a} \in R^\dagger$ and any $w, u \in \mathbf{s} \cap V'$ it holds that $\mathbf{a}[w] \in B'_w$ if and only if $\mathbf{a}[u] \in B'_u$. Let $a \in \mathsf{umax}(\mathrm{pr}_w R^\ddagger) \subseteq \mathsf{umax}(\mathrm{pr}_w R^\dagger)$ and $b \in \mathrm{pr}_w(R\cap\overline{B}')$ such that $a \sqsubseteq^{as} b$ in $\mathrm{pr}_w(R\cap\overline{B}')$. There is $\mathbf{a} \in \mathsf{umax}(R^\ddagger) = \mathsf{umax}(R'^\ddagger) \subseteq \mathsf{umax}(R^\dagger)$ with $\mathbf{a}[w] = a$. By (S6) for $R^\dagger$ there is also $\mathbf{b} \in \mathsf{umax}(R^\dagger)$ such that $\mathbf{b}[w] = b$ and $\mathbf{a} \sqsubseteq^{as} \mathbf{b}$ in $R^\dagger$. Then, for any solution $\varphi \in \mathcal{S}^\dagger_{/W'}$ with $\varphi(w) = b/_{\mu^\circ_w}$, we have $\varphi(v) \in B/_{\mu^\circ_v}$, and therefore $\mathbf{b} \in R'^\ddagger \subseteq R^\ddagger$.

Finally, for s-closedness, let $\mathbf{a} \in \mathsf{umax}(R^\ddagger)$ and $\mathbf{b} \in R \cap \overline{B}$ with $\mathbf{a} \leq \mathbf{b}$. Then by (S6) for $\mathcal{P}^\dagger$ we have $\mathbf{b} \in \mathsf{umax}(R^\dagger)$. Also, as $R^\dagger/_{\tau_C}$ is a module, $\mathbf{b} \in T$. Then as before we conclude that $\mathbf{b} \in R^\ddagger$. □

8.4 Condition (S2)

Property (S2) is more difficult to prove.

Lemma 8.8. *Condition (S2) for $\mathcal{P}^\ddagger$ holds. That is, the relations $R^{X\ddagger}$, where $R^{X\ddagger}$ is obtained from $R^{X\dagger}$ as described in (R1),(R2) for $X \subseteq V$, $|X| \leq 2$, form a nonempty $(2,3)$-strategy for $\mathcal{P}^\ddagger$.*

Proof. By (S2) for $\mathcal{P}^\dagger$ the relations $R^{X\dagger}$, $X \subseteq V$, $|X| \leq 2$, constitute a $(2,3)$-strategy for $\mathcal{P}^\dagger$. Let $x, y \in V$. As $R^{xy\ddagger}$ is generated by $R'^{xy\ddagger}$, it suffices to show that for any tuple $(a,b) \in R'^{xy\ddagger}$ and any $w \notin \{x,y\}$ there is $c \in \mathbb{A}_w$ such that $(a,c) \in R^{xw\ddagger}$, $(b,c) \in R^{yw\ddagger}$. By (R1) $R'^{xy\ddagger} \subseteq \mathsf{umax}(R^{xy\dagger})$ and so by (S2) for $\mathcal{P}^\dagger$ and the Maximality Lemma 6.5(2) applied to the relation $R^{xy\dagger}(x,y) \wedge R^{xw\dagger}(x,w) \wedge R^{yw\dagger}(y,w)$ there is $d \in \mathbb{A}_w$ such that $(a,d) \in \mathsf{umax}(R^{xw\dagger})$, $(b,d) \in \mathsf{umax}(R^{yw\dagger})$.

Let $Q_x = P[C^{xw}, W'](x, w, V^x)$, $Q_y = P[C^{yw}, W'](y, w, V^y)$, where $V^x = \{u^x \mid u \in V\}$, $V^y = \{u^y \mid u \in V\}$, and let $S = S[x,y,w,W'](x,y,w,V^x,V^y)$, $S' = S \cap \overline{B}$, and let $S^\dagger$ be defined as in Lemma 8.1(5). It suffices to show that for some $c \in R^{w\dagger}$ and $\varphi^x, \varphi^y \in \mathcal{S}^\dagger_{/W'}$ with $\varphi^x(v), \varphi^y(v) \in B/_{\mu^\circ_v}$, such that $(a,c) \in \mathsf{umax}(R^{xw\dagger})$ and $(b,c) \in \mathsf{umax}(R^{yw\dagger})$ it holds that $(a,b,c,\varphi^x,\varphi^y) \in S$. Indeed, by the definition of Q_x, Q_y and (R1) it means that $(a,c) \in R'^{xw\ddagger} \subseteq R^{xw\ddagger}$ and $(b,c) \in R'^{yw\ddagger} \subseteq R^{yw\ddagger}$. As we observed above there is $d \in R^{w\dagger}$ such that $(a,d) \in R^{xw\dagger}$, $(b,d) \in R^{yw\dagger}$, and the triple (a,b,d) extends to a tuple from $S^\dagger$.

Observe that $(a,b) \in \mathsf{umax}(\mathrm{pr}_{xy} S^\dagger) = \mathsf{umax}(R^{xy\dagger})$ by (S3) and Lemma 8.1(5). Therefore there is $d \in R^{w\dagger}$ such that $(a,b,d) \in \mathsf{umax}(\mathrm{pr}_{xyw} S^\dagger)$, that is $\mathbf{a} = (a,b,d,\varphi^x,\varphi^y) \in S^\dagger$ for some $\varphi^x, \varphi^y \in \mathcal{S}^\dagger_{/W'}$. Moreover, as $\mathrm{pr}_{xw} S^\dagger = R^{xw\dagger}$, $\mathrm{pr}_{yw} S^\dagger = R^{yw\dagger}$ and $(a,d) \in \mathsf{umax}(R^{xw\dagger})$, $(b,d) \in \mathsf{umax}(R^{yw\dagger})$, the mappings φ^x, φ^y can be chosen such that $\mathbf{a} \in \mathsf{umax}(S^\dagger) \subseteq \mathsf{umax}(S')$, where the last inclusion is by Lemma 8.1(5).

Next, we prove that there is $d_0 \in R^{w\dagger}$ such that $(a,b,d_0,\varphi^x_0,\varphi^y_0) \in S^\dagger$ for some $\varphi^x_0, \varphi^y_0 \in \mathcal{S}^\dagger_{/W'}$, $\varphi^x_0(v) \in B/_{\mu^\circ_v}$. By (S4) and Lemmas 7.4,8.1(6) the relations S, S', and $S^\dagger$ satisfy the conditions of the Congruence Lemma 6.15 for $I = \{x,y\}$ and $I_0 = \{v^x\}$. Therefore, either $\mathsf{umax}(\mathrm{pr}_{xy} S^\dagger) \times R^{v\dagger}/_\alpha \subseteq \mathrm{pr}_{xyv^x} S^\dagger/_\alpha$ (in which case we have $(a,b,e^x) \in \mathrm{pr}_{xyv^x} S^\dagger$ for some $e^x \in B$ and the result follows), or there is a congruence η of $\mathrm{pr}_{xy} S$ such that $\mathrm{pr}_{xyv^x} S^\dagger/_\alpha$ is the graph of a mapping $\psi : \mathrm{pr}_{xy} S^\dagger \to R^{v\dagger}/_\alpha$ and $\eta|_{\mathrm{pr}_{xy} S^\dagger}$ is the kernel of ψ. If $\psi(a,b) = B$, $(a,b,e^x) \in \mathrm{pr}_{xyv^x} S^\dagger$ for some $e^x \in B$ implying the desired result.

It will be more convenient now to consider relations $\mathcal{S}^\dagger_{/W'}/_\alpha, Q^x, S/_\alpha, S'/_\alpha, S^\dagger/_\alpha$, where α denotes the product congruence of the algebra which is the equality relation in every factor except for v in the case $\mathcal{S}^\dagger_{/W'}$ and v^x in the remaining cases. Note that $\mathrm{pr}_{xyv^x} S^\dagger/_\alpha = (\mathrm{pr}_{xyv^x} S^\dagger)/_\alpha$ is still the graph of the mapping ψ.

Suppose that $\psi(a,b) = D \neq B$, we show that this leads to a contradiction. Let $\mathbf{a} = (a,b,d,\varphi^x,\varphi^y) \in S^\dagger$ be as above, and $\mathbf{a}/_\alpha$ the corresponding tuple in $S^\dagger/_\alpha$. Let

$$S''/_\alpha = \{\mathbf{b} \in S'/_\alpha \mid \mathbf{b}[v^x] = D\},$$

and let $\mathbf{a}' = (a', b', d') \in \mathsf{umax}(\mathrm{pr}_{xyw} S''/_\alpha)$ be such that $\mathrm{pr}_{xyw}\mathbf{a} \sqsubseteq \mathbf{a}'$ in $\mathrm{pr}_{xyw} S''/_\alpha$. By s-closedness $\mathbf{a}' \in \mathrm{pr}_{xyw} S^\dagger = \mathrm{pr}_{xyw} S^\dagger/_\alpha$ and we can find $\mathbf{a}''_\alpha = (a', b', d', \varphi_1^x, \varphi_1^y) \in \mathsf{umax}(S''/_\alpha) \cap S^\dagger/_\alpha$ for some φ_1^x, φ_1^y. As $R^{xy\dagger}/_\eta$ is a module, (a', b') is in the same η-block with (a, b). By (R1) there is a solution φ of $\mathcal{P}^\dagger_{/W'}$ such that $a' \in \varphi(x)$, $b' \in \varphi(y)$, and $\varphi(v) \in B$. In other words, there are $(a'', c') \in R^{xw\ddagger}$ and $(b'', c'') \in R^{yw\ddagger}$ with $a'' \overset{\mu_x^\circ}{\equiv} a'$, $b'' \overset{\mu_y^\circ}{\equiv} b'$, and $c' \overset{\mu_w^\circ}{\equiv} c''$. This also means that $(a'', c', \varphi(v)) \in \mathrm{pr}_{xwv^x} Q_x$ and $(b'', c'', \varphi(v)) \in \mathrm{pr}_{ywv^y} Q_y$.

By the definition of the congruences μ_z°, $z \in \{x, y, w\}$ and Lemma 8.3(3) the interval (α, β_v) can be separated from $(\underline{0}_z, \mu_z^\circ)$ or the other way round in $S/_\alpha$. Also, by (S4) for $\mathcal{P}^\dagger$ and Lemma 6.10(2) the relation $S/_\alpha$ is chained with respect to $\overline{\beta}/_\alpha, \overline{B}/_\alpha$. Therefore, by Lemma 6.12 there exists an idempotent polynomial f of $S/_\alpha$ satisfying the following conditions:

(a) f is $\overline{B}/_\alpha$-preserving;
(b) $f(\mathbb{A}_v/_\alpha)$ in coordinate v^x contains an (α, β_v)-minimal set;
(c) $f(\mu_x^\circ|_{B_x}) \subseteq \underline{0}_x$, $f(\mu_y^\circ|_{B_y}) \subseteq \underline{0}_y$, $f(\mu_w^\circ|_{B_w}) \subseteq \underline{0}_w$.

As $B_v/_\alpha$ is a module and $\varphi_1^x(v) \in \mathsf{umax}(B_v/_\alpha)$, the set $\{\varphi_1^x(v), B\}$ is an (α, β_v)-subtrace of $\mathbb{A}_v/_\alpha$. Also, $\mathbf{a}''_\alpha \in \mathsf{umax}(S''/_\alpha) \cap S^\dagger/_\alpha$, therefore by Lemma 6.12(f) for $S/_\alpha$ the polynomial f can be chosen such that

(d) $f(\varphi_1^x(v)) = \varphi_1^x(v)$, $f(B) = B$ in the coordinate position v^x; and
(e) $f(\mathbf{a}''_\alpha) = \mathbf{a}''_\alpha$.

The appropriate restrictions of f are also polynomials of $Q_x/_\alpha, Q_y$. Therefore applying f to $(a'', c', \varphi/_\alpha)$ and (b'', c'', φ) we get $(a', c^*, \varphi^*) \in Q_x/_\alpha$, $(b', c^*, \varphi'^*) \in Q_y$, where $c^* = f(c') = f(c'')$, $\varphi^* = f(\varphi/_\alpha)$ in coordinate positions V^x, and $\varphi'^* = f(\varphi)$ in coordinate positions V^y (and so $f(\varphi/_\alpha) = \varphi^*$ does not have to be true in V^y). By (d) $\varphi^*(v) = B$. Thus, $\mathbf{b}_\alpha = (a', b', c^*, \varphi^*, \varphi'^*) \in S'/_\alpha$.

However, $(a', c^*, \varphi^*), (b', c^*, \varphi'^*)$ do not necessarily belong to $Q_x^\dagger/_\alpha = P^\dagger[C^{xw}, W']/_\alpha$ and $Q_y^\dagger = P^\dagger[C^{yw}, W']$, respectively. To fix this choose $\mathbf{a}'' \in \mathsf{umax}(S^\dagger)$ and $\mathbf{b} \in S'$ such that $\mathbf{a}''/_\alpha = \mathbf{a}''_\alpha, \mathbf{b}/_\alpha = \mathbf{b}_\alpha$. Let $\mathbf{c}$ be a tuple in $\mathrm{Sg}_{S'}(\mathbf{a}'', \mathbf{b})$ such that $\mathbf{a}''\mathbf{c}$ is a thin affine edge and $\mathbf{c}[v^x] \in B$. As is easily seen, $\mathbf{c}$ has the form $\mathbf{c} = (a', b', c^\circ, \varphi^{**}, \varphi'^{**})$. As, $\mathrm{pr}_{xwV^x}\mathbf{a}'' \in Q_x^\dagger$, $\mathrm{pr}_{ywV^y}\mathbf{a}'' \in Q_y^\dagger$, and by Lemma 8.1(4) these relations are polynomially closed in Q_x, Q_y, respectively, $(a', c^\circ, \varphi^{**}) \in Q_x^\dagger$, $(b', c^\circ, \varphi'^{**}) \in Q_y^\dagger$, as well. Therefore $\mathbf{c} \in S^\dagger$ and $\psi(a, b) = \psi(a', b') = B$, a contradiction showing that the case $\psi(a, b) = D \neq B$ is impossible.

Finally, we prove that $(a, b, e^x, e^y) \in \mathrm{pr}_{xyv^xv^y} S^\dagger$ for some $e^x, e^y \in B$ that implies the result. By what is proved above we assume that $(a, b, e^x) \in \mathrm{pr}_{xyv^x} S^\dagger$. We again apply the Congruence Lemma 6.15 to S, S', and $S^\dagger$, this time with $I = \{x, y, v^x\}$ and $I_0 = \{v^y\}$. By Lemma 6.15 either $\mathsf{umax}(\mathrm{pr}_{xyv^x} S^\dagger) \times R^{v\dagger}/_\alpha \subseteq \mathrm{pr}_{xyv^xv^y} S^\dagger/_\alpha$ (in which case we have $(a, b, e^x, e^y) \in \mathrm{pr}_{xyv^xv^y} S^\dagger$ for some $e^y \in B$ and the result follows), or there is a congruence η of $\mathrm{pr}_{xyv^x} S$ such that $\mathrm{pr}_{xyv^xv^y} S^\dagger$ is the graph of a mapping $\psi : \mathrm{pr}_{xyv^x} S^\dagger \to R^{v\dagger}/_\alpha$ and $\eta|_{\mathrm{pr}_{xyv^x} S^\dagger}$ is the kernel of ψ. If $\psi(a, b, e^x) = B$ for $e^x \in B$, then again $(a, b, e^x, e^y) \in \mathrm{pr}_{xyv^xv^y} S^\dagger$ for some $e^y \in B$ implying the desired result.

We again prove that the case $\psi(a, b, e^x) = D \neq B$ for $e^x \in B$ is impossible. As before we consider the relations $\mathcal{S}^\dagger_{/W'}/_\alpha, Q^y/_\alpha, S/_\alpha, S'/_\alpha, S^\dagger/_\alpha$, where α is defined as a product congruence which is the equality relation in every coordinate except for v, v^y, in which it is equal to α. Let $\mathbf{a} = (a, b, d, \varphi^x, \varphi^y) \in S^\dagger$ such that $\varphi^x, \varphi^y \in \mathcal{S}^\dagger_{/W'}$ and $\varphi^x(v) \in B$, and let $\mathbf{a}_\alpha$ be the corresponding tuple in $S^\dagger/_\alpha$. Let

$$S''/_\alpha = \{\mathbf{b} \in S'/_\alpha \mid \mathbf{b}[v^y] = D\},$$

and let $\mathbf{a}' = (a', b', d', e'^x) \in \mathsf{umax}(\mathrm{pr}_{xywv^x} S''/_\alpha)$ be such that $\mathrm{pr}_{xywv^x}\mathbf{a} \sqsubseteq \mathbf{a}'$ in $\mathrm{pr}_{xywv^x} S''/_\alpha$. By s-closedness $\mathbf{a}' \in \mathrm{pr}_{xywv^x} S^\dagger = \mathrm{pr}_{xywv^x} S^\dagger/_\alpha$ and we can find $\mathbf{a}''_\alpha = (a', b', d', \varphi_1^x, \varphi_1^y) \in \mathsf{umax}(S''/_\alpha) \cap S^\dagger/_\alpha$ for some φ_1^x, φ_1^y such that $\varphi_1^x(v^x) = e'_x$. As $\mathrm{pr}_{xyv^x} S^\dagger/_\eta$ is a module, (a', b', e'^x) is in the same η-block with (a, b, e^x), in particular $e'^y \in D$. For the same reason (a, b) and (a', b') are in the same $\tau_{C^{xy}}$-block. We also use a solution φ of $\mathcal{P}^\dagger_{/W'}$ such that $a' \in \varphi(x)$, $b' \in \varphi(y)$, and $\varphi(v) \in B$. In other words, there are $(a'', c', \varphi) \in Q^\dagger_x$ and $(b'', c'', \varphi) \in Q^\dagger_y$ with $a'' \stackrel{\mu^\circ_x}{\equiv} a'$, $b'' \stackrel{\mu^\circ_y}{\equiv} b'$, $c' \stackrel{\mu^\circ_w}{\equiv} c''$. Repeating the same argument as before we find a polynomial g of $S/_\alpha$ satisfying the conditions (a)–(e) using the (α, β_v)-subtrace $\{\varphi_1^y(v^y), B\}$ in the coordinate position v^y in place of $\{\varphi_1^y(v^x), B\}$ in the coordinate position v^x. Then we conclude that for some $c^\bullet \in \mathrm{Sg}_{\mathbb{A}_w}(d', g(c'))$, such that $d'c^\bullet$ is a thin affine edge it holds that $(a', b', c^\bullet, \varphi^{\bullet x}, \varphi^{\bullet y}) \in S'$ and $(a, c^\bullet, \varphi^{\bullet x}) \in Q^\dagger_x$, $(b, c^\bullet, \varphi^{\bullet y}) \in Q^\dagger_y$ for some $\varphi^{\bullet x}, \varphi^{\bullet y} \in \mathcal{S}^\dagger_{/W'}$ with $\varphi^{\bullet x}(v^x), \varphi^{\bullet y}(v^y) \in B$. This implies that $\psi(a, b, \varphi^x(v)) = \psi(a', b', \varphi^{\bullet x}(v^x)) = B$, a contradiction. □

8.5 Conditions (S3)

In this section we prove that $\mathcal{P}^\ddagger$ satisfies conditions (S3).

As before, let $W = W(v, \alpha, \beta_v, \overline{\beta})$ and $W' = W$ if $(v, \alpha, \beta_v) \in \mathcal{W}'$ and $W' = \varnothing$ otherwise. Recall also that for a coherent set $U = W(u, \gamma, \delta, \overline{\beta})$, $(u, \gamma, \delta) \notin \mathcal{W}'$ by $\overline{\mu}_{/U}$ we denote the collection of congruences $\mu_{w/U}$, $w \in V$ such that $\mu_{w/U} = \mu_w$ if $w \in \mathsf{MAX}(\mathcal{P}) - U$, and $\mu_{w/U} = \underline{0}_w$ otherwise.

Lemma 8.9. *The instance $\mathcal{P}^\ddagger$ satisfies (S3). That is, for every coherent set U the problem $\mathcal{P}^\ddagger_{/U}$ is minimal. More precisely, for every $\langle \mathbf{s}, R^\ddagger \rangle \in C^\ddagger$, and every $\mathbf{a} \in R^\ddagger$, there is a solution $\varphi \in \mathcal{S}^\ddagger_{/U}$ such that $\varphi(\mathbf{s}) = \mathbf{a}/_{\overline{\mu}_{/U}}$.*

Proof. For a coherent set U and a constraint $C = \langle \mathbf{s}, R^\ddagger \rangle$ it suffices only to check that tuples from $R'^\ddagger$ are extendable to solutions of $\mathcal{S}^\ddagger_{/U}$, because $R^\ddagger$ is generated by $R'^\ddagger$. Let $\mathbf{s} = (v_1, \dots, v_k)$.

Let $S = S[W', U]$, $S' = S \cap \overline{B}$, and $S^\dagger = S^\dagger[W', U]$ be as defined in Lemma 8.1(7). Let $C = \{C_1, \dots, C_\ell\}$, $C_j = \langle \mathbf{s}_j, R_j \rangle$, and let $\mathbf{a} \in R'^\ddagger$ and $\mathbf{a}' = \mathbf{a}/_{\overline{\mu}_{/U}}$. As before we denote by $V^i = \{u^i \mid u \in V\}$, $i \in [\ell]$, the coordinate positions of $S, S', S^\dagger$. It suffices to show that for some $\mathbf{c} \in \mathrm{pr}_{V-\mathbf{s}} \mathcal{S}^\dagger_{/U}$ such that $(\mathbf{a}', \mathbf{c}) \in \mathsf{umax}(\mathcal{S}^\dagger_{/U})$ it holds that $(\mathbf{a}', \mathbf{c}, \varphi^1, \dots, \varphi^\ell) \in S^\dagger$ for some solutions $\varphi^1, \dots, \varphi^\ell \in \mathcal{S}^\dagger_{/W'}$ such that $\varphi^i(v) \in B/_{\mu^\circ_v}$.

By (S3) for $\mathcal{P}^\dagger$ there is $\sigma \in \mathsf{umax}(\mathcal{S}^\dagger_{/U})$ with $\sigma(\mathbf{s}) = \mathbf{a}'$. Choose one for which the condition

$$\sigma(\mathbf{s}^*) \in R^{*\ddagger}/_{\overline{\mu}_{/U}} \tag{1}$$

is true for a maximal number of constraints $C^* = \langle \mathbf{s}^*, R^* \rangle$ from C. Suppose that condition (1) holds for constraints $C_1, \dots, C_{j-1}$ and does not hold for the remaining ones. We will construct another solution $\sigma_0 \in \mathcal{S}^\dagger_{/U}$ such that (1) for σ_0 is true for all constraints it is true for σ, and is also true for C_j.

We apply the Congruence Lemma 6.15 to the relation S, and its subalgebras S' and $S^\dagger$ with $I = \{v_1, \dots, v_k, v^1, \dots, v^{j-1}\}$ and $I_0 = \{v^j\}$. Let $I' = I \cup \{v^j\}$. Let α denote the product congruence of $\mathcal{S}_{/W'}, P[C_j, W'], S, S', S^\dagger$ which equals the equality relation in all coordinate positions except v for $\mathcal{S}_{/W'}$, and v^j in all the other cases. As is easily seen using Lemmas 7.4, 8.1(8) and conditions (S4), (S6), $S, S', S^\dagger$ satisfy the conditions of the Congruence Lemma 6.15. Therefore, either $\mathsf{umax}(\mathrm{pr}_I S^\dagger) \times R^{v\dagger}/_\alpha \subseteq \mathrm{pr}_{I'} S^\dagger/_\alpha$ (in which case we have $(\mathbf{a}', e^1, \dots, e^j) \in \mathrm{pr}_{I'} S^\dagger$ for some $e^1, \dots, e^j \in B/_{\mu^\circ_v}$ and the result follows), or there is a congruence η of $\mathrm{pr}_I S$ such that $\mathrm{pr}_{I'} S^\dagger/_\alpha$ is the graph of a mapping $\psi : \mathrm{pr}_I S^\dagger \to R^{v\dagger}/_\alpha$ and $\eta|_{\mathrm{pr}_I S^\dagger}$ is the kernel of ψ. If $\psi(\mathbf{a}', e^1, \dots, e^{j-1}) = B$ for some $e^1, \dots, e^{j-1} \in B/_{\mu^\circ_v}$, we obtain $(\mathbf{a}', e^1, \dots, e^{j-1}, e^j) \in \mathrm{pr}_{I'} S^\dagger$ for some $e^j \in B/_{\mu^\circ_v}$ implying the desired result.

Suppose that for any $e^1, \dots, e^{j-1} \in B/_{\mu^\circ_v}$, $\psi(\mathbf{a}', e^1, \dots, e^{j-1}) = D \neq B$, and fix some $e^1, \dots, e^{j-1} \in B/_{\mu^\circ_v}$ with $(\mathbf{a}', e^1, \dots, e^{j-1}) \in \mathrm{pr}_I S^\dagger$. Let $\mathbf{d}_\alpha = (\sigma, \varphi^1, \dots, \varphi^\ell) \in S^\dagger/_\alpha$ be an extension of σ such that $\varphi^1(v^1) = e^1, \dots, \varphi^{j-1}(v^{j-1}) =$

e^{j-1}. Let

$$S''/_\alpha = \{\mathbf{b} \in S'/_\alpha \mid \mathbf{b}[v^j] = D\},$$

and let $\mathbf{d}'_\alpha = (\sigma', \varphi_1^1, \dots, \varphi_1^\ell) \in \mathsf{umax}(S''/_\alpha) \cap S^\dagger/_\alpha$ for some $\varphi_1^1, \dots, \varphi_1^{j-1}, \varphi_1^{j+1}, \dots, \varphi_1^\ell \in \mathcal{S}^\dagger$ and $\varphi_1^j \in \mathcal{S}^\dagger/_\alpha$ be such that $\mathbf{d}_\alpha \sqsubseteq \mathbf{d}'_\alpha$ in $S''/_\alpha$. Such a tuple exists due to s-closedness. As $\mathrm{pr}_I S^\dagger/_\eta$ is a module, $\mathrm{pr}_I \mathbf{d}'_\alpha$ is in the same η-block with $(\mathbf{a}', e^1, \dots, e^{j-1})$, that is, $\psi(\mathrm{pr}_I \mathbf{d}'_\alpha) = D$. For the same reason, for any $C^* = \langle \mathbf{s}^*, R^* \rangle \in C$, it holds that $(\sigma(\mathbf{s}^*), \sigma'(\mathbf{s}^*)) \in \tau_{C^*}/_{\overline{\mu}_{/U}}$. By Lemma 6.3(1) there is $\mathbf{a}'' \in R$ such that $\mathbf{a}'' \in \sigma'(\mathbf{s})$ and $\mathbf{a} \sqsubseteq \mathbf{a}''$. By s-closedness $\mathbf{a}'' \in R^\dagger$.

By construction there is a solution φ of $\mathcal{P}^\dagger_{/W'}$ such that $\mathbf{a}''/_{\overline{\mu}^\circ} = \varphi(\mathbf{s})$ and $\varphi(v) \in B$. Since $\mathbf{a}''/_{\overline{\mu}^\circ} \in \mathsf{umax}(R^\dagger/_{\overline{\mu}^\circ})$, φ can be chosen from $\mathsf{umax}(\mathcal{S}^\dagger_{/W'})$. The existence of φ also means that for any $C^* = \langle \mathbf{s}^*, R^* \rangle \in C$ there is $\mathbf{b}_{C^*} \in R^{*\dagger}$ such that $\mathbf{b}'_{C^*} = \mathbf{b}_{C^*}/_{\overline{\mu}^\circ} = \varphi(\mathbf{s}^*)$. Again, $\mathbf{b}_{C^*}$ can be chosen from $\mathsf{umax}(R^{*\dagger})$. By the definition of $P[C_q, W']$, $q \in [\ell]$, it holds that $(\mathbf{b}_{C_j}, \varphi) \in P^\dagger[C_q, W']$, and so $(\mathbf{b}'_{C_q}, \varphi) \in P^{U\dagger}[C_q, W']$.

The reason why φ does not give rise to a solution from $\mathcal{S}_{/U}$ is that for some $z \in V$ it may be the case that $\mu_z^\circ = \mu_z$, while $\mu_{z/U} = \underline{0}_z$. By the definition of congruences μ_z° and Lemma 8.3(4) for every $z \in V$ such that $\mu_z^\circ = \mu_z$ but $\mu_{z/U} = \underline{0}_z$ the interval (α, β_v) in coordinate position v^j can be separated from $(\underline{0}_z, \mu_z^\circ)$ or the other way round in $S/_\alpha$. Therefore, by Lemma 6.12 there exists an idempotent polynomial f of $S/_\alpha$ satisfying the following conditions:

(a) f is $\overline{B}/_\alpha$-preserving;
(b) $f(\mathbb{A}_v/_\alpha)$ in the coordinate position v^j of $S/_\alpha$ contains an (α, β_v)-minimal set; and
(c) $f(\mu_x^\circ|_{B_x}) \subseteq \underline{0}_x$ for $x \in V$ such that $\mu_{x/U} = \underline{0}_x$.

As $B_v/_\alpha$ is a module and $\varphi_1^j(v^j) \in \mathsf{umax}(B_v/_\alpha)$, the set $\{\varphi_1^j(v^j), B\}$ is an (α, β_v)-subtrace of $\mathbb{A}_v/_\alpha$. Also, $\mathbf{d}'_\alpha \in \mathsf{umax}(S''/_\alpha) \cap S^\dagger/_\alpha$, therefore by Lemma 6.12(f) for $S/_\alpha$ the polynomial f can be chosen such that

(d) $f(\varphi_1^j(v^j)) = \varphi_1^j(v^j)$, $f(B) = B$ in the coordinate position v^j; and
(e) $f(\mathbf{d}'_\alpha) = \mathbf{d}'_\alpha$.

The appropriate restrictions of f are also polynomials of $P^U[C_q, W']$ and $R_q/_{\overline{\mu}_{/U}}$ for each $q \in [\ell]$. Therefore applying f to $(\mathbf{b}'_{C_q}, \varphi)$ for $q \in [\ell] - \{j\}$ and to $(\mathbf{b}'_{C_j}, \varphi/_\alpha)$ we obtain $(\sigma'(\mathbf{s}_q), \varrho^q)$ for $q \in [\ell]$. By (c) for any $C^\circ = \langle \mathbf{s}^\circ, R^\circ \rangle, C^\bullet = \langle \mathbf{s}^\bullet, R^\bullet \rangle \in C$ we have $f(\mathbf{b}'_{C^\circ}[w]) = f(\mathbf{b}'_{C^\bullet}[w])$ for each $w \in \mathbf{s}^\circ \cap \mathbf{s}^\bullet$. This means that $\sigma_0 = f(\varphi)$ is properly defined by setting $\sigma_0(w) = f(\mathbf{b}'_{C^\bullet}[w])$ for any $w \in V$ and $C^\bullet = \langle \mathbf{s}^\bullet, R^\bullet \rangle \in C$ such that $w \in \mathbf{s}^\bullet$. Also, for any constraint $C_q \in C$ for which (1) holds for σ, it also holds for σ_0, as $f(\sigma'(\mathbf{s}_q)) = \sigma'(\mathbf{s}_q) \overset{\tau_{C_q}}{\equiv} \mathbf{b}'_{C_q}$ implies $\sigma_0(\mathbf{s}_q) = f(\mathbf{b}'_{C_q}) \overset{\tau_{C_q}}{\equiv} f(\sigma'(\mathbf{s}_q)) \overset{\tau_{C_q}}{\equiv} \mathbf{b}'_{C_q}$ in this case. By (e), $\sigma_0(\mathbf{s}) = \mathbf{a}''/_{\overline{\mu}_{/U}}$. Finally, $f(B) = B$ in the coordinate position v^j of $S/_\alpha$, and so $\sigma_0(\mathbf{s}_j) \overset{\tau_{C_j}}{\equiv} \mathbf{b}'_{C_j}$, that is, (1) holds for C_j as well. Let $\mathbf{e}_\alpha = (\sigma_0, \varrho^1, \dots, \varrho^\ell)$.

The mapping σ_0 satisfies many of the desired properties, and it is a solution of $\mathcal{P}_{/U}$ because $\sigma_0(\mathbf{s}^\circ) \in R^\circ/_{\overline{\mu}_{/U}}$ for each $C^\circ = \langle \mathbf{s}^\circ, R^\circ \rangle \in C$. However, it is not necessarily a solution of $\mathcal{P}^\dagger_{/U}$, and so we need to make one more step. To convert σ_0 into a solution of $\mathcal{P}^\dagger_{/U}$ consider $\mathbf{d}'' \in \mathsf{umax}(S^\dagger)$, $\mathbf{e} \in S'$ such that $\mathbf{d}''_\alpha = \mathbf{d}''/_\alpha$, $\mathbf{e}_\alpha = \mathbf{e}/_\alpha$. In the subalgebra of S generated by $\mathbf{d}'', \mathbf{e}$ take $\mathbf{c} = (\nu, \varrho'^1, \dots, \varrho'^\ell)$ such that $\mathbf{d}''\mathbf{c}$ is a thin affine edge and $\mathbf{c}[v^j] \in B = \varrho^j(v^j)$. We prove that $\mathbf{c} \in S^\dagger$ contradicting the assumption that $\psi(\mathrm{pr}_I \mathbf{d}'') \neq B$.

For every $C^\circ = \langle \mathbf{s}^\circ, R^\circ \rangle \in C$ the relation $R^{\circ\dagger}/_{\overline{\mu}_{/U}}$ is polynomially closed in $R^\circ/_{\overline{\mu}_{/U}}$ by (S5). Since $\sigma(\mathbf{s}^\circ)\nu(\mathbf{s}^\circ)$ is a thin affine edge in the subalgebra generated by $\sigma'(\mathbf{s}^\circ), \sigma_0(\mathbf{s}^\circ)$, and $\sigma_0(\mathbf{s}^\circ)$ is the image of $\mathbf{b}'_{C^\circ} \in R^{\circ\dagger}/_{\overline{\mu}_{/U}}$ under f, we get $\nu(\mathbf{s}^\circ) \in R^{\circ\dagger}/_{\overline{\mu}_{/U}}$, as well. Thus, ν is a solution of $\mathcal{P}^\dagger_{/U}$. By the definition of $P^{U\dagger}[C_q, W']$, we get $\mathrm{pr}_{\mathbf{s}_q \cup V^q} \mathbf{c} \in P^{U\dagger}[C_q, W']$ implying $\mathbf{c} \in S^\dagger$. □

9 PROOF OF THEOREM 5.6: NON-AFFINE FACTORS

In this section we consider Case 2 of tightening instances: for every $v \in V$ and every $\alpha \in \mathsf{Con}(\mathbb{A}_v)$ with $\alpha \prec \beta_v$ it holds that $\mathsf{typ}(\alpha, \beta_v) \neq \mathbf{2}$.

Let $\mathcal{P} = (V, C)$ be a (2,3)-minimal and block-minimal instance with subdirectly irreducible domains, $\overline{\beta} = \prod_{v \in V} \beta_v$ and $\overline{B} = \prod_{v \in V} B_v$, $\beta_v \in \mathsf{Con}(\mathbb{A}_v)$, B_v is a β_v-block, and $v \in V$. Let also $\mathcal{P}^\dagger = (V, C^\dagger)$ be a $(\overline{\beta}, \overline{B})$-compressed instance, and for $C = \langle \mathbf{s}, R \rangle \in C$ there is $C^\dagger = \langle \mathbf{s}, R^\dagger \rangle \in C^\dagger$. We select $v \in V$ and $\alpha \in \mathsf{Con}(\mathbb{A}_v)$ with $\alpha \prec \beta_v$, $\mathsf{typ}(\alpha, \beta_v) \neq \mathbf{2}$, and an α-block $B \in B_v/_\alpha$ such that B is as-maximal in $R^{v\dagger}/_\alpha$. Let E denote the as-component of $R^{v\dagger}/_\alpha$ containing B. By (S6) for $\mathcal{P}^\dagger$ the relation $R^{v\dagger}$ is weakly as-closed in $B_v/_\alpha$. Therefore the α-block B is as-maximal in $B_v/_\alpha = (R^v \cap B_v)/_\alpha$. We show how $\mathcal{P}^\dagger$ can be transformed to a $(\overline{\beta}', \overline{B}')$-compressed instance such that $\beta'_w \leq \beta_w$, $B'_w \subseteq B_w$ for $w \in V$, and $\beta'_v = \alpha$, $B'_v = B$.

By Lemma 6.11 if $R^{v\dagger}/_\alpha$ contains a nontrivial as-component, there is a coherent set associated with the triple (v, α, β_v). Let $W = W(v, \alpha, \beta_v, \overline{\beta})$ in this case; note that $(v, \alpha, \beta_v) \notin \mathcal{W}'$, because $(\alpha : \beta_v) \neq \underline{1}_v$ by Lemma 6.6(2). Let also $\mathcal{S}^\dagger_{/U}$ denote the set of solutions of $\mathcal{P}^\dagger_{/U}$ for a coherent set U.

Lemma 9.1. *If $B_v/_\alpha$ contains a nontrivial as-component, then for every $w \in W$ there is a congruence $\alpha_w \in \mathsf{Con}(\mathbb{A}_w)$ with $\alpha_w < \beta_w$, and such that $R^{vw\dagger}$ is aligned with respect to (α, α_w), that is, for any $(a_1, a_2), (b_1, b_2) \in R^{vw\dagger}$, $a_1 \stackrel{\alpha}{\equiv} b_1$ if and only if $a_2 \stackrel{\alpha_w}{\equiv} b_2$.*

Proof. It suffices to show that the link congruences $\mathsf{lk}_1, \mathsf{lk}_2$ of $Q = R^{vw}$ viewed as a subdirect product of $\mathbb{A}_v \times \mathbb{A}_w$ are such that $\beta_v \wedge \mathsf{lk}_1 \leq \alpha$ and $\beta_w \wedge \mathsf{lk}_2 < \beta_w$. Since $w \in W$ there are $\gamma, \delta \in \mathsf{Con}(\mathbb{A}_w)$ such that $\gamma \prec \delta \leq \beta_w$ and (α, β_v) and (γ, δ) cannot be separated. By Lemmas 6.6(2), 6.9(2), and 7.1 it follows that $\beta_v \wedge \mathsf{lk}_1 \leq \alpha$ and $\mathsf{lk}_2 \wedge \delta \leq \gamma$. We set $\alpha_w = \beta_w \wedge \mathsf{lk}_2 < \beta_w$. ◻

Let $\mathcal{P}^\ddagger = (V, C^\ddagger)$ be constructed as follows.

(R) Let $\mathcal{P}'$ be the problem obtained from $\mathcal{P}^\dagger$ by adding extra constraint $\langle \{v\}, B \rangle$. Let $\mathcal{P}^\ddagger$ be the problem obtained from $\mathcal{P}'$ by establishing $(2,3)$-consistency, 1-minimality, and the minimality of $\mathcal{P}^\ddagger_{/U}$ for every non-central coherent set U.

Set $\beta'_v = \alpha$, $B'_v = B$. Let Z be the set of variables w such that there is a congruence $\alpha_w < \beta_w$ such that $R^{wv\dagger}/_\alpha$ is the graph of a mapping $\pi_w : R^{w\dagger} \to R^{v\dagger}/_\alpha$ and α_w is its kernel. For instance, by the Congruence Lemma 6.15 if $R^{v\dagger}/_\alpha$ contains a nontrivial as-component, then $Z = W$. For $w \in Z$ set $\beta'_w = \alpha_w$, $B'_w = \pi^{-1}(B)$. For the remaining variables w set $\beta'_w = \beta_w$, $B'_w = B_w$.

We first observe several properties of compound relations $S[x, y, w, v]$ and $S[U, w, v]$.

Lemma 9.2. (1) *For $x, y, w \in V$ let*

$$S^\dagger[x, y, w, v](x, y, w, v) = R^{xy\dagger}(x, y) \wedge R^{xw\dagger}(x, w) \wedge R^{yw\dagger}(y, w) \wedge R^{wv\dagger}(w, v).$$

Then for any $t_1, t_2 \in \{x, y, w, v\}$ it holds that $R^{t_1 t_2} \subseteq \mathrm{pr}_{t_1 t_2} S[x, y, w, v]$ and $R^{t_1 t_2 \dagger} \subseteq \mathrm{pr}_{t_1 t_2} S^\dagger[x, y, w, v]$.

(2) *$S^\dagger[x, y, w, v]$ is polynomially closed in $S[x, y, w, v]$.*

(3) *For $w \in V$, and a non-central coherent set U let*

$$S^\dagger[U, w, v](V, v') = \mathcal{S}^\dagger_{/U}(V) \wedge R^{wv\dagger}/_{\mu_{w/U} \times \underline{0}_v}(w, v').$$

Then $\mathrm{pr}_V S[U,w,v] = \mathrm{pr}_V \mathcal{S}_{/U}$, $\mathrm{pr}_V S^\dagger[U,w,v] = \mathrm{pr}_V \mathcal{S}^\dagger_{/U}$, *for any* $u \in V$ *it holds that* $R^{uv}/_{\mu_{u/U}\times \underline{0}_v} \subseteq \mathrm{pr}_{uv'} S[U,w,v]$ *and* $R^{uv\dagger}/_{\mu_{u/U}\times \underline{0}_v} \subseteq \mathrm{pr}_{uv'} S^\dagger[U,w,v]$.

(4) $S^\dagger[U,w,v]$ *is polynomially closed in* $S[U,w,v]$.

PROOF. Item (1) is straightforward from the (2,3)-minimality of $\mathcal{P}$ and (S2) for $\mathcal{P}^\dagger$. Item (2) follows from (S5) for $\mathcal{P}^\dagger$ and Lemma 6.14(5). Items (3) and (4) are similar. □

We start with an auxiliary statement. For $x \in V$ let

$$T^x = \{a \in R^{x\dagger} \mid \text{ there is } d \in B \text{ such that } (a,d) \in R^{xv\dagger}\}.$$

Similar to notation $\overline{B}, \overline{B}_U$ we use $\overline{T} = \prod_{x\in V} T^x$, $\overline{T}_U = \prod_{x \in U} T^x$ for $U \subseteq V$.

LEMMA 9.3. (1) *Let* $x,y,w \in V$, $(a,b) \in \mathsf{umax}(R^{xy\dagger} \cap \overline{T}_{\{x,y\}})$ *and let* $c \in R^{w\dagger}$ *be such that* $(a,c) \in R^{xw\dagger}$, $(b,c) \in R^{yw\dagger}$, $(c,e) \in R^{wv\dagger}$ *for some* $e \in R^{v\dagger}$ *with* $e/_\alpha \in E$. *Then there is* $d \in T^w$ *such that* $(a,d) \in R^{xw\dagger}$, $(b,d) \in R^{yw\dagger}$.

(2) *Let* U *be a non-central coherent set,* $I \subseteq V$, $w \in V - I$, *and* $\varphi \in \mathsf{umax}(\mathcal{S}^\dagger_{/U})$ *such that* $\varphi(I) \in \overline{T}_I/_{\overline{\mu}_{/U}}$ *and* $(\varphi(w), e) \in R^{wv\dagger}/_{\mu_{w/U}\times\underline{0}_v}$ *for some* $e \in R^{v\dagger}$ *such that* $e/_\alpha \in E$. *Then there is* $\psi \in \mathcal{S}^\dagger_{/U}$ *such that* $\psi(I) = \varphi(I)$ *and* $\psi(w) \in T^w/_{\mu_{w/U}}$. *Moreover, for any* $C = \langle \mathbf{s}, R\rangle$, *let* $\mathbf{a} \in \mathsf{umax}(R^\dagger) \cap \overline{T}_{\mathbf{s}}$, $\mathbf{a}' = \mathbf{a}/_{\overline{\mu}_{/U}}$ *be such that* $\mathbf{a}'$ *extends to a solution* $\varphi \in \mathcal{S}^\dagger_{/U}$ *such that* $\varphi(v)/_\alpha \in E$. *Then* $\mathbf{a}'$ *also extends to a solution* $\psi \in \mathcal{S}^\dagger_{/U} \cap \overline{T}/_{\overline{\mu}_{/U}}$.

PROOF. (1) If $|E| = 1$, then $e/_\alpha = B$ and there is nothing to prove. So, suppose that $|E| > 1$. Let $S = S[x,y,w,v]$, $S^\dagger = S^\dagger[x,y,w,v]$, $S'^\dagger = \mathrm{pr}_{xyv} S^\dagger$, and $S^{*\dagger} = S'^\dagger/_{\underline{0}_x\times\underline{0}_y\times\alpha}$. As is easily seen, it suffices to show that $(a,b,e') \in S'^\dagger$ for some $e' \in B$. We know that $(a,b,c,e) \in S^\dagger$. If $w \notin W$, then by the Congruence Lemma 6.15 $(c,B) \in R^{wv\dagger}/_\alpha$ whenever $c \in \mathsf{umax}(R^{vw\dagger}/_{\alpha\times\underline{0}_w}[E])$. Since c can be chosen from $\mathsf{umax}(R^{wv\dagger}/_{\underline{0}_w\times\alpha}[E])$, the result follows. So, assume that $w \in W$. If $x \in W$ or $y \in W$, then $e/_\alpha = B$ by (S2). Otherwise by Lemma 9.2(1) $R^{xv\dagger} \subseteq \mathrm{pr}_{xv} S^\dagger$, $R^{yv\dagger} \subseteq \mathrm{pr}_{yv} S^\dagger$, and (α, β_v) can be separated in $S^\dagger$ from any $(\gamma_x, \delta_x), (\gamma_y, \delta_y)$, where $\gamma_x \prec \delta_x \le \beta_x$, $\gamma_y \prec \delta_y \le \beta_y$, and $\gamma_x, \delta_x \in \mathrm{Con}(\mathbb{A}_x)$, $\gamma_y, \delta_y \in \mathrm{Con}(\mathbb{A}_y)$, or the other way round. By (S4) for $\mathcal{P}^\dagger$ the relation S is chained with respect to $\overline{\beta}, \overline{B}$, and by Lemma 9.2(2) $S^\dagger$ is polynomially closed in S. Let $\{e_1, e_2\} \in \mathsf{amax}(B_v)$ be an (α, β_v)-subtrace such that $e_2 \in \mathsf{as}(e_1)$ and $e_1/_\alpha, e_2/_\alpha \in E$. By Lemma 6.12(d) there is a $\overline{B}$-preserving polynomial f of S such that $f(e_1/_\alpha) = e_1/_\alpha$, $f(e_2/_\alpha) = e_2/_\alpha$, and $|f(B_x)| = |f(B_y)| = 1$. Applying the Congruence Lemma 6.15 to $S, S^\dagger$ for $I = \{x,y\}$ and $I_0 = \{v\}$ we obtain $\mathsf{umax}(F) \times E \subseteq S^{*\dagger}$, where $F = \{(d_1,d_2) \mid (d_1,d_2,e^*) \in S'^\dagger$ for some $e^* \in R^{v\dagger}$, $e^*/_\alpha \in E\}$. In particular, $(a,b,e') \in S'^\dagger$ for some $e' \in B$.

(2) If $|E| = 1$, we obtain $e/_\alpha = B$, $\varphi(w) \in T^w/_{\mu_{w/U}}$, and there is nothing to prove. So, suppose that $|E| > 1$. We will use the relations $S = S[U,w,v](V,v')$, $S^\dagger = S^\dagger[U,w,v](V,v')$, $S'^\dagger = \mathrm{pr}_{I\cup\{v'\}} S^\dagger$, and $S^{*\dagger} = S'^\dagger/_\alpha$, where α denotes the product congruence $\prod_{u\in I} \underline{0}_u \times \alpha$. As is easily seen, it suffices to show that $(\psi, e') \in S'^\dagger$ for some $e' \in B$.

If $w \notin W$, by the Congruence Lemma 6.15 $(c,B) \in R^{wv\dagger}/_{\mu_{w/U}\times\alpha}$ whenever $c \in \mathsf{umax}(R^{wv\dagger}/_{\mu_{w/U}\times\alpha}[E])$. Therefore $\varphi(w) \in T^w/_{\mu_{w/U}}$. Assume that $w \in W$. If $I \cap W \neq \varnothing$, then $e/_\alpha = B$. Otherwise by Lemma 9.2(3) $R^{v_i v\dagger}/_{\mu_{v_i/U}\times\underline{0}_v} \subseteq \mathrm{pr}_{v_i v'} S^\dagger$ and, as $v_i \notin W$, $\alpha \prec \beta_v$ can be separated in S from any $\gamma \prec \delta \le \beta_u$ or the other way round, where $\gamma, \delta \in \mathrm{Con}(\mathbb{A}_u)$, $u \in I$. By (S4) for $\mathcal{P}^\dagger$ the relation S is chained with respect to $\overline{\beta}, \overline{B}$, and by Lemma 9.2(4) $S^\dagger$ is polynomially closed in S. Similar to the proof of item (1), let $\{e_1,e_2\} \subseteq \mathsf{amax}(B_v)$ be an (α,β_v)-subtrace such that $e_2 \in \mathsf{as}(e_1)$ and $e_1/_\alpha, e_2/_\alpha \in E$. By Lemma 6.12(d) there is a polynomial f of S such that $f(e_1/_\alpha) = e_1/_\alpha$, $f(e_2/_\alpha) = e_2/_\alpha$, and $|f(B_u)| = 1$ for $u \in I$. Applying the Congruence Lemma 6.15 to S and $S^\dagger$ with I and $I_0 = \{v'\}$ we obtain $\mathsf{umax}(F) \times E \subseteq S^{*\dagger}$, where $F = \{\chi \in \mathrm{pr}_I \mathcal{S}^\dagger_{/U} \mid (\chi, e^*) \in S'^\dagger$ for some $e^* \in R^{v\dagger}$, $e^*/_\alpha \in E\}$. In particular, $(\psi, e') \in S'^\dagger$ for some $e' \in B$.

To prove the second claim of item (2) apply what is proved above to $I = \mathbf{s}, \varphi$ and $w = v$, assuming $R^{vv\dagger}$ to be the equality relation on $R^{v\dagger}$. We obtain $\psi \in \mathcal{S}^{\dagger}_{/U}$ such that $(\psi(x), \psi(v)) \in R^{xv\dagger}/_{\overline{\mu}_{/U}}$ for each $x \in V$, implying $\psi \in \mathcal{S}^{\dagger}_{/U} \cap \overline{T}/_{\overline{\mu}_{/U}}$. □

REMARK 9.4. *By inspecting the proof of Lemma 9.3 it can be observed that if* $(a, b) \in \mathsf{amax}(R^{xy\dagger})$ *in item (1) and* $\varphi(I) \in \mathsf{amax}(\mathcal{S}^{\dagger}_{/U})$ *or* $\mathbf{a} \in \mathsf{amax}(R^{\dagger})$ *in item (2) then* $e' \in B$ *can be chosen such that* $(a, b, e') \in \mathsf{amax}(\mathrm{pr}_{xyv} S^{\dagger}[x, y, w, v])$, *and* $(\psi(I), e') \in \mathsf{amax}(\mathrm{pr}_{I\cup\{v'\}} S^{\dagger}[U, w, v])$ *and* $(\mathbf{a}', e') \in \mathrm{pr}_{\mathbf{s}\cup\{v'\}} S^{\dagger}[U, w, v])$, *respectively.*

The following two lemmas show that the constraints of $\mathcal{P}^{\ddagger}$ are not empty. We do it by identifying a set of tuples in every constraint relation that withstand the propagation algorithms. We start with constructing such sets for $(2, 3)$-consistency. Set

$$Q^x = \{a \in \mathsf{amax}(R^{x\dagger}) \mid \text{ there is } d \in B \text{ such that } (d, a) \in R^{vx\dagger}\}.$$

LEMMA 9.5. *The collection of sets* $Q^{xy} = \mathsf{amax}(R^{xy\dagger}) \cap (Q^x \times Q^y)$, $x, y \in V$, *is a* $(2, 3)$*-strategy for* $\mathcal{P}'$.

PROOF. We need to show that for any $x, y, w \in V$ and $(a, b) \in Q^{xy}$ there is $c \in R^{w\dagger}$ such that $(a, c) \in Q^{xw}$, $(b, c) \in Q^{yw}$. Note first that by (S2),(S3) for $\mathcal{P}^{\dagger}$ the relation Q^{xy} is nonempty and subdirect in $Q^x \times Q^y$. By (S2) for $\mathcal{P}^{\dagger}$ there is c with $(a, c) \in R^{xw\dagger}$, $(b, c) \in R^{yw\dagger}$. Let $S = S[x, y, w, v]$, $S^{\dagger} = S^{\dagger}[x, y, w, v]$, $S'^{\dagger} = \mathrm{pr}_{xyv} S^{\dagger}$, and $S^{*\dagger} = S'^{\dagger}/_{\underline{0}_x \times \underline{0}_y \times \alpha}$. As is easily seen, it suffices to show that $(a, b, d, e) \in \mathsf{amax}(S^{\dagger})$ for some $d \in R^{w\dagger}$ and $e \in B$. Condition (S2) for $\mathcal{P}^{\dagger}$ also implies that $\mathbf{a} = (a, b, e') \in S'^{\dagger}$ for some $e' \in B_v$, and $\mathbf{a}$ can be chosen to be as-maximal in $S'^{\dagger}$, as $(a, b) \in \mathsf{amax}(R^{xy\dagger})$. We use the Quasi-2-Decomposition Theorem 6.7. By Lemma 9.2(1) $(a, b) \in \mathrm{pr}_{xy} S^{*\dagger}$, $(a, B) \in \mathrm{pr}_{xv} S^{*\dagger}$, and $(b, B) \in \mathrm{pr}_{yv} S^{*\dagger}$. By Theorem 6.7 $(a, b, D) \in S^{*\dagger}$ for some $D \in B_v/_{\alpha}$ with $B \sqsubseteq^{as} D$ in $B_v/_{\alpha}$. By Lemma 9.3(1) and Remark 9.4 $(a, b, e) \in \mathsf{amax}(S'^{\dagger})$ for some $e \in B$. Therefore there is $d \in S^w$ with $(a, b, d, e) \in \mathsf{amax}(S^{\dagger})$. □

Let $\mathcal{Q} = \{Q^x \mid x \in V\}$. We say that a tuple $\mathbf{a} \in \prod_{w\in U} \mathbb{A}_w$, $U \subseteq V$, is *$\mathcal{Q}$-compatible* if $\mathbf{a}[w] \in Q^w$ for any $w \in U$.

LEMMA 9.6. *Let* $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$. *Then for any non-central coherent set* U *and any* $\mathcal{Q}$*-compatible tuple* $\mathbf{a} \in \mathsf{amax}(R^{\dagger})$ *there is a* $\mathcal{Q}$*-compatible solution* $\varphi \in \mathcal{S}^{\dagger}_{/U}$ *such that* $\varphi(\mathbf{s}) = \mathbf{a}/_{\overline{\mu}_{/U}}$.

PROOF. Note first that R contains a $\mathcal{Q}$-compatible tuple. Let $R^* = \mathrm{pr}_{\mathbf{s}\cup\{v\}} \mathcal{S}^{\dagger}_{/U}$. Since B is as-maximal in $R^{v\dagger}$, by the Maximality Lemma 6.5(2) and by (S3) there exists $\mathbf{a} \in \mathsf{amax}(\mathrm{pr}_{\mathbf{s}} R^*)$ such that $(\mathbf{a}, a) \in R^*$ for some $a \in B$. As is easily seen, $\mathbf{a}$ is $\mathcal{Q}$-compatible.

The proof of this lemma follows the same lines as the proof of Lemma 9.5. We show by induction that for every I, $\mathbf{s} \subseteq I \subseteq V$, there is $\psi \in \mathsf{amax}(\mathrm{pr}_I \mathcal{S}^{\dagger}_{/U})$ such that $\mathbf{a}' = \psi(\mathbf{s})$, where $\mathbf{a}' = \mathbf{a}/_{\overline{\mu}_{/U}}$ and $\psi(w) \in Q^x/_{\mu_{w/U}}$ for all $w \in I$. The base case, $I = \mathbf{s}$ is given by (S3) for $\mathcal{P}^{\dagger}$.

Suppose the claim is proved for some I, $\mathbf{s} \subseteq I \subseteq V$, and $w \in V - I$. Let $\psi \in \mathsf{amax}(\mathrm{pr}_I \mathcal{S}^{\dagger}_{/U})$ be a partial solution for this set, $\psi(u) \in Q^u/_{\mu_{u/U}}$ and $I' = I \cup \{w\}$. We will use the relations $S = S[U, w, v](V, v')$, $S^{\dagger} = S^{\dagger}[U, w, v](V, v')$, $S'^{\dagger} = \mathrm{pr}_{I\cup\{v'\}} S^{\dagger}$, and $S^{*\dagger} = S'^{\dagger}/_{\alpha}$, where α is viewed as the product congruence with the only non-equality component α in the coordinate position v'. As is easily seen, it suffices to show that $(\psi, d, e) \in \mathsf{amax}(\mathrm{pr}_{I\cup\{w,v'\}} S^{\dagger})$ for some $d \in R^{w\dagger}$ and $e \in B$. Firstly, $\psi \in \mathrm{pr}_I S^{*\dagger}$ by the induction hypothesis, as any value of w can be extended to a pair from $R^{wv\dagger}$. For $u \in I$, we have $(\psi(u), b) \in R^{uv\dagger}/_{\mu_{u/U} \times \underline{0}_v}$ for some $b \in B$. By (S3) for $\mathcal{P}^{\dagger}$ this pair can be extended to a solution from $\mathcal{S}^{\dagger}_{/U}$. This implies $(\psi(u), B) \in \mathrm{pr}_{uv} S^{*\dagger}$. By the Quasi-2-Decomposition Theorem 6.7 $(\psi, D) \in S^{*\dagger}$ for some $D \in \mathsf{as}(B)$ in $B_v/_{\alpha}$. By Lemma 9.3(2) and Remark 9.4 $(\psi, e) \in \mathsf{amax}(S'^{\dagger})$ for some $e \in B$. Therefore there is $d \in S^w$ with $(\psi, d, e) \in \mathsf{amax}(\mathrm{pr}_{I\cup\{w,v'\}} S^{\dagger})$. In particular, $(\psi, d) \in \mathsf{amax}(\mathrm{pr}_{I\cup\{w\}} \mathcal{S}^{\dagger}_{/U})$ and $d \in Q^w$. □

Conditions (S2), (S3) hold for $\mathcal{P}^\ddagger$ by construction. The instance $\mathcal{P}^\ddagger$ does not contain empty constraint relations by Lemmas 9.5 and 9.6, and is 1-minimal by construction, implying (S1). Since B'_w is u-maximal in $B_w/_{\beta'_w}$ for each $w \in Z$, condition (S4) for $\mathcal{P}^\ddagger$ holds by Lemma 6.10(5).

LEMMA 9.7. *For any $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and any non-central coherent set U, the relations $R^\ddagger$, $R^\ddagger/_{\overline{\mu}_{/U}}$ are as-closed in $R'^\dagger$ and $R'^\dagger/_{\overline{\mu}_{/U}}$, respectively, where $R'^\dagger = R^\dagger \cap \overline{B}'$.*

PROOF. Let $R'' = R^\dagger \cap \prod_{x\in\mathbf{s}} T^x$, we show that R'' is as-closed in $R'^\dagger$. Note that it suffices to show that T^x is as-closed in $\mathrm{pr}_x R^\dagger \cap B'_x$ for $x \in \mathbf{s}$. If E is a trivial as-component and $a \in T^x$, $b \in B_x$ with $a \sqsubseteq^{as} b$ in B_x, then $b \in R^{x\dagger}$ by (S6) and $(b, e) \in R^{xv\dagger}$ for some $e \in B$ by the Maximality Lemma 6.5(1) implying $b \in T^x$. If E is nontrivial and the first option of the Congruence Lemma 6.15 holds for $R^{xv\dagger}$ then $\mathsf{umax}(T^x) = \mathsf{umax}(T'^x)$, where

$$T'^x = \{a \in R^{x\dagger} \mid \text{ there is } d \in B_v \text{ such that } d/_\alpha \in E, \text{ and } (a, d) \in R^{xv\dagger}\},$$

and the result follows by (S6). Finally, if E is nontrivial and the second option of the Congruence Lemma 6.15 holds for $R^{xv\dagger}$, then $T^x = B'_x$. In all cases the claim holds.

Next we show that $R^\ddagger$ is as-closed in $R''^\dagger$. The instance $\mathcal{P}^\ddagger$ is obtained from $\mathcal{P}^\dagger$ by adding an extra constraint, whose relation is equal and as-closed in T^v, and establishing various sorts of consistency and minimality. We track this process and make sure by induction that for every $C_0^\dagger = \langle \mathbf{s}_0, R_0^\dagger\rangle \in \mathcal{C}^\dagger$ all the interim relations R_0^* produced from $R_0^\dagger$ remain as-closed in R_0''. The first step of this process can be thought of as replacing every relation $R_0^\dagger$ with R_0''. Every relation R_0'' is as-closed in itself.

Let $\mathcal{P}^*$ be the instance obtained after some steps of establishing consistency and minimality, let $C_0^* = \langle \mathbf{s}_0, R_0^*\rangle$ denote the constraint produced from $C_0 = \langle \mathbf{s}_0, R_0''\rangle$, and suppose that R_0^* is as-closed in R_0''. The next transformation of $\mathcal{P}^*$ can be of one of the 3 types, we show that in each case the as-closedness of the resulting relation is maintained.

(1) For some $x, y, w \in V$ and every $(a, b) \in R^{xy*}$, if $(a, c) \in R^{xw*}$, $(b, c) \in R^{yw*}$ for no $c \in \mathbb{A}_w$, the pair (a, b) is removed from R^{xy*}. Let R^{xy**} denote the resulting relation, $(a, b) \in \mathsf{umax}(R^{xy**}) \subseteq \mathsf{umax}(R^{xy''})$ (the second inclusion is by Lemmas 9.5 and 9.6), and $(a', b') \in R^{xy''}$ such that $(a, b)(a', b')$ is a thin semilattice or affine edge in $R^{xy''}$. Then there is $c \in R^{w\dagger} \cap T^w$ such that $(a, c) \in R^{xw*}$, $(b, c) \in R^{yw*}$. Also, as $(a', b') \in R^{xy\dagger}$, there is $d \in R^{w\dagger}$ such that $(a', d) \in R^{xw\dagger}$, $(b', d) \in R^{yw\dagger}$. Applying the Maximality Lemma 6.5(1) to the relation $R^{xy\dagger}(x, y) \wedge R^{xw\dagger}(x, w) \wedge R^{yw\dagger}(y, w)$ we may assume that cd is a thin edge in $R^{w\dagger}$ of the same type as $(a, b)(a', b')$. Since $(c, e) \in R^{wv\dagger}$ for some $e \in B$, $(d, e') \in R^{wv\dagger}$ for some $e' \in R^{v\dagger}$ such that $e'/_\alpha \in E$. Therefore we are in the conditions of Lemma 9.3(1) and d can be assumed from T^w. By the induction hypothesis for R^{xw*}, R^{yw*} it follows that $(a', b') \in R^{xy**}$.

(2) For $C = \langle \mathbf{s}_0, R_0^*\rangle \in \mathcal{C}^*$ the relation R_0^* is replaced with

$$R_0^{**}(\mathbf{s}_0) = R_0^*(\mathbf{s}_0) \wedge \bigwedge_{x\in\mathbf{s}_0} R^{x*}(x).$$

Since R_0^* and R^{x*} are as-closed in the corresponding relations R_0'' and $R^{x''}$ the result follows by Lemma 6.14(6).

(3) For a non-central coherent set U let $\mathcal{S}^*_{/U}$ denote the set of solutions of the instance $\mathcal{P}^*_{/U}$. For $C = \langle \mathbf{s}_0, R_0^*\rangle \in \mathcal{C}^*$ and every $\mathbf{a} \in R_0^*$, if $\mathbf{a}' = \mathbf{a}/_{\overline{\mu}_{/U}}$ cannot be extended to a solution from $\mathcal{S}^*_{/U}$, the tuple $\mathbf{a}$ is removed from R_0^*. Let R_0^{**} denote the resulting relation, $\mathbf{a} \in \mathsf{umax}(R_0^{**})$, and $\mathbf{b} \in R_0''$ such that $\mathbf{ab}$ is a thin semilattice or affine edge in R_0''. There is $\varphi \in \mathcal{S}^*_{/U}$ and $\psi \in \mathcal{S}^\dagger_{/U}$ such that $\mathbf{a} \in \varphi(\mathbf{s}_0), \mathbf{b} \in \psi(\mathbf{s}_0)$. By the Maximality Lemma 6.5(1) we can assume that $\varphi\psi$ is a thin edge in $\mathcal{S}^\dagger_{/U}$. Since $\varphi(v)/_\alpha = B$, we have $\psi(v)/_\alpha \in E$. Therefore by Lemma 9.3(2) $\psi(v)/_\alpha$ can be assumed to be B and so $\psi(\mathbf{s}_1) \in R_1''$ for every constraint $C_1 = \langle \mathbf{s}_1, R_1\rangle \in \mathcal{C}$. By Lemma 6.14(6) and the induction hypothesis $\psi \in \mathcal{S}^*_{/U}$, and the result follows.

The statement about $R^{\ddagger}/_{\overline{\mu}_{/U}}$ and $R'^{\dagger}/_{\overline{\mu}_{/U}}$ now follows from Lemma 6.14(4). □

LEMMA 9.8. *$\mathcal{P}^{\ddagger}$ satisfies condition (S5). In other words, for every $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and a non-central coherent set U, the relations $R^{\ddagger}, R^{\ddagger}/_{\overline{\mu}_{/U}}$ are polynomially closed in $R, R/_{\overline{\mu}_{/U}}$, respectively.*

PROOF. Let $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and U a non-central coherent set and $R'^{\dagger} = R^{\dagger} \cap \overline{B}'$. Let $\mathbf{a} \in \mathsf{umax}(R^{\ddagger})$ and $\mathbf{b} \in \mathsf{Block}(R^{\ddagger})$ be such that $\mathbf{a} \sqsubseteq^{as} \mathbf{b}$ in $\mathsf{Block}(R^{\ddagger})$. Then, as $\mathsf{umax}(R^{\ddagger}) \subseteq \mathsf{umax}(R'^{\dagger})$, we also have $\mathbf{a} \in \mathsf{umax}(R'^{\dagger})$. By Lemma 6.14(2,3) and (S5) for $\mathcal{P}^{\dagger}$ the relation $R'^{\dagger}$ is polynomially closed, as $R \cap \overline{B}'$ is a congruence block. Moreover, as $R^{\ddagger} \subseteq R'^{\dagger}$, it holds that $\mathsf{Block}(R^{\ddagger}) \subseteq \mathsf{Block}(R'^{\dagger})$ and $\mathbf{b} \in \mathsf{Block}(R'^{\dagger})$. Therefore $\mathbf{b} \in R'^{\dagger}$. Then by Lemma 9.7 we obtain $\mathbf{b} \in R^{\ddagger}$. To prove the result for $R^{\ddagger}/_{\overline{\mu}_{/U}}$ it suffices to replace every relation above with its quotient modulo $\overline{\mu}_{/U}$ and observe that $R/_{\overline{\mu}_{/U}} \cap \overline{B}'/_{\overline{\mu}_{/U}}$ is still a congruence block of $\overline{\beta}/_{\overline{\mu}_{/U}}$. □

Finally, we verify condition (S6).

LEMMA 9.9. *Condition (S6) for $\mathcal{P}^{\ddagger}$ holds.*

PROOF. Let $C = \langle \mathbf{s}, R\rangle \in \mathcal{C}$ and $w \in \mathbf{s}$. If $B'_w = B_w$, the result follows from (S6) for $\mathcal{P}^{\dagger}$, so assume $B'_w \neq B_w$.

By Lemma 9.7 $R^{\ddagger}$ is as-closed in $R'^{\dagger}$. Note that $\mathsf{umax}(\mathrm{pr}_w R^{\ddagger}) \subseteq \mathsf{umax}(\mathrm{pr}_w(R \cap \overline{B}'))$, because Lemma 9.5 implies that $\mathrm{pr}_w R^{\ddagger}$ contains an element of $\mathsf{amax}(\mathrm{pr}_w(R \cap \overline{B}')$ and then we obtain the desired inclusion by Lemma 6.4(3). Now, if $a \in \mathsf{umax}(\mathrm{pr}_w R^{\ddagger}) \subseteq \mathsf{umax}(\mathrm{pr}_w(R \cap \overline{B}'))$ and $b \in \mathrm{pr}_w(R \cap \overline{B}')$ are such that $a \sqsubseteq^{as} b$ for $w \in \mathbf{s}$, then let $\mathbf{a} \in R^{\ddagger}$ with $a = \mathbf{a}[w]$. Since $\mathbf{a} \in R^{\dagger}$, by (S6) for $R^{\dagger}$, $b \in \mathrm{pr}_w R^{\dagger}$, and therefore $b = \mathbf{b}[w]$ for some $\mathbf{b} \in R^{\dagger}$. As $a \sqsubseteq^{as} b$, the tuple $\mathbf{b}$ can be chosen such that $\mathbf{a} \sqsubseteq^{as} \mathbf{b}$ in $R^{\dagger}$. As by Lemma 9.3(2) T^w is as-closed in $R^{w\dagger}$, we have $b \in T^w$. By (S2) for $\mathcal{P}^{\dagger}$ for any $u_1, u_2 \in \mathbf{s} \cap Z$ and any $\mathbf{a} \in R^{\dagger}$, if $\mathbf{a}[u_1] \in B'_w$ then $\mathbf{a}[u_2] \in B'_u$. This implies $\mathbf{b} \in R''$. Hence, $\mathbf{b} \in R^{\ddagger}$ by Lemma 9.7, confirming the claim.

For the s-closedness part of (S6), if $\mathbf{a} \in \mathsf{umax}(R^{\ddagger})$ and $\mathbf{b} \in R \cap \overline{B}'$ with $\mathbf{a} \leq \mathbf{b}$, then $\mathbf{b} \in R^{\dagger}$ by (S6) for $\mathcal{P}^{\dagger}$ and therefore $\mathbf{b} \in R''$. By Lemma 9.7 $\mathbf{b} \in R^{\ddagger}$. □

REFERENCES

[1] Eric Allender, Michael Bauland, Neil Immerman, Henning Schnoor, and Heribert Vollmer. The complexity of satisfiability problems: Refining Schaefer's theorem. In *MFCS*, pages 71–82, 2005.

[2] K.A. Baker and A.F. Pixley. Polynomial interpolation and the Chinese remainder theorem. *Mathematische Zeitschrift*, 143:165–174, 1975.

[3] Libor Barto. The dichotomy for conservative constraint satisfaction problems revisited. In *LICS*, pages 301–310, 2011.

[4] Libor Barto. The collapse of the bounded width hierarchy. *J. Log. Comput.*, 26(3):923–943, 2016.

[5] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Logical Methods in Computer Science*, 8(1), 2012.

[6] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):3:1–3:19, 2014.

[7] Libor Barto, Marcin Kozik, and Todd Niven. The CSP dichotomy holds for digraphs with no sources and no sinks (A positive answer to a conjecture of Bang-Jensen and Hell). *SIAM J. Comput.*, 38(5):1782–1802, 2009.

[8] Libor Barto, Marcin Kozik, and Ross Willard. Near unanimity constraints have bounded pathwidth duality. In *LICS*, pages 125–134, 2012.

[9] Libor Barto, Andrei A. Krokhin, and Ross Willard. Polymorphisms, and how to use them. In *The Constraint Satisfaction Problem: Complexity and Approximability*, pages 1–44. 2017.

[10] Joel Berman, Paweł Idziak, Petar Marković, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Varieties with few subalgebras of powers. *Trans. Amer. Math. Soc.*, 362(3):1445–1473, 2010.

[11] Jonah Brown-Cohen and Prasad Raghavendra. Correlation decay and tractability of CSPs. In *ICALP*, volume 55 of *LIPIcs*, pages 79:1–79:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[12] Andrei A. Bulatov. A dichotomy theorem for constraints on a three-element set. In *FOCS'02*, pages 649–658, 2002.

[13] Andrei A. Bulatov and Peter G. Jeavons. Algebraic structures in combinatorial problems. Technical Report MATH-AL-4-2001, Technische universität Dresden, Dresden, Germany, 2001. available at http://web.comlab.ox.ac.uk/oucl/research/areas/constraints/publications/index.html.

[14] Andrei A. Bulatov. Tractable conservative constraint satisfaction problems. In *LICS*, page 321, 2003.

[15] Andrei A. Bulatov and Peter Jeavons. An Algebraic Approach to Multi-sorted Constraints. In *CP*, Lecture Notes in Computer Science, vol.2833, page 183–198, 2003.

[16] Andrei A. Bulatov. A graph of a relational structure and constraint satisfaction problems. In *LICS*, pages 448–457, 2004.

[17] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006.

[18] Andrei A. Bulatov. Complexity of conservative constraint satisfaction problems. *ACM Trans. Comput. Log.*, 12(4):24, 2011.

[19] Andrei A. Bulatov. Conservative constraint satisfaction re-revisited. *Journal of Computer and System Sciences*, 82(2):347–356, 2016.

[20] Andrei A. Bulatov. Graphs of finite algebras: edges, and connectivity. *Algebra universalis*, 85(4), 45, 2024.

[21] Andrei A. Bulatov. Graphs of relational structures: restricted types. In *LICS*, 2016.

[22] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *FOCS*, pages 319–330, 2017.

[23] Andrei A. Bulatov. Graphs of finite algebras: maximality, rectangularity, and decomposition Algebra universalis 85 (4), 1-33, 2024.

[24] Andrei A. Bulatov. Graphs of relational structures: restricted types. CoRR, abs/2006.11713, 2020.

[25] Andrei A. Bulatov. Separation of congruence intervals and implications. *CoRR*, abs/2007.07237, 2020.

[26] Andrei A. Bulatov and Víctor Dalmau. A simple algorithm for Mal'tsev constraints. *SIAM J. Comput.*, 36(1):16–27, 2006.

[27] Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.

[28] Andrei A. Bulatov, Andrei A. Krokhin, and Benoit Larose. Dualities for constraint satisfaction problems. In *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar].*, pages 93–124, 2008.

[29] Andrei A. Bulatov and Matthew Valeriote. Recent results on the algebraic approach to the CSP. In *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar].*, pages 68–92, 2008.

[30] Stanley Burris and H.P. Sankappanavar. *A course in universal algebra*, volume 78 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1981.

[31] Hubie Chen and Benoît Larose. Asking the metaquestions in constraint tractability. *TOCT*, 9(3):11:1–11:27, 2017.

[32] Víctor Dalmau and Andrei A. Krokhin. Majority constraints have bounded pathwidth duality. *Eur. J. Comb.*, 29(4):821–837, 2008.

[33] Rina Dechter. *Constraint processing*. Morgan Kaufmann Publishers, 2003.

[34] Tomas Feder and Moshe Vardi. Monotone monadic SNP and constraint satisfaction. In *STOC*, pages 612–622, 1993.

[35] Tomas Feder and Moshe Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM Journal of Computing*, 28:57–104, 1998.

[36] Ralph Freese and Ralph McKenzie. *Commutator theory for congruence modular varieties*, volume 125 of *London Math. Soc. Lecture Notes*. London, 1987.

[37] Ralph Freese and Matthew Valeriote. On the complexity of some Maltsev conditions. *IJAC*, 19(1):41–77, 2009.

[38] Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. Treewidth and hypertree width. In *Tractability: Practical Approaches to Hard Problems*, pages 3–38. Cambridge University Press, 2014.

[39] Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, 2004.

[40] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Ser.B*, 48:92–110, 1990.

[41] David Hobby and Ralph McKenzie. *The Structure of Finite Algebras*, volume 76 of *Contemporary Mathematics*. American Mathematical Society, Providence, R.I., 1988.

[42] Pawel M. Idziak, Petar Markovic, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM J. Comput.*, 39(7):3023–3037, 2010.

[43] Peter G. Jeavons, David A. Cohen, and Marc Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997.

[44] Peter G. Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200:185–204, 1998.

[45] Peter G. Jeavons, David A. Cohen, and Martin C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, 1998.

[46] Phokion G. Kolaitis. Constraint satisfaction, databases, and logic. In *IJCAI'03*, pages 1587-1595, 2003.

[47] Phokion G. Kolaitis and Moshe Y. Vardi. A game-theoretic approach to constraint satisfaction. In *AAAI'00*, pages 175–181, 2000.

[48] Marcin Kozik. Weak consistency notions for all the CSPs of bounded width. In *LICS*, pages 633–641, 2016.

[49] Andrei A. Krokhin and Stanislav Zivny, editors. *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[50] Gábor Kun and Mario Szegedy. A new line of attack on the dichotomy conjecture. *Eur. J. Comb.*, 52:338–367, 2016.

[51] Benoit Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3(4), 2007.

[52] Petar Markovic. The complexity of CSPs on a 4-element set. Personal communication, 2011.

[53] Miklós Maróti. Malcev on top. Manuscript, available at http://www.math.u-szeged.hu/~mmaroti/pdf/200x%20Maltsev%20on%20top.pdf, 2011.

[54] Miklós Maróti. Tree on top of Malcev. Manuscript, available at http://www.math.u-szeged.hu/~mmaroti/pdf/200x%20Tree%20on%20top%20of%20Maltsev.pdf, 2011.

[55] Miklós Maróti and Ralph McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3-4):463–489, 2008.

[56] Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4):17:1–17:24, 2008.

[57] Tomas J. Schaefer. The complexity of satisfiability problems. In *STOC'78*, pages 216–226, 1978.
[58] Ágnes Szendrei. *Clones in Universal Algebra.* Séminaire de Mathématiques Supérieures, vol. 99., Les Presses de l'Université de Montréal, Montréal, 1986.
[59] Thomas A. Whitelaw *Introduction to Abstract Algebra.* Springer, 1990.
[60] Ross Willard. Personal communication, 2019.
[61] Dmitriy Zhuk. On key relations preserved by a weak near-unanimity function. In *ISMVL*, pages 61–66, 2014.
[62] Dmitriy Zhuk. On CSP dichotomy conjecture. In *Arbeitstagung Allgemeine Algebra AAA'92*, page 32, 2016.
[63] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In *FOCS*, pages 331–342, 2017.
[64] Dmitriy Zhuk. A Proof of the CSP Dichotomy Conjecture. *J. ACM*, 67(5):30:1–30:78, 2020.