

Equations over Finite Monoids with Infinite Promises

ALBERTO LARRAURI, University of Zaragoza, Spain

ANTOINE MOTTET, Hamburg University of Technology, Germany

STANISLAV ŽIVNÝ, University of Oxford, United Kingdom

Larrauri and Živný [ICALP'24/ACM ToCL'24] recently established a complete complexity classification of the problem of solving a system of equations over a monoid N assuming that a solution exists over a monoid M , where both monoids are finite and M admits a homomorphism to N . Using the algebraic approach to promise constraint satisfaction problems, we extend their complexity classification in two directions: we obtain a complexity dichotomy in the case where arbitrary relations are added to the monoids, and we moreover allow the monoid M to be finitely generated.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms; Problems, reductions and completeness.**

Additional Key Words and Phrases: constraint satisfaction, promise constraint satisfaction, equations, minions, monoids

ACM Reference Format:

Alberto Larrauri, Antoine Mottet, and Stanislav Živný. 2026. Equations over Finite Monoids with Infinite Promises. *ACM Trans. Comput. Logic* 1, 1, Article 1 (January 2026), 23 pages.

1 Introduction

Solving Equations. Deciding the solvability of systems of equations is a fundamental problem in computer science and mathematics. In general, the equation satisfiability problem for an algebraic structure A takes as input a set of equations of the form

$$s(x_1, \dots, x_n) = t(x_1, \dots, x_n)$$

where $x_1, \dots, x_n$ are variables and s, t are terms over the algebraic symbols of the structure. For example, over a monoid the terms would be simply words over the alphabet $\{x_1, \dots, x_n\}$, over groups the terms would be words over the alphabet $\{x_1, x_1^{-1}, \dots, x_n, x_n^{-1}\}$, and over rings the terms would be multivariate polynomials. For now, the reader can imagine that the terms may also contain constants from the algebraic structure A , but we will revisit this assumption in the following. The question is to decide whether there exists an assignment $h: \{x_1, \dots, x_n\} \rightarrow A$ such that every equation becomes true under this assignment.

Besides their intrinsic mathematical appeal, equation satisfiability problems are very natural for computer science, in particular in the case where the equations are to be solved over a finitely generated free structure. Indeed, the elements of such structures are themselves words or terms over the generators and the equation satisfiability problem then coincides with the unification problem, where on an input consisting of equations as above, the goal is to determine the existence

Authors' Contact Information: Alberto Larrauri, llarrauri@unizar.es, University of Zaragoza, Zaragoza, Spain; Antoine Mottet, antoine.mottet@tuhh.de, Hamburg University of Technology, Hamburg, Germany; Stanislav Živný, standa.zivny@cs.ox.ac.uk, University of Oxford, Oxford, United Kingdom.

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.

© 2026 ACM.

ACM 1557-945X/2026/1-ART1

<https://doi.org/>

of a substitution of the variables into terms that would make the equations true, possibly modulo a background equational theory. For example, unification problems over words with a finite alphabet Σ correspond to solving equations in the free monoid generated by Σ . Unification is relevant to areas such as automated reasoning and description logics [2, 3], among many others.

Depending on the underlying algebraic structure over which the equations are to be solved, and depending on the type of equations to be solved, the complexity of the problem varies wildly. While systems of linear equations over a field, where each term s, t in the input is a linear polynomial, are known to be efficiently solvable to any undergraduate student, solving arbitrary polynomial equations is NP-hard and is not considered to be efficiently solvable. It is also known that the equation satisfiability problem in the group $(\mathbb{Z}, +)$ (or any finitely generated free *commutative* group) is solvable in polynomial time [23], while the problem becomes much more complex over arbitrary finitely generated free groups or monoids, for which only algorithms requiring polynomial space are known [28].

Over *finite* structures, it has been known for several decades that the equation satisfiability problem for finite groups is solvable in polynomial time when the group is commutative, and is NP-complete otherwise [18]. For finite monoids, a similar dichotomy has been obtained by Klíma, Tesson, and Thérien: the equation satisfiability problem is solvable in polynomial time for *regular* commutative monoids, and is NP-complete otherwise [24].

Generalized Equations. From here on we do not necessarily allow constants to appear in the instances of the equation satisfiability problem. We outline two possible approaches to obtaining polynomial-time algorithms solving more general problems than equation satisfiability.

On the one hand, one can start with a tractable equation satisfiability problem and additionally allow in the input constraints that are *not* equations. A natural example for this is to allow e.g. constraints of the form $s(x_1, \dots, x_n) \neq t(x_1, \dots, x_n)$ or $(x_1, \dots, x_r) \in R$, where $R \subseteq A^r$ is a fixed subset. The appropriate framework to study these problems is that of *constraint satisfaction problems* whose templates are expansions $A = (A, R)$ of algebraic structures by a relation; we introduce this terminology formally in Section 2. For finite commutative groups, Feder and Vardi [15] showed that allowing any constraint that is not itself expressible by a system of equations (so-called cosets) yields an NP-hard problem. Interestingly, this is not true for finitely generated commutative groups [8] and in particular it is possible in this setting to decide the satisfiability of a system of equations together with disequation constraints as above.

On the other hand, one can start with an NP-hard equation satisfiability problem, and *restrict* the instances allowed as input. This has been investigated e.g. in [26] in the following form. Fix two finite monoids M, N . The *promise* equation problem only allows systems of equations that are promised to be satisfiable in the monoid M or unsatisfiable in N , and the problem is to decide which case applies. When $M = N$, we see that we recover exactly the classical equation satisfiability problem. Note that to specify this problem formally, one needs to also specify a translation between the constant symbols corresponding to elements of M and those corresponding to elements of N ; this is more than a technical detail, as the complexity of the problem also depends on this translation. The suitable framework to study such problems is that of *promise* constraint satisfaction problems [4, 9] whose templates are monoids M and N .

Main Result. In this work, we generalize the state-of-the-art in both aforementioned directions simultaneously. Let M, N be monoids, and let $R(M) \subseteq M^r$ and $R(N) \subseteq N^r$ be relations of some arity r . The problem that we consider takes as input a system of constraints that are either monoid or group equations of the form $s(x_1, \dots, x_n) = t(x_1, \dots, x_n)$ as considered above, or constraints of the form $(x_1, \dots, x_r) \in R$. The problem is to decide whether this system is satisfiable in M (with R

interpreted as $R(M)$), or not satisfiable in N (with R interpreted as $R(N)$). Until proper terminology is introduced below, we call this problem the *generalized equation problem*.

For example, consider M to be the monoid $(\mathbb{Z}_{\geq 0}, +)$ of non-negative integers, and N to be the monoid $(\mathbb{Z}/n\mathbb{Z}, +)$ for an arbitrary $n \geq 2$. Let $R(M)$ contain all the triples obtained by permuting entries of tuples of the form $(a, a, a + 1)$ for $a \in \mathbb{N}$. Let $R(N)$ be the relation containing all triples except for (a, a, a) for all $a \in \mathbb{Z}/n\mathbb{Z}$. An example of an instance to this problem is $\{x + y = u + v, (x, y, u) \in R, (u, v, x) \in R, (u, v, y) \in R\}$, which is not satisfiable in $\mathbf{M} = (M, R(M))$, but is satisfiable in $\mathbf{N} = (N, R(N))$ for all $n \geq 2$.

We note that deciding the satisfiability of such instances in $\mathbf{M}$ is NP-complete.¹ Indeed, 1-in-3-SAT readily reduces to instances of $\mathbf{M}$ that do not even use monoid equations over M and only use the constraints of the form $R(M)$. Similarly, deciding satisfiability of such instances in $\mathbf{N}$ is NP-complete for all $n \geq 2$ by a reduction from Not-All-Equal-SAT. However, when the problem is to distinguish between instances satisfiable in $\mathbf{M}$ and those unsatisfiable in $\mathbf{N}$, then the problem becomes solvable in polynomial time when n is a multiple of 3 and remains NP-hard otherwise.²

We are able to obtain a dichotomy result for the class of such problems where M is finitely generated and where N is a finite monoid.

THEOREM 1.1 (MAIN THEOREM, INFORMAL VERSION). *Let M be a finitely generated monoid, and let N be a finite monoid. Let $R(M) \subseteq M^r$ and $R(N) \subseteq N^r$ be arbitrary relations of some arity r . Then the generalized equation problem parameterized by $\mathbf{M} = (M, R(M))$ and $\mathbf{N} = (N, R(N))$ is either solvable in polynomial time or is NP-hard.*

We remark that this result generalizes the main dichotomy from [26]. At a high level, the classification in [26] involves templates consisting of pairs of finite monoids M , and N^3 enriched with constants $c_1(M), \dots, c_\ell(M)$, and $c_1(N), \dots, c_\ell(N)$. Such a problem can be expressed as in the statement of Theorem 1.1 by letting $R(M) = \{(c_1(M), \dots, c_\ell(M))\}$ and $R(N) = \{(c_1(N), \dots, c_\ell(N))\}$.

Other Related Work. In [13], the authors consider the complexity of the *uniform* CSP over expansions of finitely generated free groups. Since the constraint relations are here part of the input, a fixed encoding of the constraints needs to be agreed upon, and the authors represent constraints by regular expressions over the generators of the group. In particular, not all possible constraints are allowed as input. Under those conditions, the problem is shown to be solvable in polynomial space.

The literature on the equation satisfiability problem over monoids that are *not* finitely generated is sparse as such problems are almost always intractable; see e.g. [17] where the problem over $(\mathbb{N}; \times)$, the free commutative monoid with countably many generators, is investigated.

Finally, an important related problem is the problem of trying to maximize the number of satisfied equations, in cases where the input system cannot be fully satisfied. The promise version of this problem has recently been considered in [11], where it is shown that beating the random assignment is NP-hard, just as in the non-promise setting for systems of equations over the commutative [20] and non-commutative groups [14]. A natural and interesting open problem is to consider the complexity of this problem when constraints other than equations are also allowed on the input.

¹NP containment holds, since for example this problem reduces to the existential theory of Presburger arithmetic.

²The regularization of $\mathbf{M}$ is $(\mathbb{Z}, +)$ and any monoid homomorphism to $\mathbf{N}$ contains in its kernel the set $n\mathbb{Z}$ of multiples of n . The coset $\text{Cos}(R(M))$ consists of all the triples (a, b, c) such that $a + b + c = 1 \pmod 3$. In particular, for all n not divisible by 3, it contains a tuple (a, b, c) such that $a = b = c \pmod n$, and therefore there is no homomorphism from $(\mathbb{Z}, +, \text{Cos}(R(M)))$ to $(\mathbb{Z}/n\mathbb{Z}, +, R(N))$. This implies, by our main result (cf. Theorem 3.1), that $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is NP-hard. In contrast, if n is divisible by 3 then no tuple (a, b, c) such that $a = b = c \pmod 3$ is in $\text{Cos}(R(M))$, and therefore the canonical projection $x \mapsto x \pmod 3$ is a homomorphism from $(\mathbb{Z}, +, \text{Cos}(R(M)))$ to $(\mathbb{Z}/3\mathbb{Z}, +, R(N))$. By Theorem 3.1, in this case $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is solvable in polynomial time.

³The phrase “infinite promises” in the title of this article alludes to the fact that M can be infinite (if finitely generated).

2 Background

We adopt the convention that the set of natural numbers $\mathbb{N}$ begins at 1. Given an integer $n \geq 1$, we write $[n]$ for the set $\{1, \dots, n\}$. We implicitly extend functions $f: U \rightarrow V$ to arbitrary Cartesian powers by coordinate-wise application. In other words, we write $f(u_1, \dots, u_n)$ for the tuple $(f(u_1), \dots, f(u_n))$ for all $(u_1, \dots, u_n) \in U^n$. Given a tuple $\mathbf{u} = (u_1, \dots, u_m) \in U^m$ and a map $\sigma: [n] \rightarrow [m]$, we write $\mathbf{u} \circ \sigma$ for the tuple $(u_{\sigma(1)}, \dots, u_{\sigma(n)}) \in U^n$.

We assume basic familiarity with the notions of semigroup, group, monoid, and their homomorphisms (see, for instance [21]). Nevertheless, we recall some of the definitions in order to introduce notation.

Semigroups, Monoids, and Groups. We introduce here several basic algebraic notions. We warn the reader that our notion of inverse in a monoid deviates from standard references (e.g., [12, 19, 21]).

A *semigroup* S is a set S equipped with a binary associative operation $a \cdot^S b$. In most cases the semigroup is clear from the context and we write ab for $a \cdot^S b$. A *monoid* M is a semigroup that additionally contains an element $e_M \in M$ satisfying $e_M a = a e_M = a$ for all $a \in M$, which we call the *identity element*. A *group* G is a monoid where for each element $a \in G$ there is another element $a^{-1} \in G$ satisfying $aa^{-1} = a^{-1}a = e_G$. We call a semigroup S *commutative* if $ab = ba$ for each $a, b \in S$. An element $a \in S$ is called *idempotent* if $a^2 = a$. A *semilattice* is a commutative monoid where each element is idempotent. Given a semigroup (resp., monoid, group) S and an integer $n \geq 1$, the Cartesian power S^n inherits the semigroup (resp., monoid, group) structure from S in the natural way.

A monoid homomorphism f from a monoid M to a monoid N , denoted $f: M \rightarrow N$, is a map $f: M \rightarrow N$ satisfying both $f(e_M) = e_N$ and $f(ab) = f(a)f(b)$ for all $a, b \in M$. We write $M \rightarrow N$ to denote the fact that M maps homomorphically to N .

Let M be a monoid and $U, V \subseteq M$ be sets. We define $U \otimes V$ as the set $\{st \mid s \in U, t \in V\}$. Similarly, given $n \geq 1$, we define $U^{\otimes n}$ as $\bigotimes_{i=1}^n U$, whereas U^n denotes the n -th Cartesian power of U . Finally, we write $\langle U \rangle$ to denote the submonoid of M generated by U , which is the smallest submonoid of M containing all elements of U . We say that S *generates* M if $\langle S \rangle = M$.

Let S be a semigroup. A subsemigroup $T \leq S$ is called a *subgroup* of S if there is an element $e_T \in T$ that acts as the identity element in T , and each element $s \in T$ has an inverse $t \in T$ satisfying $st = ts = e_T$. Observe that even if S is a monoid, a subgroup $T \leq S$ is not necessarily a submonoid: T does not need to contain the identity e_S . A commutative semigroup S is called *regular* if every element in S belongs to a subgroup⁴. Under $P \neq NP$, systems of equations over a finite monoid M can be solved in polynomial time if, and only if, M is commutative and regular [24, 26]. Given a semigroup S , the preorder $a \leq b$ contains all the pairs $a, b \in S$ for which either $a = b$ or there exist elements c_1, c_2 satisfying $c_1 b = b c_2 = a$. We write $a \mathcal{H} b$ whenever both $a \leq b$ and $b \leq a$ hold. It can be seen that $\mathcal{H}$ is an equivalence relation. This way, we denote the $\mathcal{H}$ -class of an element $a \in S$ by H_a . The following is sometimes referred to as Green's Theorem:

THEOREM 2.1 ([21, THEOREM 2.2.5]). *Let S be a semigroup, and $T \subseteq S$ be a $\mathcal{H}$ -class. Then T is a subgroup of S if, and only if, T contains some idempotent element. In particular, no $\mathcal{H}$ -class of S contains more than one idempotent element.*

This result implies that the maximal subgroup containing an idempotent element $d \in S$ is precisely H_d . Additionally, an element $a \in S$ belongs to a subgroup if, and only if, $a \mathcal{H} d$ for some

⁴We remark that in the non-commutative setting, there is a difference between regular semigroups S , where it is required that for each $a \in S$ there is some $b \in S$ such that $aba = a$ and $bab = b$, and completely regular semigroups S (also called unions of subgroups [19]), where each element belongs to a subgroup. For commutative semigroups both notions coincide, so there should be no confusion regarding this.

idempotent element d . In this situation we say that a is a *group element* of S and write a^{-1} for its inverse in H_a . Equivalently, a^{-1} is the only element satisfying $aa^{-1} = a^{-1}a$ together with $a^2a^{-1} = a$ and $a^{-2}a = a^{-1}$, where we use a^{-k} to denote $(a^{-1})^k$.⁵ Given a set $S \subseteq M$ of group elements we write S^{-1} for the set $\{a^{-1} \mid a \in S\}$.

Let M be a commutative monoid. A *coset* of M is a set $U \subseteq M$ of group elements satisfying $U \otimes (U^{-1} \otimes U) = U$. Given a set $U \subseteq M$ of group elements, the *coset generated by U* , denoted by $\text{Cos}(U)$, is defined as $U \otimes \langle U^{-1} \otimes U \rangle$. It can be seen that $\text{Cos}(U)$ is a coset for all $U \subseteq M$ and it is also the intersection of all cosets containing U . We remark that in the particular case where M is a group our notion of coset coincides with the usual one.

Relational Structures. A *relational signature* Σ is a set of symbols where each $R \in \Sigma$ has a number $\text{ar}(R)$ associated to it, called its arity. In this work we assume relational signatures to always be finite.

Given a relational signature Σ , a Σ -structure $\mathbf{A}$ consists of a set A , called the universe of $\mathbf{A}$, and a set $R(\mathbf{A}) \subseteq A^{\text{ar}(R)}$ for each symbol $R \in \Sigma$, which is called the interpretation of R on $\mathbf{A}$. A *relational structure* is a Σ -structure for some relational signature Σ .

A homomorphism f from a relational structure $\mathbf{A}$ to another relational structure $\mathbf{B}$ with the same signature Σ is a map $f : A \rightarrow B$ satisfying $f(R(\mathbf{A})) \subseteq R(\mathbf{B})$ for each $R \in \Sigma$. We write $f : \mathbf{A} \rightarrow \mathbf{B}$ to denote that f is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$, and $\mathbf{A} \rightarrow \mathbf{B}$ to denote that there exists such a homomorphism without specifying one.

Let Σ_{mon} be the relational signature consisting of a unary symbol R_e and a ternary symbol $R_{\otimes}$. Given a monoid M , there is a natural Σ_{mon} -structure $\mathbf{M}$ associated to it whose universe is M , and where $R_e(\mathbf{M}) = \{e_M\}$, and $R_{\otimes}(\mathbf{M}) = \{(a, b, c) \in M^3 \mid ab = c\}$. Note however that not every Σ_{mon} -structure is associated to a monoid. Observe that given two monoids M, N , a map $f : M \rightarrow N$ is a homomorphism from M to N if, and only if, it is a homomorphism from $\mathbf{M}$ to $\mathbf{N}$. If M is a monoid and $R \subseteq M^n$, we denote by (M, R) the *relational structure* with the three relations $R_e(\mathbf{M})$, $R_{\otimes}(\mathbf{M})$, and R .

Promise Constraint Satisfaction Problems. Let $\mathbf{A}, \mathbf{B}$ be relational structures such that $\mathbf{A} \rightarrow \mathbf{B}$. The *promise constraint satisfaction problem* $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is the computational problem of distinguishing, for a given input $\mathbf{X}$, between the following two cases: (1) there exists a homomorphism $\mathbf{X} \rightarrow \mathbf{A}$, and (2) there does not exist any homomorphism $\mathbf{X} \rightarrow \mathbf{B}$. Instances falling into the first case are called Yes-instances, instances falling into the second case are called No-instances. We write $\text{CSP}(\mathbf{A})$ for the problem $\text{PCSP}(\mathbf{A}, \mathbf{A})$.

Note that by our convention, if $\mathbf{M}$ and $\mathbf{N}$ are expansions of monoids, the input to $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is *not* assumed to be an expansion of a monoid, but a relational structure. This differs, e.g., from [5, 7], where the CSPs of expansions of algebraic structures are investigated and where the inputs are assumed to be themselves algebras. An input $\mathbf{X}$ to $\text{PCSP}(\mathbf{M}, \mathbf{N})$ can then be seen as a finite set of equations of the form $s(x_1, \dots, x_n) = t(x_1, \dots, x_n)$ where s, t are words over $\{x_1, \dots, x_n\}$, together with additional constraints involving the extra relations.

3 Dichotomy and Finite Tractability

We are now in position to state our main result (Theorem 1.1) formally.

THEOREM 3.1. *Let $\mathbf{M} = (M, R(M))$, $\mathbf{N} = (N, R(N))$ be expansions of monoids such that $\mathbf{M} \rightarrow \mathbf{N}$. Assume that M is finitely generated and N is finite. Then, either (1) there exists a homomorphism*

⁵This notion of inverse is less general than the commonly adopted notion of inverse for semigroups, where it is only required that $aba = a$ and $bab = b$ (e.g., [21]). In that setting inverses are not necessarily unique.

$h: \mathbf{M} \rightarrow \mathbf{N}$ whose image is a regular commutative monoid and such that $\text{Cos}(h(R(\mathbf{M}))) \subseteq R(\mathbf{N})$, in which case $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is solvable in polynomial time, or (2) $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is NP-hard.

We note that our result also captures templates where $\mathbf{M} = (M, R_1(M), \dots, R_\ell(M))$ and $\mathbf{N} = (N, R_1(N), \dots, R_\ell(N))$ are expansions of monoids by finitely many non-empty relations. To see this, observe that $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is polynomial time equivalent to $\text{PCSP}(\mathbf{M}', \mathbf{N}')$ where $\mathbf{M}', \mathbf{N}'$ are the expansions of M, N by the relations $R(M), R(N)$ obtained by concatenating the tuples in $R_1(M), \dots, R_\ell(M)$ and in $R_1(N), \dots, R_\ell(N)$ respectively.

Interestingly, our proof of NP-hardness in Theorem 3.1 makes use of the algebraic approach to promise constraint satisfaction [4] albeit in a setting where the left-hand side structure is infinite and the right-hand side structure is finite. In [4], many of the statements for finite structures hold for templates where the *right-hand* side template is infinite and the *left-hand* side template is finite, however the direction that we take here is essentially new and works for arbitrary algebraic structures (i.e., not only groups and monoids) under the natural condition that the structures and their powers are finitely generated. Such structures with finitely generated left-hand side and finite right-hand side also appear in the study of (non-promise) constraint satisfaction problems for infinite structures with a large automorphism group [27].

The polynomial-time algorithm that we use to prove Theorem 3.1 is inspired by the work of [24] on finite regular commutative monoids. We show that the CSP of an expansion of a finitely generated regular commutative monoid by any finite number of cosets is solvable in polynomial time. Although the monoid M in Theorem 3.1 is only assumed to be finitely generated, we can apply our algorithm by showing that every finitely generated monoid M admits a *commutative regularization* $(M^{\text{r.c.}}, \pi^{\text{r.c.}})$, as given by the following result. This fact and this proof might be folklore; we include the proof in the appendix (see Appendix B) for the convenience of the reader.

Lemma 3.2. *Let M be a monoid. Then there exists a regular commutative monoid $M^{\text{r.c.}}$ and a homomorphism $\pi_M^{\text{r.c.}}: M \rightarrow M^{\text{r.c.}}$ satisfying that for any monoid homomorphism $f: M \rightarrow N$ where N is commutative and regular, there is a homomorphism $f': M^{\text{r.c.}} \rightarrow N$ such that $f = f' \circ \pi_M^{\text{r.c.}}$. Moreover, if $S \subseteq M$ is a set generating M , then $\pi_M^{\text{r.c.}}(S)$ generates $M^{\text{r.c.}}$.*

Let $\mathbf{M}^{\text{r.c.}} = (M^{\text{r.c.}}, \text{Cos}(\pi^{\text{r.c.}}(R(\mathbf{M}))))$. Then $\text{CSP}(\mathbf{M}^{\text{r.c.}})$ is solvable in polynomial time by our algorithm briefly discussed above, cf. Section 6 for all details. Theorem 3.1 together with the universal property of $\mathbf{M}^{\text{r.c.}}$ imply that this structure serves as a classifier for the complexity of problems of the form $\text{PCSP}(\mathbf{M}, _)$.

Corollary 3.3. *Assume $P \neq \text{NP}$. Suppose that $\mathbf{M} = (M, R(\mathbf{M})), \mathbf{N} = (N, R(\mathbf{N}))$ are expansions of monoids such that M is finitely generated and N is finite. Then $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is solvable in polynomial time if, and only if, there exists a homomorphism $\mathbf{M}^{\text{r.c.}} \rightarrow \mathbf{N}$.*

Another implication of Theorem 3.1 is that a similar dichotomy holds if the templates $\mathbf{M}, \mathbf{N}$ are expansions of *groups* instead of monoids. In that case, the terms appearing in an input to $\text{PCSP}(\mathbf{M}, \mathbf{N})$ can also use the inverse symbol. However, up to introducing a new variable x_- and constraints $x_-x = e = xx_-$ for each variable x , we reduce “group instances” to “monoid instances”. Moreover, a group is generated by U if, and only if, it is generated as a monoid by $U \cup U^{-1}$.

Corollary 3.4. *Let $\mathbf{G} = (G, R^G), \mathbf{H} = (H, R^H)$ be expansions of groups such that there exists a homomorphism $\mathbf{G} \rightarrow \mathbf{H}$. Assume that G is finitely generated and H is finite. Then, either (1) there exists a homomorphism $h: \mathbf{G} \rightarrow \mathbf{H}$ with a commutative image such that $[h(R^G)] \subseteq R^H$, in which case $\text{PCSP}(\mathbf{G}, \mathbf{H})$ is solvable in polynomial time, or (2) $\text{PCSP}(\mathbf{G}, \mathbf{H})$ is NP-hard.*

The problem $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is said to be *finitely tractable* if there exist a finite structure $\mathbf{C}$ and homomorphisms $\mathbf{A} \rightarrow \mathbf{C}$ and $\mathbf{C} \rightarrow \mathbf{B}$ such that $\text{CSP}(\mathbf{C})$ is solvable in polynomial time. The

study of finite tractability in promise constraint satisfaction is important, as it highlights some essential differences between the non-promise setting (where $A = B$). It is known that there exist problems $\text{PCSP}(A, B)$ with A, B finite that are solvable in polynomial time but that are not finitely tractable [1, 4]. A simple consequence of Theorem 3.1 is that such problems do not appear in the class of problems under consideration.

Corollary 3.5. *Assume $P \neq NP$. Suppose that M, N are expansions of monoids M, N such that M is finitely generated and N is finite. If $\text{PCSP}(M, N)$ is solvable in polynomial time, then it is finitely tractable.*

4 The Structure of Regular Commutative Monoids

In this section we introduce some elementary facts about monoids that will be used in our main results. We refer to Appendix A for the short proofs, which we include for completeness although they might be known. We start with the following characterization of regular commutative monoids, which is a consequence of Green's theorem.

Proposition 4.1 ([21, Proposition 4.1.1]). *Let M be a commutative monoid. Then M is regular if and only if, each of its $\mathcal{H}$ -classes contains an idempotent element.*

Given a commutative monoid M , we write M_I for its set of idempotent elements. Proposition 4.1 implies that if M is regular we can write $M = \bigsqcup_{d \in M_I} H_d$, where H_d is the $\mathcal{H}$ -class of $d \in M_I$, which also is the maximal subgroup containing this element, by 2.1. The set M_I has the following nice properties.

Lemma 4.2. *Let M be a commutative monoid and $M_I \subseteq M$ be the subset of its idempotent elements. Then, (1) M_I is a semilattice, and (2) if M is regular and finitely generated, then M_I is finite.*

It is well-known that for any finite group G there is a constant $C \in \mathbb{N}$ such that $a^C = e_G$ for all $a \in G$. For monoids the following analogue holds.

Lemma 4.3. *Let M be a finite monoid. There is an integer $C > 1$ such that a^C is idempotent for all $a \in M$. If $a \in M$ is a group element, then it also holds that $a^{C-1} = a^{-1}$. Additionally, given $a \in M$, there is a unique idempotent element d_a that satisfies $d_a = a^n$ for some $n \in \mathbb{N}$.*

This way, given an element a in a finite monoid M , we write d_a for the only idempotent element that can be expressed as a^n for some $n \in \mathbb{N}$.

Lemma 4.4. *Let M be a finite commutative monoid, $M_I \subseteq M$ its subset of idempotent elements, and $M_{\dagger} \subseteq M$ its subset of group elements. The following hold: (1) The map $\pi_I : M \rightarrow M_I$ given by $a \mapsto d_a$ is a surjective monoid homomorphism satisfying $\pi_I \circ \pi_I = \pi_I$. (2) $M_{\dagger}$ is a regular submonoid of M . (3) The map $\pi_{\dagger} : M \rightarrow M_{\dagger}$ given by $a \mapsto d_a a$ is a surjective monoid homomorphism satisfying $\pi_{\dagger} \circ \pi_{\dagger} = \pi_{\dagger}$.*

Given a finite commutative monoid M and an element $a \in M$, we write $a_{\dagger}$ for the product ad_a and, if $S \subseteq M$, we write $S_{\dagger}$ for $\{a_{\dagger} \mid a \in S\}$.

Lemma 4.5. *Let M be a monoid, and N be a commutative monoid. Let F be a finite family of monoid homomorphisms $f : M \rightarrow N$, and $g : M \rightarrow N$ be the homomorphism given by $g(a) = \prod_{f \in F} f(a)$.⁶ Then for each $S \subseteq M$, it holds that $\text{Cos}(g(S)_{\dagger}) \subseteq \bigotimes_{f \in F} \text{Cos}(f(S)_{\dagger})$.*

⁶Observe that this notation is well-defined because products commute in a commutative monoid.

5 Hardness

We prove here the hardness side of Theorem 3.1. That is, we show that in the absence of a suitable homomorphism $\mathbf{M} \rightarrow \mathbf{N}$ we have that $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is NP-hard. For this, we first introduce an additional decision problem from [4] whose NP-hardness can be proved under a certain algebraic condition ($\star$), then exhibit a reduction from this problem to $\text{PCSP}(\mathbf{M}, \mathbf{N})$ which holds beyond the setting of groups and monoids, and finally show that ($\star$) holds for the templates under consideration in this paper.

Polymorphisms. Given a relational structure $\mathbf{A}$ and a number $n \in \mathbb{N}$, the n -th Cartesian power of $\mathbf{A}$, denoted $\mathbf{A}^n$, is defined in the natural way. That is, the universe of $\mathbf{A}^n$ is A^n , and for each symbol R in the signature, $R(\mathbf{A}^n)$ is the set of tuples $(\mathbf{a}_1, \dots, \mathbf{a}_{\text{ar}(R)})$ where $\mathbf{a}_i \in R(\mathbf{A})$ for each $i \in [\text{ar}(R)]$.

Let $\mathbf{A} \rightarrow \mathbf{B}$ be two relational structures. Given $n \in \mathbb{N}$, a n -ary *polymorphism* of the pair $(\mathbf{A}, \mathbf{B})$ is a homomorphism $f : \mathbf{A}^n \rightarrow \mathbf{B}$. We denote by $\text{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$ the set of n -ary polymorphisms of $(\mathbf{A}, \mathbf{B})$, and use $\text{Pol}(\mathbf{A}, \mathbf{B})$ to represent the disjoint union $\bigsqcup_{n \in \mathbb{N}} \text{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$.

Given $n, m \in \mathbb{N}$, $f \in \text{Pol}^{(n)}(\mathbf{A}, \mathbf{B})$, and $\sigma : [n] \rightarrow [m]$ the element $f^\sigma \in \text{Pol}^{(m)}(\mathbf{A}, \mathbf{B})$ is given by $f^\sigma(\mathbf{a}) = f(\mathbf{a} \circ \sigma)$ for each $\mathbf{a} \in A^m$. In this situation we say that f^σ is a *minor* of f . The *unary minor* of f is the polymorphism f^τ , where τ is the constant map from $[n]$ to $[1]$. In this situation, it is easy to see that $(f^{\sigma_1})^{\sigma_2} = f^{\sigma_2 \circ \sigma_1}$ for all suitable maps σ_1, σ_2 , and that $f^\sigma = f$ whenever $f \in \text{Pol}^{(n)}$, and σ is the identity over $[n]$ for some $n \in \mathbb{N}$. This sort of algebraic structure has been called a *minion* in the literature [4].

Minor Conditions. A minor condition Φ is a tuple $(U, V, E, (\phi_{(u,v) \in E}))$ where U, V are finite disjoint sets where each element $u \in U \sqcup V$ has an arity $\text{ar}(u) \in \mathbb{N}$ associated to it, $E \subseteq U \times V$, and $\phi_{u,v}$ is a map from $[\text{ar}(u)]$ to $[\text{ar}(v)]$ for each $(u, v) \in E$. We say that Φ is *trivial* if for each $u \in V \sqcup U$ there is an index i_u satisfying that $\phi_{u,v}(i_u) = i_v$ for each $(u, v) \in E$. Given relational structures $\mathbf{A} \rightarrow \mathbf{B}$, we say that Φ is *satisfiable* over $\text{Pol}(\mathbf{A}, \mathbf{B})$ if for each $u \in V \sqcup U$ there is an element $f_u \in \text{Pol}^{(\text{ar}(u))}(\mathbf{A}, \mathbf{B})$ and there exists $f_v \in \text{Pol}^{(\text{ar}(v))}(\mathbf{A}, \mathbf{B})$ satisfying $f_u^{\phi_{u,v}} = f_v$ for each $(u, v) \in E$. It can be seen that if Φ is trivial, then it is satisfiable over $\text{Pol}(\mathbf{A}, \mathbf{B})$ for any choice of $\mathbf{A} \rightarrow \mathbf{B}$.

Let $\mathbf{A}, \mathbf{B}$ be relational structures such that $\mathbf{A} \rightarrow \mathbf{B}$, and let $\ell \in \mathbb{N}$. The *promise minor condition problem*, denoted by $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$,⁷ is the computational problem of distinguishing, given a finite minor condition Σ with symbols of arity at most ℓ , between the following two cases: (1) Σ is trivial, and (2) Σ is not satisfiable in $\text{Pol}(\mathbf{A}, \mathbf{B})$.

Let $\mathbf{A} \rightarrow \mathbf{B}$ be relational structures, and $\mathcal{N} \subseteq \text{Pol}(\mathbf{A}, \mathbf{B})$ be a subset. A *selection function* over $\mathcal{N}$ is a function $\mathcal{I}$ with domain $\mathcal{N}$ such that for every $n \geq 1$, and every $f \in \mathcal{N}$ of arity n , $\mathcal{I}(f)$ is a subset of $[n]$, and such that whenever $f \in \text{Pol}^{(n)}(\mathbf{A}, \mathbf{B}) \cap \mathcal{N}$ and $\sigma : [n] \rightarrow [m]$ are such that $f^\sigma \in \mathcal{N}$, then $\sigma(\mathcal{I}(f)) \cap \mathcal{I}(f^\sigma) \neq \emptyset$. We say that $\mathcal{I}$ is *bounded* if $\{|\mathcal{I}(f)| : f \in \mathcal{N}\} \subseteq \mathbb{N}$ is bounded. The proof of Theorem 5.21 in [4] implies the following result (cf. also [6] and [25, Corollary 4.5]).

THEOREM 5.1 ([4]). *Let $\mathbf{A} \rightarrow \mathbf{B}$ be relational structures. Suppose that*

$\text{Pol}(\mathbf{A}, \mathbf{B})$ is the union of subsets $\mathcal{M}_1, \dots, \mathcal{M}_k$ such that for each $i \in [k]$ there exists a bounded selection function for $\mathcal{M}_i$. ($\star$)

Then there exists some $\ell \in \mathbb{N}$ such that $\text{PMC}_{\ell'}(\mathbf{A}, \mathbf{B})$ is NP-hard for all $\ell' \geq \ell$.

Algebraic Approach in the Finitely Generated Case. For *finite* relational structures, the problem $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$ is related to $\text{PCSP}(\mathbf{A}, \mathbf{B})$ by the following result, which lies at the heart of the so-called algebraic approach to promise constraint satisfaction.

⁷For the sake of readability, we differ slightly here from the notation introduced in [4].

THEOREM 5.2 ([4]). *Let $A \rightarrow B$ be finite relational structures. Then $\text{PMC}_\ell(A, B)$ and $\text{PCSP}(A, B)$ are equivalent under logspace reductions for all large enough ℓ .*

Let M, N be monoids, where M is finitely generated and N is finite, and $\mathbf{M}, \mathbf{N}$ be expansions of M, N by a single relation $R(M), R(N)$ such that $\mathbf{M} \rightarrow \mathbf{N}$. The hardness side of Theorem 3.1 is shown by reducing $\text{PMC}_\ell(\mathbf{M}, \mathbf{N})$ to $\text{PCSP}(\mathbf{M}, \mathbf{N})$ and then showing that $\text{PMC}_\ell(\mathbf{M}, \mathbf{N})$ is NP-hard for ℓ sufficiently large. The fact that $\mathbf{M}$ may be infinite means that Theorem 5.2 does not yield the reduction we are looking for. In general, Theorem 5.2 is known to fail for infinite structures, since $\text{PMC}_L(A, B)$ is always in NP (since the computationally harder problem to check whether a given minor condition Σ is trivial or not is itself in NP), while the (P)CSP of infinite structures can have arbitrary high complexities, even under strong structural assumptions [16].

Our reduction is more general and applies to a wider context than for groups and monoids. In order to present it, we will need the standard notions from universal algebra to put our result in its natural habitat. An *algebraic signature* τ is a list of function symbols, each having a finite arity $n \in \mathbb{N}$. As for relational structures, we assume in this work that algebraic signatures are finite. A τ -*algebra* A is a tuple consisting of a set A together with, for each n -ary operation symbol $f \in \tau$, an operation $f^A: A^n \rightarrow A$. The operations f^A are called the basic operations of the algebra. A *term* (in the signature τ) is a formal operation obtained by composing the symbols in τ in a way that respects the arities of the symbols. Every term t can be evaluated naturally in any τ -algebra, yielding an operation t^A . For example, if τ consists of a single binary symbol $\cdot$, then $t_1(x, y, z) = (x \cdot y) \cdot z$ and $t_2(x, y, z) = x \cdot (y \cdot z)$ are both terms. Any set with a binary operation $A = (A, \cdot^A)$ is a τ -algebra. We have $t_1^A = t_2^A$ in any algebra where $\cdot^A$ is associative, e.g., in monoids and groups.

We can treat any algebra as a relational structure, by replacing a basic operation of arity n by a relation of arity $n + 1$. The notion of homomorphism between algebras can be defined as homomorphisms between the associated relational structures. Alternatively, a homomorphism between two τ -algebras A and B is a function $h: A \rightarrow B$ such that

$$h(f^A(a_1, \dots, a_n)) = f^B(h(a_1), \dots, h(a_n))$$

holds for every symbol f of arity n in τ and every $a_1, \dots, a_n \in A$. By a simple induction, one sees that this equality must also hold for every term t .

For $\ell \geq 1$, the ℓ th *power* of A , denoted by A^ℓ , is the algebra with base set A^ℓ and whose fundamental operations are defined component-wise, i.e., by

$$f^{A^\ell}(\mathbf{a}^1, \dots, \mathbf{a}^n) = (f^A(a_1^1, \dots, a_1^n), \dots, f^A(a_L^1, \dots, a_L^n)),$$

where $\mathbf{a}_j^i$ denotes the j th component of the L -tuple $\mathbf{a}^i$. This coincides with the definition of products for groups and monoids.

We say that an algebra A is *finitely generated* if there exists a finite subset $A' = \{a_1, \dots, a_r\} \subseteq A$, called a *generating set*, such that every $a \in A$ is of the form $t^A(a_1, \dots, a_r)$ for some term t .

In general, an algebra A can be finitely generated while its powers are not. A simple example is the algebra with a single unary function $N = (\mathbb{N}; s^N)$ where $s^N(n) = n + 1$. One sees that $\{1\}$ is a generating set of N , while any generating set of N^2 must contain the elements $(1, 1), (1, 2), (1, 3), \dots$. However, if A is a finitely generated group (or a monoid) with identity element e , then A^ℓ is finitely generated for all $L \in \mathbb{N}$. Indeed, if $G = \{g_1, \dots, g_r\}$ is a generating set of A , then one sees that elements of the form $(e, \dots, e, g_i, e, \dots, e)$ for all i and all possible positions of g_i generate A^ℓ .

Finally, we note the following. Suppose that A, B are τ -algebras and $U = \{\alpha_1, \dots, \alpha_r\}$ is a generating set of A . Let $h: U \rightarrow B$ be a function. Then h can be extended to a homomorphism $\tilde{h}: A \rightarrow B$ if, and only if, for all terms s, t of arity r , we have that if $s^A(\alpha_1, \dots, \alpha_r) = t^A(\alpha_1, \dots, \alpha_r)$, then $s^B(h(\alpha_1), \dots, h(\alpha_r)) = t^B(h(\alpha_1), \dots, h(\alpha_r))$. Indeed, we can simply define $\tilde{h}(s^A(\alpha_1, \dots, \alpha_r))$

as $s^B(h(\alpha_1), \dots, h(\alpha_r))$, which is well-defined by assumption and is a total function by assumption that G generates A .

Proposition 5.3. *Let A, B be algebras such that all finite powers of A are finitely generated and B is finite. Fix $m \in \mathbb{N}$ and $R^A \subseteq A^m, R^B \subseteq B^m$ such that $\mathbf{A} = (A, R^A)$ admits a homomorphism to $\mathbf{B} = (B, R^B)$. Then there is a logspace reduction from $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$ to $\text{PCSP}(\mathbf{A}, \mathbf{B})$ for any $\ell \in \mathbb{N}$.*

PROOF. Consider an instance Σ of $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$. Without loss of generality, we can assume that all symbols in Σ have arity ℓ (otherwise lower-arity symbols are padded by dummy variables). We define an instance $\mathbf{I} = \mathbf{I}_\Sigma$ of $\text{PCSP}(\mathbf{A}, \mathbf{B})$ that is constructible in logspace from Σ . Let U be a finite generating set of A^L . We write $U = \{\alpha_1, \dots, \alpha_r\}$. Without loss of generality, we can assume that for every $\alpha \in U$ and $\sigma: [L] \rightarrow [L]$, we have $\alpha \circ \sigma \in U$ as well.

For each $x \in X$ and $\alpha \in U$, let $x(\alpha)$ be a variable of $\mathbf{I}$. Our goal is to obtain that in every solution of $\mathbf{I}$ in $\mathbf{B}$, the values $x(\alpha)$ define a map $U \rightarrow B$ that extends to a homomorphism $\tilde{x}: A^L \rightarrow B$. Note that there are only finitely many functions $f: U \rightarrow B$. For each such function that does not extend to a homomorphism $\tilde{f}: A^L \rightarrow B$, there must exist a pair (s, t) of terms such that $s^{A^L}(\alpha_1, \dots, \alpha_r) = t^{A^L}(\alpha_1, \dots, \alpha_r)$ while $s^B(f(\alpha_1), \dots, f(\alpha_r)) \neq t^B(f(\alpha_1), \dots, f(\alpha_r))$. Make a finite list $(s_1, t_1), \dots, (s_k, t_k)$ of such pairs of terms corresponding to maps $U \rightarrow B$ that do not extend to a homomorphism. For each such pair, and each variable x , add the constraint

$$s(x(\alpha_1), \dots, x(\alpha_r)) = t(x(\alpha_1), \dots, x(\alpha_r))$$

to $\mathbf{I}$. Formally, for each variable x and pair (s, t) , this involves the introduction of a bounded number of variables for each subterm of s and t and corresponding equations. All in all, in every solution to $\mathbf{I}$, we are guaranteed that $x: U \rightarrow B$ extends to a homomorphism $\tilde{x}: A^L \rightarrow B$.

We now additionally want to make sure that the extension $\tilde{x}: A^L \rightarrow B$ of x is a homomorphism $A^L \rightarrow \mathbf{B}$, i.e., that it is an element of $\text{Pol}(\mathbf{A}, \mathbf{B})$. Note that there are only finitely many homomorphisms $A^L \rightarrow B$. Let $f_1, \dots, f_k$ be those homomorphisms that are not homomorphisms $A^L \rightarrow \mathbf{B}$. For each $i \in [k]$, proceed as follows. First, pick elements $r^1, \dots, r^L \in R^A$ such that $f_i(r^1, \dots, r^L) \notin R^B$. Each $(r_j^1, \dots, r_j^L)$ for $j \in [r]$ is equal to $t_j^{A^L}(\alpha_1, \dots, \alpha_r)$ for some term t_j . Add to $\mathbf{I}$ the constraint

$$(t_1(x(\alpha_1), \dots, x(\alpha_r)), \dots, t_{\text{ar}(R)}(x(\alpha_1), \dots, x(\alpha_r))) \in R.$$

Finally, we want that the resulting homomorphisms satisfy the minor condition Σ . Let $x = y^\sigma$ be a minor identity in Σ . We add the constraints

$$x(\alpha) = y(\alpha \circ \sigma)$$

for every $\alpha \in U$, where we recall that $\alpha \circ \sigma$ is also an element of U . We prove that if $\tilde{x}$ and $\tilde{y}$ are the resulting homomorphisms $A^L \rightarrow B$, then $\tilde{x}(a_1, \dots, a_\ell) = \tilde{y}(a_{\sigma(1)}, \dots, a_{\sigma(N)})$ holds for all $a_1, \dots, a_\ell \in A$. There exists a term t such that $t^{A^L}(\alpha_1, \dots, \alpha_r) = (a_1, \dots, a_\ell)$. Thus, $t^A(\alpha_{1i}, \dots, \alpha_{ri}) = a_i$ for all $i \in [L]$. It follows that $t^A(\alpha_{1\sigma(i)}, \dots, \alpha_{r\sigma(i)}) = a_{\sigma(i)}$ for all $i \in [L]$ and therefore $t^{A^L}(\alpha_1 \circ \sigma, \dots, \alpha_r \circ \sigma) = (a_{\sigma(1)}, \dots, a_{\sigma(N)})$. By definition of $\tilde{x}$ and $\tilde{y}$, we have

$$\tilde{x}(a_1, \dots, a_\ell) = t^B(x(\alpha_1), \dots, x(\alpha_r))$$

and

$$\tilde{y}(a_{\sigma(1)}, \dots, a_{\sigma(N)}) = t^B(y(\alpha_1 \circ \sigma), \dots, y(\alpha_r \circ \sigma)).$$

Since $\mathbf{I}$ contains the constraint $x(\alpha_i) = y(\alpha_i \circ \sigma)$ for all $i \in [r]$, we obtain that $\tilde{x}(a_1, \dots, a_\ell) = \tilde{y}(a_{\sigma(1)}, \dots, a_{\sigma(N)})$ is satisfied.

It remains to prove that this is a correct reduction. We have already described above that if $\mathbf{I}$ is not a No-instance of $\text{PCSP}(\mathbf{A}, \mathbf{B})$, i.e., if there is a solution ξ to $\mathbf{I}$, then for each $x \in X$ the function $U \rightarrow B$ defined by $\alpha \mapsto \xi(x(\alpha))$ extends to an element $\tilde{x} \in \text{Pol}(\mathbf{A}, \mathbf{B})$ showing that Σ is

satisfiable in $\text{Pol}(\mathbf{A}, \mathbf{B})$ and therefore not a No-instance of $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$. Therefore, No-instances of $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$ are mapped to No-instances of $\text{PCSP}(\mathbf{A}, \mathbf{B})$. Conversely, suppose that Σ is a Yes-instance of $\text{PMC}_\ell(\mathbf{A}, \mathbf{B})$. Then there is a map $x \mapsto i_x$ that satisfies $\sigma(i_x) = i_y$ for any constraint of the form $x^\sigma = y$ in Σ . Then, the map that sends each variable $x(\alpha_j)$ to the i_x -th component of α_j is a solution over $\mathbf{A}$ to our instance. $\square$

Polymorphisms of Expansions of Monoids. For the NP-hardness part of Theorem 3.1, we are left to show in this section that if there does not exist a homomorphism $h: \mathbf{M} \rightarrow \mathbf{N}$ whose image is a regular commutative submonoid of N and such that $\text{Cos}(h(R(M))) \subseteq R(N)$, then the requirement in Theorem 5.1 is met.

The following two auxiliary results show that in a finite commutative monoid M , for a sufficiently large ℓ and any set $U \subseteq M$, any product of the form $\prod_{i=1}^\ell a_i$, where $a_i \in \text{Cos}(U_+)$ for each i , can be obtained alternatively as $\prod_{i=1}^\ell b_i$, where $b_i \in U$ for each i .

Lemma 5.4. *Let M be a finite regular commutative monoid. There exists $\ell(M) \in \mathbb{N}$ such that for all $n \geq \ell(M)$ and for all $U \subseteq M$ we have $\text{Cos}(U)^{\otimes n} = U^{\otimes n}$.*

PROOF. Recall the decomposition of regular commutative monoids described in Section 4. Let M_I be the set of idempotent elements in M , and H_d the maximal subgroup of M containing d for each $d \in M_I$.

For all n , we have $|U^{\otimes n}| \leq |U^{\otimes n+1}|$. To see this, we simply show that $|U^{\otimes n} \cap H_d| \leq |U^{\otimes n+1} \cap H_d|$ for each $d \in M_I$. Suppose that $|U^{\otimes n} \cap H_d| \geq 1$ (otherwise we are done). Then there must be some element $a \in U$ with $d \leq a$. It follows that $ad \in H_d$, so $ad \in H_d$. This way,

$$a \otimes H_d = a \otimes (d \otimes H_d) = ad \otimes H_d \subseteq H_d,$$

and therefore $a \otimes (U^{\otimes n} \cap H_d) \subseteq U^{\otimes n+1} \cap H_d$. Moreover, $|a \otimes (U^{\otimes n} \cap H_d)| = |U^{\otimes n} \cap H_d|$ holds, proving the statement. Thus, since M is finite, there exists a natural number $\ell(M)$ such that $|U^{\otimes n}| = |U^{\otimes \ell(M)}| = \lambda$ holds for all $n \geq \ell(M)$.

Now let us show that $|U \otimes (U^{-1} \otimes U)^{\otimes n}| \geq |U^{\otimes n}|$ for all $n \in \mathbb{N}$. To do this we define an injective map from $U^{\otimes n} \cap H_d$ to $U \otimes (U^{-1} \otimes U)^{\otimes n} \cap H_d$ for each $d \in M_I$. Suppose that $U^{\otimes n} \cap H_d$ is not empty. In particular, there must be some element $a \in U$ satisfying $d \leq a$. We claim the map $b \mapsto a^{1-n}b$ is an injective map from $U^{\otimes n} \cap H_d$ to $U \otimes (U^{-1} \otimes U)^{\otimes n} \cap H_d$. This is a well-defined map: if $b \in U^{\otimes n} \cap H_d$, then we can write $b = s_1 \cdots s_n$ for some $s_1, \dots, s_n \in U$, and then $a^{1-n}b = a(a^{-1}s_1) \cdots (a^{-1}s_n)$, which belongs to $U \otimes (U^{-1} \otimes U)^{\otimes n}$. The fact that $b \in H_d$ and $b \leq a$ implies $a^{1-n}b \in H_d$. Finally, let us show that the map $b \mapsto a^{1-n}b$ is injective. Observe that $b = db$, for all $b \in H_d$, so $a^{1-n}b = a^{1-n}b'$ if, and only if, $(a^{1-n}d)b = (a^{1-n}d)b'$ for each $b, b' \in U \otimes (U^{-1} \otimes U)^{\otimes n} \cap H_d$. However, the fact that $d \leq a$ implies that $a^{1-n}d \in H_d$, so $(a^{1-n}d)b = (a^{1-n}d)b'$ holds if, and only if, $b = b'$.

Now we can see that $\text{Cos}(U) = U^{\otimes n}$ for some n . Indeed, $\text{Cos}(U) = U \otimes (U^{-1}U)^{\otimes n}$ for all large enough n , so in particular $|\text{Cos}(U)| \geq |U^{\otimes n}| = \lambda$. Since $U^{-1} \subseteq U^{\otimes k_M-1}$ for some constant k_M , we get $\text{Cos}(U) \subseteq U^{\otimes(1+nk_M)}$ and this set has size at most λ , so the two sets are equal.

Finally, for all n , we trivially have $U^{\otimes n} \subseteq \text{Cos}(U)^{\otimes n}$ and since for $n \geq \ell(M)$ the two sets have size λ , we get $U^{\otimes n} = \text{Cos}(U)^{\otimes n}$. $\square$

Lemma 5.5. *Let M be a finite commutative monoid. There exists $\ell(M) \in \mathbb{N}$ such that for all $n \geq \ell(M)$ and for all $U \subseteq M$ we have $\text{Cos}(U_+)^{\otimes n} \subseteq U^{\otimes n}$.*

PROOF. Let $k \in \mathbb{N}$ be such that a^k is idempotent for all $a \in M$, as given by Lemma 4.3. Let λ be a constant witnessing Lemma 5.4 with respect to M_+ . Let $\ell = \max(\lambda, 2k|M|)$. We claim that ℓ satisfies the lemma's statement. Fix $n \geq \ell$. Then, it is enough to show that $\text{Cos}(U_+)^{\otimes n} \subseteq U^{\otimes n}$.

Let $a \in \text{Cos}(U_{\dagger})^{\otimes n}$. By Lemma 5.4, $a = \prod_{i \in [n]} (b_i)_{\dagger}$ for some $b_1, \dots, b_n \in U$. Another way of writing this identity is $a = \prod_{b \in U'} b_{\dagger}^{n_b}$, where $U' \subseteq U$, and the n_b denote positive integers satisfying $\sum_{b \in U'} n_b = n$. As $n \geq 2k|M|$, and $|U'| \leq |M|$, there are integers $(m_b)_{b \in U'}$ such that $\sum_{b \in U'} m_b = n$ that satisfy both $m_b = n_b \pmod k$, and $m_b \geq k$ for all $b \in U'$. Let us construct this sequence. For each $b \in U'$, let $0 < n'_b \leq k$ be the only integer satisfying $n'_b = n_b \pmod k$. Define $m'_b = n'_b + k$ for each $b \in U'$. Define

$$\tau = \sum_{b \in U'} m'_b - \sum_{b \in U'} n_b.$$

As $n \geq 2k|M|$ and $m'_b \leq 2k$ for each $b \in U'$, we have that $\tau \leq 0$. Moreover, by construction $m'_b = n_b \pmod k$ for each $b \in U'$. This way, $\tau = 0 \pmod k$. In order to obtain the desired $(m_b)_{b \in U'}$, define $m_b = m'_b + \tau$ for one $b \in U'$, and $m_c = m'_c$ for all $c \in U' \setminus \{b\}$. This way,

$$a = \prod_{b \in U'} b_{\dagger}^{n_b} = \prod_{b \in U'} b_{\dagger}^{m_b} = \prod_{b \in U'} b^{m_b}.$$

The last term belongs to $U^{\otimes n}$, so this completes the proof. $\square$

With this in hand, we are now able to provide a bounded selection function for the subset of the polymorphisms of $(\mathbf{M}, \mathbf{N})$ with commutative image, and whose unary minor has a regular image.

Let M be a monoid with identity e . We say that two sets, $S, T \subseteq M$ commute if $ab = ba$ for every $a \in S, b \in T$. Let M, N be monoids and $f: M^n \rightarrow N$ be a homomorphism. For $i \in [n]$, we define the homomorphism $f_i: M \rightarrow N$ by $f_i(x) = f(e, \dots, e, x, e, \dots, e)$, where x is at position i . This way, $f(x_1, \dots, x_n) = \prod_{i=1}^n f_i(x_i)$. Moreover, $\text{Im}(f_i)$ and $\text{Im}(f_j)$ commute for each $i \neq j$. We say that $i, j \in [n]$ are f -equivalent if $f_i = f_j$. This clearly defines an equivalence relation on $[n]$ whose equivalence classes are called f -equivalence classes. If I is an f -equivalence class, we define f_I as the homomorphism f_I for an arbitrary $i \in I$.

Observation 5.6. Let M, N be monoids, $f: M^n \rightarrow N$ a homomorphism, and $g = f^{\sigma}$ for some $\sigma: [n] \rightarrow [m]$. Then, $g_i(a) = \prod_{j \in \sigma^{-1}(i)} f_j(a)$ for each $i \in [m]$, $a \in M$.

Observation 5.7. Let $\mathbf{M}, \mathbf{N}$ be expansions of monoids M, N by a single relation $R(M), R(N)$ such that $\mathbf{M} \rightarrow \mathbf{N}$. Then a map $p: M^n \rightarrow N$ is a polymorphism in $\text{Pol}(\mathbf{M}, \mathbf{N})$ if, and only if, p is a monoid homomorphism from M^n to N and $p((R(M))^n) \subseteq p(R(N))$.

Proposition 5.8. Let M, N be monoids with $M \rightarrow N$, where N is finite. Fix $r \in \mathbb{N}$ and relations $R \subseteq M^r, S \subseteq N^r$. Let $\mathcal{M} \subseteq \text{Pol}(M, N)$ be the subset of polymorphisms f such that (1) $\text{Im}(f)$ is commutative, (2) the unary minor h of f has a regular image, and (3) $\text{Cos}(h(R)) \not\subseteq S$. Given $f \in \mathcal{M}$, let $I(f) \subseteq [n]$ be the union of all f -equivalence classes of size smaller than $\ell(N^r)$, where ℓ is given by Lemma 5.5. Suppose that $I(f^{\sigma}) \cap \sigma(I(f)) = \emptyset$ for some $f, f^{\sigma} \in \mathcal{M}$. Then $f(R^n) \not\subseteq S$, where n denotes the arity of f .

PROOF. Let $f, g \in \mathcal{M}$ be polymorphisms satisfying $f^{\sigma} = g$, and $I(g) \cap \sigma(I(f)) = \emptyset$. Suppose the arities of f and g are n and m respectively. Let $J_1, \dots, J_p \subseteq [m]$ be a list of the g -equivalence classes. By assumption $\text{Cos}(h(R)) \not\subseteq S$, where h is the unary minor of g and f . Let $c \in \text{Cos}(h(R)) \setminus S$. By Lemma 4.5, there exist elements $b_i \in \text{Cos}(g_i(R)_{\dagger})$ for each $i \in [m]$ satisfying $\prod_{i \in [m]} b_i = c$.

In this paragraph, we construct a second sequence $(b'_i)_{i \in [m]}$ whose product is also c , but where $b'_i \in g_i(R)$ for each $i \notin I(g)$. First, we define $b'_i = b_i$ for each $i \in I(g)$. Now let $J \not\subseteq I(g)$ be a g -equivalence class. Then

$$\prod_{j \in J} b_j \in \text{Cos}(g_J(R)_{\dagger})^{\otimes |J|}.$$

We know that $|J| \geq \ell(N^r)$, so by Lemma 5.5 we conclude there are elements $\mathbf{b}'_i \in g_J(R)$ for each $i \in J$ satisfying $\prod_{j \in J} \mathbf{b}_j \mathbf{d}_{\mathbf{b}_j} = \prod_{j \in J} \mathbf{b}'_j$. Applying this reasoning to each g -equivalence class $J \not\subseteq I(g)$, we define an element $\mathbf{b}'_i \in g_i(R)$ for each $i \notin I(g)$. This way, $\prod_{i \in [m]} \mathbf{b}'_i = \mathbf{c}$.

We now construct another sequence $(\mathbf{a}_j)_{j \in [n]}$ whose product equals $\mathbf{c}$. For each $i \in I(g)$, by Observation 5.6 it holds that $g_i = \prod_{j \in \sigma^{-1}(i)} f_j$. Thus, by Lemma 4.5,

$$\text{Cos}(g_i(R)_{\dagger}) \subseteq \bigotimes_{j \in \sigma^{-1}(i)} \text{Cos}(f_j(R)_{\dagger}).$$

We use this fact to obtain elements $\mathbf{a}_j \in \text{Cos}(f_j(R)_{\dagger})$ for each $j \in \sigma^{-1}(i)$ satisfying $\prod_{j \in \sigma^{-1}(i)} \mathbf{a}_j = \mathbf{b}'_i$. Similarly, for each $i \in [m] \setminus I(g)$, we have that $\mathbf{b}'_i = g_i(\alpha)$ for some $\alpha \in R$, so we define $\mathbf{a}_j = f_j(\alpha)$ for each $j \in \sigma^{-1}(i)$. This way, $\prod_{j \in \sigma^{-1}(i)} \mathbf{a}_j = \mathbf{b}'_i$ as well. Hence, we have obtained a sequence $(\mathbf{a}_j)_{j \in [n]}$ satisfying $\prod_{j \in [n]} \mathbf{a}_j = \mathbf{c}$. Observe that, according to our construction, if $j \notin \sigma^{-1}(I(g))$, then $\mathbf{a}_j \in f_j(R)$. Finally, we construct a fourth sequence $(\mathbf{a}'_j)_{j \in [n]}$ whose product equals $\mathbf{c}$ and where $\mathbf{a}'_j \in f_j(R)$ for each $j \in [n]$. Given $j \in I(f)$, we define $\mathbf{a}'_j = \mathbf{a}_j$. Observe that by the assumption of the lemma, in this case $j \notin \sigma^{-1}(I(g))$, so $\mathbf{a}'_j \in f_j(R)$. Now let $J \not\subseteq I(f)$ be an f -equivalence class. Then

$$\prod_{j \in J} \mathbf{a}_j \mathbf{d}_{\mathbf{a}_j} \in \text{Cos}(f_J(R)_{\dagger})^{\otimes(|J|)}.$$

By the definition of I , we know that $|J| \geq \ell(N^r)$, so by Lemma 5.5 we conclude there are elements $\mathbf{a}'_i \in f_J(R)$ for each $j \in J$ satisfying $\prod_{j \in J} \mathbf{a}_j \mathbf{d}_{\mathbf{a}_j} = \prod_{j \in J} \mathbf{a}'_j$. Applying this reasoning to each f -equivalence class $J \not\subseteq I(f)$, we define an element $\mathbf{a}'_j \in f_j(R)$ for each $j \notin I(f)$. This way,

$$\prod_{j \in [n]} \mathbf{a}'_j = \left(\prod_{j \in I(f)} \mathbf{a}_j \right) \left(\prod_{j \in [n] \setminus I(f)} \mathbf{a}_j \mathbf{d}_{\mathbf{a}_j} \right) = \prod_{j \in [n]} \mathbf{a}_j = \mathbf{c}.$$

To see the second equality, observe that $\mathbf{c}$ is a group element, and $\mathbf{c} \leq \mathbf{a}_j$ for all $j \in [n]$, so $\mathbf{c} \mathbf{d}_{\mathbf{a}_j} = \mathbf{c}$. Now, by construction $\mathbf{a}'_j = f_j(\alpha_j)$ for some $\alpha_j \in R$, for each $j \in [n]$. In particular, $f(\alpha_1, \dots, \alpha_n) = \mathbf{c}$. However, $f(\alpha_1, \dots, \alpha_n) \in f(R^n)$, and $\mathbf{c} \notin S$, so this proves the result. $\square$

We conclude this section proving the NP-hardness part of our main theorem (Theorem 3.1).

Proposition 5.9. *Let M be a finitely generated monoid, N be a finite monoid, and $R(M) \subseteq M^r$, $R(N) \subseteq N^r$ be relations of arity r satisfying that $\mathbf{M} = (M, R(M))$ maps homomorphically to $\mathbf{N} = (N, R(N))$. Suppose that there is no homomorphism $f : \mathbf{M} \rightarrow \mathbf{N}$ with regular commutative image such that $\text{Cos}(f(R(M))) \subseteq R(N)$. Then $\text{PCSP}(\mathbf{M}, \mathbf{N})$ is NP-hard.*

PROOF. We show that $\text{PMC}_k(\mathbf{M}, \mathbf{N})$ is NP-hard for sufficiently large k . By Proposition 5.3 this proves the result.

We describe $\text{Pol}(\mathbf{M}, \mathbf{N})$ as the union of three subsets (1) the set $\mathcal{M}_1$ of polymorphisms f such that $\text{Im}(f)$ is not a commutative submonoid of N , (2) the set $\mathcal{M}_2$ of polymorphisms f such that $\text{Im}(f)$ is commutative and whose unary minor g satisfies that $\text{Im}(g)$ is not a regular submonoid of N , and (3) $\mathcal{M}_3$, the complement of $\mathcal{M}_1 \cup \mathcal{M}_2$ in $\mathcal{M}$.

Using that N is finite, the proof of [26, Theorem 3] on pages 3:14–3:15 therein shows the existence of a selection function for $\mathcal{M}_1$ and $\mathcal{M}_2$. It remains therefore to construct a selection function for $\mathcal{M}_3$.

Since there is no homomorphism $h : M \rightarrow N$ with regular commutative image and such that $\text{Cos}(h(R(M))) \subseteq R(N)$, all the elements of $\mathcal{M}_3$ satisfy the assumption of Proposition 5.8. Thus, by Proposition 5.8, the map I where for an n -ary $f \in \mathcal{M}_3$ we define $I(f)$ to be the union of all f -equivalence classes of size smaller than the constant $\ell(N^r)$ given in Lemma 5.5 is a selection

function. It remains to show that it is bounded. Since M is finitely generated and N is finite, there is a finite number λ of homomorphisms $M \rightarrow N$, and therefore for an n -ary $f \in \mathcal{M}_3$ there are only λ possibilities for f_i with $i \in [n]$. Thus, the number of distinct f -equivalence classes is λ for each $f \in \mathcal{M}_3$. Therefore, at most λ of those f -equivalence classes have size smaller than $\ell(N^r)$. It follows that $\mathcal{I}$ is a bounded selection function, which concludes the proof. $\square$

6 Tractability

Let $\mathbf{A} \rightarrow \mathbf{B}$ be two relational structures, and $i \in \mathbb{N}$. A 2-block symmetric polymorphism $f \in \text{Pol}(\mathbf{A}, \mathbf{B})$ of arity $2i + 1$ is a homomorphism $f : \mathbf{A}^{2i+1} \rightarrow \mathbf{B}$ such that $f^\sigma = f$ for all bijections $\sigma : [2i + 1] \rightarrow [2i + 1]$ satisfying $\sigma([i + 1]) = [i + 1]$ (i.e., σ preserves the sets $\{1, \dots, i + 1\}$ and $\{i + 2, \dots, 2i + 1\}$). We use the following characterization of solvability of PCSPs via the BLP+AIP algorithm defined in [10].

THEOREM 6.1 ([10]). *Let $\mathbf{A} \rightarrow \mathbf{B}$ be finite relational structures. Then $\text{PCSP}(\mathbf{A}, \mathbf{B})$ is solvable in polynomial time via BLP+AIP if, and only if, $\text{Pol}(\mathbf{A}, \mathbf{B})$ contains 2-block symmetric polymorphisms of all odd arities.*

The following shows the tractability side of our main result, Theorem 3.1.

Proposition 6.2. *Let $\mathbf{M} = (M, R(M))$, $\mathbf{N} = (N, R(N))$ be expansions of monoids such that there exists a homomorphism $\mathbf{M} \rightarrow \mathbf{N}$. Assume that M is finitely generated and N is finite. The following are equivalent.*

- (1) *There is a homomorphism $h : \mathbf{M} \rightarrow \mathbf{N}$ with regular commutative image, satisfying that $\text{Cos}(h(R(M))) \subseteq R(N)$.*
- (2) *$\text{Pol}(\mathbf{M}, \mathbf{N})$ contains 2-block symmetric polymorphisms of all odd arities.*
- (3) *There is a finite structure $\mathbf{A}$ satisfying $\mathbf{M} \rightarrow \mathbf{A} \rightarrow \mathbf{N}$ such that $\text{CSP}(\mathbf{A})$ is solvable via BLP+AIP.*
- (4) *There is a homomorphism from $(M^{\text{r.c.}}, T)$ to $\mathbf{N}$, where $T = \text{Cos}(\pi_M^{\text{r.c.}}(R(M)))$, and $M^{\text{r.c.}}, \pi_M^{\text{r.c.}}$ are given by Lemma 3.2.*

PROOF. We show that (1) is equivalent to each of the other items. The arguments are similar to those in [26], but the implication (2) $\implies$ (1) is significantly more complex.

(1) $\implies$ (2) Let $i \in \mathbb{N}$. We construct a $2i + 1$ -ary 2-block symmetric polymorphism $f \in \text{Pol}(\mathbf{A}, \mathbf{B})$. We define

$$f(a_1, \dots, a_{i+1}, b_1, \dots, b_i) = \prod_{j=1}^{i+1} h(a_j) \prod_{j=1}^i h(b_j)^{-1}$$

for each $(a_1, \dots, a_{i+1}, b_1, \dots, b_i) \in M^{2i+1}$. The fact that f is a monoid homomorphism from M^{2i+1} to N follows from $\text{Im}(h) \leq N$ being commutative and regular. We only need to show that $f(R(M)^{2i+1}) \subseteq R(N)$. This can be seen as follows:

$$\begin{aligned} f(R(M)^{2i+1}) &= h(R(M))^{\otimes i+1} \otimes (h(R(M))^{-1})^{\otimes i} = \\ &h(R(M)) \otimes (h(R(M))^{-1} \otimes h(R(M)))^{\otimes i} \subseteq \text{Cos}(h(R(M))) \subseteq R(N). \end{aligned}$$

(2) $\implies$ (1) Suppose, for the sake of a contradiction, that there is no homomorphism $h : \mathbf{M} \rightarrow \mathbf{N}$ satisfying (1). In Proposition 5.9 we show that in this case $\text{Pol}(\mathbf{M}, \mathbf{N})$ is in the scope of Theorem 5.1. This means that $\text{Pol}(\mathbf{M}, \mathbf{N})$ is the union of subsets $\mathcal{M}_1, \dots, \mathcal{M}_k$ such that there exists a bounded selection function $\mathcal{I}_i$ for $\mathcal{M}_i$ for each $i \in [k]$. Let ℓ be a bound on $|\mathcal{I}_i(f)|$ for all $i \in [k]$, $f \in \mathcal{M}_i$. Let f be a 2-block symmetric polymorphism of arity $4\ell + 1$, and let i be such that $f \in \mathcal{M}_i$. Because $|\mathcal{I}_i(f)| \leq \ell$, and each block of f has size at least 2ℓ , there is a bijection $\sigma : [4\ell + 1] \rightarrow [4\ell + 1]$ that preserves each block and satisfies $\mathcal{I}_i(f) \cap \sigma(\mathcal{I}_i(f)) = \emptyset$. However, $f = f^\sigma$, so this contradicts the

fact that I_i is a selection function on $\mathcal{M}_i$. We derived this contradiction from the assumption that (1) does not hold, so this completes the proof.

(1) $\implies$ (3) Let $\mathbf{A}$ be the expansion of the monoid $\text{Im}(h)$ by the relation $\text{Cos}(h(R))$. Observe that h is a homomorphism from $\mathbf{M}$ to $\mathbf{A}$ and the inclusion is a homomorphism from $\mathbf{A}$ to $\mathbf{N}$. We show that $\text{CSP}(\mathbf{A})$ is solvable via BLP+AIP. By Theorem 6.1, we just need to find a 2-block symmetric polymorphism $f \in \text{Pol}(\mathbf{A})$ of arity $2i + 1$ for each $i \in \mathbb{N}$. Here $\text{Pol}(\mathbf{A})$ denotes $\text{Pol}(\mathbf{A}, \mathbf{A})$. We define

$$f(a_1, \dots, a_{i+1}, b_1, \dots, b_i) = \prod_{j=1}^{i+1} a_j \prod_{j=1}^i b_j^{-1},$$

for each $a_1, \dots, a_{i+1}, b_1, \dots, b_i \in \text{Im}(h)$. The fact that f is a well-defined monoid homomorphism from $\text{Im}(h)^{2i+1}$ to $\text{Im}(h)$ follows from the fact that $\text{Im}(h)$ is commutative and regular. Let $T = \text{Cos}(h(R))$. To see that $f(T^{2i+1}) \subseteq T$, observe that $f(T) = T^{\otimes i+1} \otimes (T^{-1})^{\otimes i} = T \otimes (T^{-1} \otimes T)^{\otimes i} = T$, where the last equality uses the fact that T is a coset.

(3) $\implies$ (2) Suppose there is a finite structure $\mathbf{A}$ such that $\text{CSP}(\mathbf{A})$ is solvable via BLP+AIP, and there are homomorphisms $g_1 : \mathbf{M} \rightarrow \mathbf{A}$ and $g_2 : \mathbf{A} \rightarrow \mathbf{N}$. Let $i \in \mathbb{N}$. We show that $\text{Pol}(\mathbf{M}, \mathbf{N})$ has a 2-block symmetric polymorphism f of arity $2i + 1$. By Theorem 6.1, $\text{Pol}(\mathbf{A})$ contains a 2-block symmetric polymorphism h of arity $2i + 1$. Then we can define $f = g_2 \circ h \circ g_1$. (1) $\implies$ (4) Follows directly from Lemma 3.2.

(4) $\implies$ (1) Follows from the fact that the homomorphic image of a regular commutative monoid must be commutative and regular, and the homomorphic image of a coset is a coset. $\square$

Algorithm for Infinite Templates. Let $\mathbf{M}$ be the expansion of a finitely generated regular commutative monoid M by a coset $R(M) \subseteq M^r$. We sketch a polynomial-time algorithm solving $\text{CSP}(\mathbf{M})$. This way, using item (4) in Proposition 6.2, we obtain that all the tractable problems $\text{PCSP}(\mathbf{A}, \mathbf{B})$ within the scope of our main theorem, Theorem 3.1, can be solved by an algorithm that does not depend on the second structure $\mathbf{B}$ of the template. Our algorithm is an extension of the algorithm outlined in [24, Lemma 23], which is a polynomial-time algorithm to solve systems of equations over a finite regular commutative monoid, although we observe the algorithm can also be applied to finitely generated monoids. The full details can be found in Section 7

By Proposition 4.1, we can describe M as the disjoint union $\bigsqcup_{d \in M_I} H_d$, where $M_I \leq M$ is the semilattice formed by the idempotent elements in M and H_d is the $\mathcal{H}$ -class of $d \in M_I$, which is a subgroup of M . Using the fact that M is finitely generated, Lemma 4.2 yields that M_I is finite. We define $\pi_I : M \rightarrow M_I$ as the natural projection that sends each element $a \in M$ to the only idempotent element δ_a in its $\mathcal{H}$ -class. Finally, let $R(M_I) \subseteq M_I^r$ be the set $\{\pi_I(a) \mid a \in R(M)\}$, and $\mathbf{M}_I$ be the expansion of M_I by $R(M_I)$.

Let $\mathbf{X}$ be an instance of $\text{CSP}(\mathbf{M})$. We can see this instance as a set of variables X and constraints of the form $x_1 x_2 = x_3$, $x = e$, and $(x_1, \dots, x_r) \in R$. If there is a solution $f : \mathbf{X} \rightarrow \mathbf{M}$, then $\pi_I \circ f$ is a solution for $\mathbf{X}$ in $\text{CSP}(\mathbf{M}_I)$. It is not difficult to see that homomorphisms $h : \mathbf{X} \rightarrow \mathbf{M}_I$ are closed under point-wise products. Hence, if there exists such a homomorphism, there must be a minimal homomorphism $h_{\min}$. “Minimal” here means that $h_{\min}(x) \leq h(x)$ for any $h : \mathbf{X} \rightarrow \mathbf{M}_I$, $x \in X$. It can also be seen that if $f : \mathbf{X} \rightarrow \mathbf{M}$ and $h : \mathbf{X} \rightarrow \mathbf{M}_I$ are homomorphisms, then their point-wise product fh is still a homomorphism from $\mathbf{X}$ to $\mathbf{M}$. This way, we can conclude that if there is such a homomorphism f , then there must be one satisfying $\pi_I \circ f = h_{\min}$. Our algorithm aims to find a homomorphism f with this property.

Firstly, arc-consistency (e.g., [4]) or a slight modification of the algorithm in [24, Lemma 19] allows us to find $h = h_{\min}$ if $\mathbf{X}$ is a satisfiable instance. Using the homomorphism h we reduce the problem of finding $f : \mathbf{X} \rightarrow \mathbf{M}$ to the problem of solving a system of linear equations Σ_X over $\mathbb{Z}$, which can be solved in polynomial time using Gaussian elimination. The key insight is that

for each $d \in M_I$, the group H_d is commutative and finitely generated, so it is isomorphic to the homomorphic image of some power of $\mathbb{Z}$ (e.g., [22, Theorem 1.4, Chapter II]). In other words, H_d is isomorphic to $\mathbb{Z}^{k_d}/G_d$ for some number $k_d \in \mathbb{N}$ and some subgroup $G_d \leq \mathbb{Z}^{k_d}$. This way, we can choose a big enough number $k \in \mathbb{N}$ such that all the elements in $\mathbf{M}$ can be represented with pairs of the form (d, v) , where $d \in M_I$ is an idempotent element, and $v \in \mathbb{Z}^k$ is an integer vector that is mapped to the quotient $\mathbb{Z}^{k_d}/G_d$. We refer to Section 7 for the details. In Σ_X we represent each variable $x \in X$ by a vector of integer variables $v^x = (v_1^x, \dots, v_k^x)$. Constraints of the form $x = e$ are translated to $v^x = \mathbf{0}$. Given a constraint of the form $xy = z$ in $\mathbf{X}$, we add to Σ_X equations ensuring $v^x + v^y = v^z$ inside H_d , where $d = h(z)$. Finally, we observe that for each $d \in R(M_I)$, the intersection $R(M) \cap H_d$ is a coset⁸. A coset U in a commutative group G can be expressed (using additive notation) as $a + H$, where $a \in G$ is an element, and $H \leq G$ is a subgroup. This way, the condition $a \in R(M_I)$ can be expressed via a set of linear equations inside each group of the form H_d with $d \in M_I^r$. Using this insight, we represent each constraint $(x_1, \dots, x_r) \in R$ in $\mathbf{X}$ with a set of linear equations in Σ_X ensuring $(v^{x_1}, \dots, v^{x_r}) \in R(M_I) \cap H_d$, where $d = (h(x_1), \dots, h(x_r))$. This completes the reduction.

7 Algorithm For Infinite Templates

Let N be a semilattice with identity element $e(N)$, Q a set, and let $\lambda : N \rightarrow 2^Q$ be a map satisfying $\lambda(a) \subseteq \lambda(b)$ whenever $b \leq a$, for any $a, b \in N$. The monoid $M = N \times_\lambda \mathbb{Z}$ consists of the set

$$\bigsqcup_{a \in N} \{a\} \times \mathbb{Z}^{\lambda(a)}.$$

We implicitly identify $\mathbb{Z}^{\lambda(a)}$ with the subspace of $\mathbb{Z}^Q$ consisting of all vectors v satisfying $v_\alpha = 0$ for all $\alpha \notin \lambda(a)$. This way, given $a, b \in N$, $u \in \mathbb{Z}^{\lambda(a)}$, $v \in \mathbb{Z}^{\lambda(b)}$ it holds that $u + v \in \mathbb{Z}^{\lambda(ab)}$. Hence, we define $(a, u) \cdot (b, v) = (ab, u + v)$ for all $(a, u), (b, v)$. The identity element of the resulting monoid M is the element $(e(N), \mathbf{0})$. It is easy to see that M is commutative and regular. The idempotent elements of M are precisely $(d, \mathbf{0})$ for each $d \in N$, so M_I is isomorphic to N , and the $\mathcal{H}$ -class of $(d, \mathbf{0})$ is the subgroup $\{d\} \times \mathbb{Z}^{\lambda(d)}$.

Let $\Xi : N \rightarrow 2^{2^Q}$ be such that $\Xi(d)$ is a subgroup of $\mathbb{Z}^{\lambda(d)}$ for each $d \in N$, and $\Xi(a) \subseteq \Xi(b)$ whenever $b \leq a$. Then we define $N \times_\lambda^\Xi \mathbb{Z}$ as the quotient of $N \times_\lambda \mathbb{Z}$ by the equivalence relation $\cong$ defined by $(a, u) \cong (b, v)$, whenever $a = b$ and $u - v \in \Xi(a)$. Given $(d, v) \in N \times_\lambda \mathbb{Z}$, we write $[d, v]$ for its corresponding $\cong$ -class in $N \times_\lambda^\Xi \mathbb{Z}$. It is not difficult to see that $N \times_\lambda^\Xi \mathbb{Z}$ equipped with the product

$$[a, u][b, v] = [ab, u + v]$$

is a well-defined monoid. Moreover, $N \times_\lambda^\Xi \mathbb{Z}$ must be commutative and regular because it is a homomorphic image of $N \times_\lambda \mathbb{Z}$.

Lemma 7.1. *Let M be a regular commutative monoid generated by a finite set $Q \subseteq M$. Then there are suitable maps $\lambda : M_I \rightarrow 2^Q$ and $\Xi : M_I \rightarrow 2^{2^Q}$ such that M is isomorphic to $M_I \times_\lambda^\Xi \mathbb{Z}$.*

PROOF. Let $\lambda : M_I \rightarrow 2^Q$ be the map defined by $d \mapsto \{\alpha \in Q \mid d \leq \alpha\}$. We first prove that M is a homomorphic image of $M_I \times_\lambda \mathbb{Z}$. Consider the map $\rho : M_I \times_\lambda \mathbb{Z} \rightarrow M$ given by

$$(d, v) \mapsto d \prod_{\alpha \in \lambda(d), v_\alpha \neq 0} \alpha^{v_\alpha}. \quad (1)$$

⁸We remind the reader that the notion of coset introduced in Section 2 generalizes cosets in Abelian groups, so there should be no confusion here.

We claim that ρ is a monoid homomorphism. By construction, ρ preserves the identity element. Let us show that it also preserves products. Another way of writing (1) is

$$(d, \mathbf{v}) \mapsto d \prod_{\alpha \in \lambda(d)} \alpha^{v_\alpha},$$

where we adopt the convention that $\alpha^0 = d_\alpha$. To see that this is equivalent to (1), observe that, by definition of λ , the fact that $\alpha \in \lambda(d)$ implies that $d \leq d_\alpha$, so $dd_\alpha = d$. Let $(d_1, \mathbf{u}), (d_2, \mathbf{v}) \in M_I \times_\lambda \mathbb{Z}$. Then

$$\begin{aligned} \rho(d_1, \mathbf{u})\rho(d_2, \mathbf{v}) &= d_1 d_2 \prod_{\alpha \in \lambda(d_1) \cup \lambda(d_2)} \alpha^{u_\alpha + v_\alpha} \\ &= d_1 d_2 \prod_{\alpha \in \lambda(d_1 d_2)} \alpha^{(u+v)_\alpha} = \rho(d_1 d_2, \mathbf{u} + \mathbf{v}). \end{aligned}$$

This shows that ρ is a monoid homomorphism. Now, it is also not hard to see that ρ is surjective. Let $a \in M$. Then a can be expressed as $\prod_{\alpha \in Q_a} \alpha^{n_\alpha}$ for some subset of generators $Q_a \subseteq Q$ and some integers $n_\alpha > 0$ for each $\alpha \in Q_a$. In particular, it must be that $d_a \leq \alpha$ for each $\alpha \in Q_a$, so $Q_a \subseteq \lambda(d_a)$. Additionally, the fact that $a = d_a a$ implies $a = d_a \prod_{\alpha \in Q_a} \alpha^{n_\alpha}$. Define $\mathbf{v} \in \mathbb{Z}^{\lambda(d_a)}$ by letting $v_\alpha = n_\alpha$ if $\alpha \in Q_a$ and $v_\alpha = 0$ otherwise. Then $a = \rho(d_a, \mathbf{v})$.

The fact that $M = \text{Im}(\rho)$ implies that M is isomorphic to the quotient of $M_I \times_\lambda \mathbb{Z}$ by the congruence $\cong$ given by $(a, \mathbf{u}) \cong (b, \mathbf{v})$ whenever $\rho(a, \mathbf{u}) = \rho(b, \mathbf{v})$. Let us analyze this congruence. The first observation is that if $c = \rho(d, \mathbf{v})$, then $d_c = d$. This shows that $(a, \mathbf{u}) \cong (b, \mathbf{v})$ implies that necessarily $a = b$. For each $d \in M_I$, define $\Xi(d) \subseteq \mathbb{Z}^{\lambda(d)}$ as the submodule consisting of the vectors $\mathbf{v}$ satisfying $\rho(d, \mathbf{v}) = \rho(d, \mathbf{0})$. By construction $(a, \mathbf{u}) \cong (b, \mathbf{v})$ if and only if $a = b$, and $\mathbf{u} - \mathbf{v} \in \Xi(a)$. This shows that M is isomorphic to $M_I \times_\lambda^\Xi \mathbb{Z}$. $\square$

In the following result we will consider the expansion of a monoid M of the form $N \times_\lambda^\Xi \mathbb{Z}$ by a relation $R \subseteq M^r$. It will be convenient to represent tuples $([d_1, \mathbf{v}_r], \dots, [d_r, \mathbf{v}_r]) \in M^r$ as pairs $[\mathbf{d}, \mathbf{u}]$ where $\mathbf{d} = (d_1, \dots, d_r)$ and $\mathbf{u} \in \prod_{i \in [r]} \mathbb{Z}^{\lambda(d_i)}$ is a vector whose projection to $\mathbb{Z}^{\lambda(d_i)}$ is $\mathbf{v}_i$ for each $i \in [r]$. This way, we identify M^r with the set

$$\bigsqcup_{\mathbf{d} \in N^r} \{\mathbf{d}\} \times \prod_{i \in [r]} \mathbb{Z}^{\lambda(d_i)} / \Xi(d_i).$$

THEOREM 7.2. *If $\mathbf{M}$ is an expansion of a finitely generated regular commutative monoid M by a coset $R \subseteq M^r$, then $\text{CSP}(\mathbf{M})$ is solvable in polynomial time.*

PROOF. Let $\mathbf{X}$ be an instance of $\text{CSP}(\mathbf{M})$. We can see $\mathbf{X}$ as a set of variables X together with expressions of the form (1) $x_1 x_2 = x_3$ for $x_1, x_2, x_3 \in X$, (2) $x = e$ for $x \in X$, and (2) $(x_1, \dots, x_r) \in R$ for $x_1, \dots, x_r \in X$.

By Lemma 4.2, the fact that M is finitely generated implies that the submonoid $M_I \leq M$ of idempotent elements is finite. Hence, by Lemma 7.1, we can assume that M is of the form $N \times_\lambda^\Xi \mathbb{Z}$, where N is a finite semilattice, $\lambda : N \rightarrow 2^Q$ for some finite set Q , and $\Xi : N \rightarrow 2^{\mathbb{Z}^Q}$. Recall that $M = \bigsqcup_{\delta \in M_I} H_\delta$, where H_δ denotes the $\mathcal{H}$ -class of the idempotent M_I , which is a subgroup of M . As $M = N \times_\lambda^\Xi \mathbb{Z}$, we have that $M_I = \{[\delta, \mathbf{0}] \mid \delta \in N\}$. We abuse the notation and identify N with M_I via the isomorphism $\delta \mapsto [\delta, \mathbf{0}]$. This way, given $\delta \in M_I$, the subgroup H_δ is defined as $\{[\delta, \mathbf{v}] \mid \mathbf{v} \in \mathbb{Z}^{\lambda(\delta)}\}$.

Given $a \in M$, we write δ_a for the unique idempotent element satisfying $\delta_a \mathcal{H} a$. Similarly, if $\mathbf{a} \in M^r$, we write $\delta_{\mathbf{a}}$ for the unique idempotent element in the $\mathcal{H}$ -class of $\mathbf{a}$ (such element exists because M^r is a regular commutative monoid). Let $\mathbf{M}_I$ be the expansion of M_I by the relation $R(\mathbf{M}_I) = \{\delta_{\mathbf{a}} \mid \mathbf{a} \in R\}$.

Then the projection $\pi_I : \mathbf{M} \rightarrow \mathbf{M}_I$ defined by $a \mapsto \delta_a$ is a homomorphism. Suppose that $f : \mathbf{X} \rightarrow \mathbf{M}$ is a homomorphism. Then $\pi_I \circ f$ is a homomorphism from $\mathbf{X}$ to $\mathbf{M}_I$. Additionally, if $h : \mathbf{X} \rightarrow \mathbf{M}_I$ is a homomorphism, then the product fh is also a homomorphism from $\mathbf{X}$ to $\mathbf{M}$. Indeed, fh clearly preserves products and sends $e(X)$ to $e(M)$. To see that fh preserves the relation R , observe that if $R(X)$ contains a tuple $\mathbf{x}$, then $f(\mathbf{x}) \in R(M)$ and $h(\mathbf{x}) \in R(M_I) \subseteq (R(M))^{-1} \otimes R(M)$. This way, using the fact that $R(M)$ is a coset, we obtain

$$fh(\mathbf{x}) = f(\mathbf{x})h(\mathbf{x}) \in R(M) \otimes ((R(M))^{-1} \otimes R(M)) \subseteq R(M).$$

Additionally, this new solution fh satisfies $\pi_I \circ (fh)(\mathbf{x}) \leq h(\mathbf{x})$ for all $\mathbf{x} \in X$. Now, observe that if $\mathbf{X} \rightarrow \mathbf{M}_I$ then there must be a *minimal* homomorphism $h : \mathbf{X} \rightarrow \mathbf{M}_I$, meaning that for any $h' : \mathbf{X} \rightarrow \mathbf{M}_I$ it holds that $h(\mathbf{x}) \leq h'(\mathbf{x})$ for all $\mathbf{x} \in X$. To see this define $h(\mathbf{x}) = \prod_{h': \mathbf{X} \rightarrow \mathbf{M}_I} h'(\mathbf{x})$. This way, if there is a solution $f' : \mathbf{X} \rightarrow \mathbf{M}$ then, by defining $f = f'h$, we can obtain a solution $f : \mathbf{X} \rightarrow \mathbf{M}$ such that $\pi_I \circ f$ is the minimal homomorphism $h : \mathbf{X} \rightarrow \text{CSP}(\mathbf{M}_I)$. Our algorithm aims to find such a solution f .

By [24, Lemma 19], we can find a minimal solution $h : \mathbf{X} \rightarrow \mathbf{M}_I$ in polynomial time if one exists. Otherwise we reject the instance $\mathbf{X}$. Now we look for a map $f : \mathbf{X} \rightarrow \mathbf{M}$ such that $f(\mathbf{x})$ is of the form $[\delta, \mathbf{v}] \in H_\delta$ for some, whenever $h(\mathbf{x}) = \delta$. Next we construct a system Σ_X of linear equations over $\mathbb{Z}$. For each variable $x \in X$, we include a vector $\mathbf{v}^x$ of integer variables $(v_\alpha^x)_{\alpha \in \lambda(h(\mathbf{x}))}$. The intuition is that our intended solution f should be defined as $f(\mathbf{x}) = [h(\mathbf{x}), \mathbf{v}^x]$. For each $x \in X$ and each $\alpha \notin \lambda(h(\mathbf{x}))$ we add to Σ_X an equation $v_\alpha^x = 0$. For each constraint in $\mathbf{X}$ of the form $\mathbf{x} = e$, we add to Σ_X equations ensuring $\mathbf{v}^x = \mathbf{0}$.

We make use of the well-known result that any subgroup of a finitely generated commutative group is itself finitely generated [22, Theorem 1.6, Chapter II]. For each $\delta \in M_I$, let W_δ be a finite set generating the subgroup $\Xi(\delta) \subseteq \mathbb{Z}^{\lambda(\delta)}$. For each constraint in $\mathbf{X}$ of the form $xy = z$, we add to Σ_X equations ensuring

$$\mathbf{v}^x + \mathbf{v}^y = \mathbf{v}^z + \sum_{\mathbf{u} \in W_{h(z)}} n_{\mathbf{u}} \mathbf{u}$$

where $n_{\mathbf{u}}$ is a fresh variable for each $\mathbf{u} \in W_{h(z)}$. This simply ensures that

$$[h(\mathbf{x}), \mathbf{v}^x] + [h(\mathbf{y}), \mathbf{v}^y] = [h(\mathbf{z}), \mathbf{v}^z].$$

For each $\delta \in \pi_I(R)$, let R_δ be $R \cap H_\delta$. Using the fact that R is a coset of M^r , we obtain that R_δ must be a coset of H_δ . This way, there must be a coset $U_\delta \subseteq \prod_{i \in [r]} \mathbb{Z}^{\lambda(\delta_i)}$ such that $R_\delta = \{\delta\} \times U_\delta$. It is not hard to show that a coset U in a commutative group G must be of the form $a + H$ (using additive notation), where $a \in G$ and $H \leq G$ is a subgroup. Using again the fact that subgroups of finitely generated commutative groups must be finitely generated, we conclude there is a vector $\mathbf{o}_\delta \in \prod_{i=1}^r \mathbb{Z}^{\lambda(\delta_i)}$ and $V_\delta \subseteq \prod_{i=1}^r \mathbb{Z}^{\lambda(\delta_i)}$ a finite subset such that

$$R_\delta = \{[\delta, \mathbf{u}] \mid \mathbf{u} \in \mathbf{o}_\delta + \langle V_\delta \rangle\}.$$

Then, for each constraint in $\mathbf{X}$ of the form $\mathbf{x} = (x_1, \dots, x_r) \in R$, we add to Σ_X equations ensuring

$$\mathbf{v}^x = \mathbf{o}_{h(\mathbf{x})} + \sum_{\mathbf{u} \in V_{h(\mathbf{x})}} n_{\mathbf{u}} \mathbf{u},$$

where $\mathbf{v}^x$ denotes the vector $(v^{x_1}, \dots, v^{x_r})$, and $n_{\mathbf{u}}$ is a fresh variable for each $\mathbf{u} \in V_{h(\mathbf{x})}$. This ensures that $[h(\mathbf{x}), \mathbf{v}^x] \in R(M)$. Now, by construction, Σ_X has a solution over the integers if and only if $\mathbf{X}$ has a solution in $\text{CSP}(\mathbf{M})$. Systems of linear equations over the integers can be solved in polynomial time via Gaussian elimination, so this completes the proof. $\square$

A Auxiliary Results About Monoids

Lemma 4.2. *Let M be a commutative monoid and $M_I \subseteq M$ be the subset of its idempotent elements. Then, (1) M_I is a semilattice, and (2) if M is regular and finitely generated, then M_I is finite.*

PROOF. Item (1) is straightforward. Let us prove (2). If M is regular, then the quotient $M/\mathcal{H}$ is isomorphic to M_I through the map that sends each $\mathcal{H}$ -class to the single idempotent element belonging to it. Hence if M is finitely generated, so is M_I . Since M_I is a finitely generated semilattice, it is finite. $\square$

The following is a characterization of group elements in finite regular commutative monoids.

Lemma A.1 ([26, Lemma 1]). *Let M be a finite commutative monoid, and $a \in M$ an element. Then the following are equivalent.*

- (1) $a^2b = a$ for some $b \in M$,
- (2) $a^k = a$ for some $k > 1$,
- (3) a is a group element,
- (4) $a \leq a^2$.

Lemma 4.3. *Let M be a finite monoid. There is an integer $C > 1$ such that a^C is idempotent for all $a \in M$. If $a \in M$ is a group element, then it also holds that $a^{C-1} = a^{-1}$. Additionally, given $a \in M$, there is a unique idempotent element d_a that satisfies $d_a = a^n$ for some $n \in \mathbb{N}$.*

PROOF. Let $a \in M$. As M is finite, there must be some $\ell_a, k_a \in \mathbb{N}$ satisfying that

$$a^{\ell_a+n} = a^{\ell_a+m}, \quad (2)$$

for any integers $n, m \geq 0$ satisfying $m = n \pmod{k_a}$ (see e.g., [12][Section 1.6]). This way, the element $a^{k_a\ell_a}$ is idempotent. Indeed,

$$a^{2k_a\ell_a} = a^{\ell_a+(2k_a-1)\ell_a} = a^{\ell_a+(k_a-1)\ell_a} = a^{k_a\ell_a}.$$

Here, the second equality uses (2). The constant C that witnesses the result can be defined as $C = \prod_{a \in M} k_a\ell_a$. If $a \in M$ is a group element, then it belongs to a subgroup whose identity must be a^C , so $a^{-1} = a^{C-1}$.

Fix $a \in M$. We show that there is a unique idempotent element d_a of the form a^n for some $n \in \mathbb{N}$. We have shown that there is at least one idempotent a^n of this form. Suppose that there is another positive integer $m \neq n$ such that a^m is idempotent as well. It holds that $a^m\mathcal{H}a^n$. Indeed, there are numbers $k, \ell \in \mathbb{N}$ satisfying $km > n$ and $\ell n > m$, so

$$a^n = a^{\ell n} = a^m a^{\ell n - m} = a^{\ell n - m} a^m,$$

and

$$a^m = a^{km} = a^n a^{km-n} = a^{km-n} a^n.$$

By Theorem 2.1, there is at most a single idempotent element in each $\mathcal{H}$ -class of M , so $a^m = a^n$. This completes the proof. $\square$

Lemma 4.4. *Let M be a finite commutative monoid, $M_I \subseteq M$ its subset of idempotent elements, and $M_{\dagger} \subseteq M$ its subset of group elements. The following hold: (1) The map $\pi_I : M \rightarrow M_I$ given by $a \mapsto d_a$ is a surjective monoid homomorphism satisfying $\pi_I \circ \pi_I = \pi_I$. (2) $M_{\dagger}$ is a regular submonoid of M . (3) The map $\pi_{\dagger} : M \rightarrow M_{\dagger}$ given by $a \mapsto d_a a$ is a surjective monoid homomorphism satisfying $\pi_{\dagger} \circ \pi_{\dagger} = \pi_{\dagger}$.*

PROOF. (1) The map π_I preserves the identity element, because $d_{e_M} = e_M$. To see that π_I preserves the product, let $k \in \mathbb{N}$ be such that $a^k = d_a$ for each $a \in M$, as given by Lemma 4.3. Then $d_a d_b = a^k b^k = (ab)^k = d_{ab}$ for each $a, b \in M$. (2) Let $k \in \mathbb{N}$ be such that a^k is idempotent for all $a \in M$, as given by Lemma 4.3. We have that $a \in M$ is a group element if and only if $a^{k+1} = a$. Indeed, if $a^{k+1} = a$ then a is a group element by Lemma A.1. Conversely, if a is a group element, by Theorem 2.1 its $\mathcal{H}$ -class is a subgroup of M . The identity of this subgroup must be a^k , so $a^{k+1} = a$. (3) The map $\pi_{\dagger}$ clearly preserves the identity element. To see that it also preserves products, observe that by (1), $d_a d_b = d_{ab}$ for every $a, b \in M$, so $ad_a b d_b = ab d_{ab}$. Finally, by (1), it holds that $d_{ad_a} = d_a d_{d_a} = d_a^2 = d_a$ for all $a \in M$. This way, $\pi_{\dagger} \circ \pi_{\dagger}(a) = (ad_a) d_{ad_a} = ad_a d_a = ad_a$ for all $a \in M$. $\square$

Lemma 4.5. *Let M be a monoid, and N be a commutative monoid. Let F be a finite family of monoid homomorphisms $f : M \rightarrow N$, and $g : M \rightarrow N$ be the homomorphism given by $g(a) = \prod_{f \in F} f(a)$.⁹ Then for each $S \subseteq M$, it holds that $\text{Cos}(g(S)_{\dagger}) \subseteq \bigotimes_{f \in F} \text{Cos}(f(S)_{\dagger})$.*

PROOF. Let $a \in \text{Cos}(g(S)_{\dagger})$. Then we can express a as a product of the form

$$g(b)_{\dagger} \prod_{i \in [k]} g(s_i)_{\dagger} g(t_i)_{\dagger}^{-1},$$

for some elements $b, s_i, t_i \in S$. By Lemma 4.4, for each $s \in M$ it holds that $g(s)_{\dagger} = \prod_{f \in F} f(s)_{\dagger}$. Hence, we can write

$$a = \prod_{f \in F} \left(f(b)_{\dagger} \prod_{i \in [k]} f(s_i)_{\dagger} f(t_i)_{\dagger}^{-1} \right).$$

This shows that $a \in \bigotimes_{f \in F} \text{Cos}(f(S)_{\dagger})$, as we wanted to prove. $\square$

B Commutative Regularization

The aim of this section is to prove Lemma B.1. More explicitly, given a monoid M , we construct its *commutative regularization*. This is a second monoid $M^{\text{r.c.}}$ (r.c. standing for *regular commutative*) together with a homomorphism $f : M \rightarrow M^{\text{r.c.}}$ satisfying the following universal property: for every regular commutative monoid N and every homomorphism $g : M \rightarrow N$, there exists a unique homomorphism $h : M^{\text{r.c.}} \rightarrow N$ satisfying $g = f \circ h$.

We use the notion of a monoid presented by a set of generators and relations [21]. This way, in a monoid M presented by a set of generators S and some relations, the elements of M are equivalence classes of S^* , where S^* denotes the set of non-empty words over S . Given a word $\alpha \in S^*$, we write $[\alpha]$ to denote its equivalence class in M .

Let M be a monoid. Informally, to define $M^{\text{r.c.}}$ we add inverses to each element $a \in M$ and we impose relations that make the resulting monoid commutative. We define $M^{\text{r.c.}}$ as the monoid presented by the set of generators Ω that contains the symbols $\hat{a}, \check{a}, 1_a$ for each element $a \in M$, and the set of relations containing all identities of the form (1) $\prod_{i \in [k]} \hat{a}_i = \prod_{j \in [\ell]} \hat{b}_j$ for each equality $\prod_{i \in [k]} a_i = \prod_{j \in [\ell]} b_j$ that holds in M , (2) $\hat{a} 1_a = \hat{a}$ for each $a \in M$, (3) $\hat{a} \check{a} = 1_a$ for each $a \in M$, (4) $\widehat{e_M} = e$, and (5) $xy = yx$ for each $x, y \in \Omega$.

Lemma B.1. *The monoid $M^{\text{r.c.}}$ is commutative and regular. Moreover, for each $a \in M$, the idempotent element in the $\mathcal{H}$ -class of $[\hat{a}]$ is $[1_a]$, and $[\hat{a}]^{-1} = [\check{a}]$.*

⁹Observe that this notation is well-defined because products commute in a commutative monoid.

PROOF. Commutativity follows from (5). In order to see that $M^{\text{r.c.}}$ is regular, first observe that the elements $[1_a]$ are idempotent for all $a \in M$:

$$[1_a^2] = [1_a \hat{a} \tilde{a}] = [\hat{a} \tilde{a}] = [1_a].$$

Now let us see that $M^{\text{r.c.}}$ is regular. By Proposition 4.1 we just need to show that there is an idempotent element in each $\mathcal{H}$ -class of $M^{\text{r.c.}}$. As $M^{\text{r.c.}}$ is a commutative monoid and $\{[a] \mid a \in \Omega\}$ is a set of generators, any non-identity element c can be expressed as $\prod_{i \in [k]} [b_i]^{n_i}$ for some elements $b_1, \dots, b_k \in \Omega$ and some integers $n_1, \dots, n_k \geq 1$. For each $i \in [k]$ define $b'_i = \tilde{a}$ if $b_i = \hat{a}$, $b'_i = \hat{a}$ if $b_i = \tilde{a}$, or $b'_i = b_i$ if $b_i = 1_a$. This way $[b_i][b'_i]$ is an idempotent element, and $[b_i]^2[b'_i] = [b_i]$. We define $c' = \prod_{i \in [k]} [b'_i]^{n_i}$. It holds that cc' is idempotent, and $c^2c' = c$, proving that $M^{\text{r.c.}}$ is regular.

To prove the last part of the statement, observe that both $[\hat{a}][\tilde{a}] = [1_a]$ and $[\hat{a}]^2[\tilde{a}] = [\hat{a}]$ hold for all $a \in M$, and $[1_a]$ is an idempotent element. $\square$

Lemma B.2. *Let $S \subseteq M$ be a set of generators of a monoid M . Then the set*

$$\tilde{S} = \{[\hat{a}] \mid a \in S\} \cup \{[\tilde{a}] \mid a \in S\}$$

is a set of generators of $M^{\text{r.c.}}$.

PROOF. We show that every element of the form $[b]$ for $b \in \Omega$ can be generated by these elements. Let $c \in M$ be an arbitrary element. By assumption $c = t^M(a_1, \dots, a_k)$ for some term t and some elements $a_1, \dots, a_k \in S$. Suppose that $b = \hat{c}$. Then, because $M^{\text{r.c.}}$ satisfies all identities in (1), it must hold that $[b] = t^{M^{\text{r.c.}}}([\hat{a}_1], \dots, [\hat{a}_k])$. Suppose that $b = \tilde{c}$. We know that $[\tilde{a}] = [\hat{a}]^{-1}$ for all $a \in M$, and that the map that sends each element to its inverse is an endomorphism of $M^{\text{r.c.}}$. Hence it must be that $[b] = t^{M^{\text{r.c.}}}([\tilde{a}_1], \dots, [\tilde{a}_k])$. Finally, suppose that $b = 1_c$. In other words, $[b]$ is the idempotent element in $H_{[\hat{c}]}$. Recall that the map sending each element $a \in M^{\text{r.c.}}$ to the idempotent element in H_a is a homomorphism. Hence, $t^{M^{\text{r.c.}}}([1_{a_1}], \dots, [1_{a_k}])$ must be the idempotent element in $H_{[\hat{c}]}$, which equals $[b]$. Observe that $[1_a] = [\hat{a}][\tilde{a}]$ for each $a \in M$, so each $[1_{a_i}]$ is generated by the elements in $\tilde{S}$. This completes the proof. $\square$

Consider the map $\pi_M^{\text{r.c.}} : M \rightarrow M^{\text{r.c.}}$ defined by $a \mapsto [\hat{a}]$. By definition $M^{\text{r.c.}}$ satisfies all identities in (1), so $\pi_M^{\text{r.c.}}$ preserves products. By the identity (4), $[\hat{e}_M]$ is the identity element in $M^{\text{r.c.}}$, so $\pi_M^{\text{r.c.}}$ is a monoid homomorphism.

Lemma B.3. *Let $f : M \rightarrow N$ be a monoid homomorphism, where N is a regular commutative monoid. Then there is another monoid homomorphism $g : M^{\text{r.c.}} \rightarrow N$ such that $g \circ \pi_M^{\text{r.c.}} = f$.*

PROOF. Let Ω be the set of generators in the definition of $M^{\text{r.c.}}$. We define a map $h : \Omega \rightarrow N$ as follows. For each $a \in M$ let $h(\hat{a}) = f(a)$, $h(\tilde{a}) = f(a)^{-1}$, and $h(1_a) = d_{f(a)}$. It is easy to see that for each identity $t(x_1, \dots, x_k) = s(y_1, \dots, y_\ell)$ in the presentation of $M^{\text{r.c.}}$, the identity $t^N(h(x_1), \dots, h(x_k)) = s^N(h(y_1), \dots, h(y_\ell))$ holds. This way, by the universal property of a semigroup given by a presentation, the map $g : M^{\text{r.c.}} \rightarrow N$ given by $[x] \mapsto h(x)$ is a well-defined semigroup homomorphism from $M^{\text{r.c.}}$ to N . To see that it is also a monoid homomorphism, observe that $h(\hat{e}) = f(e)$, and $f(e)$ must be the identity element of N .

Finally, the fact that $g \circ \pi_M^{\text{r.c.}} = f$ is straightforward. Given $s \in M$, we have $\pi_M^{\text{r.c.}}(s) = [\hat{s}]$, and $g([\hat{s}]) = h(\hat{s}) = f(s)$, as we wanted to show. This completes the proof. $\square$

Acknowledgments

We thank the anonymous reviewers. This work was supported by UKRI EP/X024431/1. Work done while Alberto Larrauri was at the University of Oxford. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. All data is provided in full in the results section of this paper.

References

- [1] Kristina Asimi and Libor Barto. 2021. Finitely Tractable Promise Constraint Satisfaction Problems. In *Proc. 46th International Symposium on Mathematical Foundations of Computer Science (MFCS'21) (LIPIcs, Vol. 202)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 11:1–11:16. arXiv:2010.04618 doi:10.4230/LIPIcs.MFCS.2021.11
- [2] Franz Baader, Pavlos Marantidis, Antoine Mottet, and Alexander Okhotin. 2020. Extensions of unification modulo ACUI. *Math. Struct. Comput. Sci.* 30, 6 (2020), 597–626. doi:10.1017/S0960129519000185
- [3] Franz Baader and Wayne Snyder. 2001. Unification Theory. In *Handbook of Automated Reasoning*. Elsevier, 447–533.
- [4] Libor Barto, Jakub Bulín, Andrei A. Krokhin, and Jakub Opršal. 2021. Algebraic approach to promise constraint satisfaction. *J. ACM* 68, 4 (2021), 28:1–28:66. arXiv:1811.00970 doi:10.1145/3457606
- [5] Libor Barto, William J. DeMeo, and Antoine Mottet. 2021. Constraint Satisfaction Problems over Finite Structures. In *Proc. 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'21)*. IEEE, 1–13. doi:10.1109/LICS52264.2021.9470670
- [6] Libor Barto and Marcin Kozik. 2022. Combinatorial Gap Theorem and Reductions between Promise CSPs. In *Proc. 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA'22)*. 1204–1220. arXiv:2107.09423 doi:10.1137/1.9781611977073.50
- [7] Libor Barto and Antoine Mottet. 2025. Finite Algebras with Hom-Sets of Polynomial Size. *Trans. Amer. Math. Soc.* 378 (2025), 569–596. doi:10.1090/tran/9262
- [8] Manuel Bodirsky, Barnaby Martin, Marcello Mamino, and Antoine Mottet. 2018. The Complexity of Disjunctive Linear Diophantine Constraints. In *Proc. 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS'18) (LIPIcs, Vol. 117)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 33:1–33:16. doi:10.4230/LIPIcs.MFCS.2018.33
- [9] Joshua Brakensiek and Venkatesan Guruswami. 2021. Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy. *SIAM J. Comput.* 50, 6 (2021), 1663–1700. arXiv:1704.01937 doi:10.1137/19M128212X
- [10] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. 2020. The power of the combined basic LP and affine relaxation for promise CSPs. *SIAM J. Comput.* 49 (2020), 1232–1248. Issue 6. arXiv:1907.04383 doi:10.1137/20M1312745
- [11] Silvia Butti, Alberto Larrauri, and Stanislav Živný. 2025. Optimal Inapproximability of Promise Equations over Finite Groups. In *Proc. 52nd International Colloquium on Automata, Languages, and Programming (ICALP'25)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. arXiv:2411.01630
- [12] Alfred H Clifford and Gordon B Preston. 1961. The algebraic theory of semigroups, vol. 1. *Mathematical surveys* 7 (1961).
- [13] Volker Diekert, Claudio Gutierrez, and Christian Hagenah. 2005. The existential theory of equations with rational constraints in free groups is PSPACE-complete. *Inf. Comput.* 202, 2 (2005), 105–140. doi:10.1016/J.IC.2005.04.002
- [14] Lars Engebretsen, Jonas Holmerin, and Alexander Russell. 2004. Inapproximability results for equations over finite groups. *Theor. Comput. Sci.* 312, 1 (2004), 17–45. doi:10.1016/S0304-3975(03)00401-8
- [15] Tomás Feder and Moshe Y. Vardi. 1998. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. *SIAM J. Comput.* 28, 1 (1998), 57–104. doi:10.1137/S0097539794266766
- [16] Pierre Gillibert, Julius Jonas, Michael Kompatscher, Antoine Mottet, and Michael Pinsker. 2022. When Symmetries Are Not Enough: A Hierarchy of Hard Constraint Satisfaction Problems. *SIAM J. Comput.* 51, 2 (2022), 175–213. arXiv:2002.07054 doi:10.1137/20M1383471
- [17] Christian Glaßer, Peter Jonsson, and Barnaby Martin. 2017. Circuit satisfiability and constraint satisfaction around Skolem Arithmetic. *Theor. Comput. Sci.* 703 (2017), 18–36. doi:10.1016/j.tcs.2017.08.025
- [18] Mikael Goldmann and Alexander Russell. 2002. The Complexity of Solving Equations over Finite Groups. *Inf. Comput.* 178, 1 (2002), 253–262. doi:10.1006/INCO.2002.3173
- [19] Pierre A Grillet. 2017. *Semigroups: an introduction to the structure theory*. Routledge.
- [20] Johan Håstad. 2001. Some optimal inapproximability results. *J. ACM* 48, 4 (2001), 798–859. doi:10.1145/502090.502098
- [21] John M Howie. 1995. *Fundamentals of semigroup theory*. Oxford University Press.
- [22] Thomas W Hungerford. 2012. *Algebra*. Vol. 73. Springer Science & Business Media.
- [23] Ravindran Kannan and Achim Bachem. 1979. Polynomial Algorithms for Computing the Smith and Hermite Normal Forms of an Integer Matrix. *SIAM J. Comput.* 8, 4 (1979), 499–507. doi:10.1137/0208040
- [24] Ondřej Klima, Pascal Tesson, and Denis Thérien. 2007. Dichotomies in the Complexity of Solving Systems of Equations over Finite Semigroups. *Theory Comput. Syst.* 40, 3 (2007), 263–297. doi:10.1007/S00224-005-1279-2
- [25] Andrei Krokhin and Jakub Opršal. 2022. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News* 9, 3 (2022), 30–59. arXiv:2208.13538
- [26] Alberto Larrauri and Stanislav Živný. 2024. Solving promise equations over monoids and groups. *ACM Trans. Comput. Log.* 26, 1 (2024), 1–24. arXiv:2402.08434 doi:10.1145/3698106

- [27] Antoine Mottet. 2025. Algebraic and algorithmic synergies between promise and infinite-domain CSPs. In *Proc. 40th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'25)*. IEEE. arXiv:2501.13740
- [28] Wojciech Plandowski. 2004. Satisfiability of word equations with constants is in PSPACE. *J. ACM* 51, 3 (2004), 483–496. doi:10.1145/990308.990312