

Symmetric Operations on Domains of Size at Most 4

Zarathustra Brady and Holden Mui

January 15, 2021

Abstract

To convert a fractional solution to an instance of a constraint satisfaction problem into a solution, a rounding scheme is needed, which can be described by a collection of symmetric operations with one of each arity. An intriguing possibility, raised in a recent paper by Carvalho and Krokhin, would imply that any clone of operations on a set D which contains symmetric operations of arities $1, 2, \dots, |D|$ contains symmetric operations of all arities in the clone. If true, then it is possible to check whether any given family of constraint satisfaction problems is solved by its linear programming relaxation.

We characterize all idempotent clones containing symmetric operations of arities $1, 2, \dots, |D|$ for all sets D with size at most four and prove that each one contains symmetric operations of every arity, proving the conjecture above for $|D| \leq 4$.

1 Introduction

The Constraint Satisfaction Problem, commonly abbreviated as the CSP, is the decision problem where we are given a list of variables and a list of constraints on the variables, and we must determine whether or not there exists an assignment of the variables which satisfies every constraint. While this decision problem is NP-hard in general, certain classes of CSPs can be solved in polynomial time. Among those are the CSPs that are solved by their linear programming relaxation; that is, a “fractional solution” to an instance of such a CSP can be rounded to a solution. Such CSPs have been characterized as the CSPs for which the clone of operations preserving its relations contains symmetric operations of every arity. In this paper, we investigate a conjecture that, if true, gives a sufficient condition for a clone to contain symmetric operations of every arity.

Conjecture 1. *Suppose a clone over a domain D contains symmetric operations of arities $1, 2, \dots, |D|$. Then it contains symmetric operations of every arity.*

This conjecture is a weak form of an open problem mentioned in section 6 of [7]: its authors speculate that a finite algebraic structure \mathbb{A} has symmetric terms of every arity if and only if \mathbb{A} has no subquotient \mathbb{B} such that the automorphism group of \mathbb{B} contains a pair of automorphisms with no common fixed point.

In this paper, we prove Conjecture 1 for $|D| \leq 4$.

Theorem 2. *Suppose a clone over a domain D with $|D| \leq 4$ contains symmetric operations of arities $1, 2, \dots, |D|$. Then it contains symmetric operations of every arity.*

In the appendix, we sketch a more ambitious conjecture about the solvability of certain types of weakly consistent constraint satisfaction problems attached to an algebraic structure satisfying the assumption of Conjecture 1.

1.1 Motivation

If Conjecture 1 is true, then it gives an efficient way of determining whether or not combinatorial puzzles can be solved using systems of linear inequalities.

To understand Conjecture 1’s implications, it is helpful to characterize the types of combinatorial puzzles we want to look at; a good example of such a puzzle is Sudoku. Sudoku is a constraint satisfaction problem over the domain $D = \{1, 2, 3, 4, 5, 6, 7, 8, 9\}$ with ten types of constraints. There is the 9-ary constraint which asserts that its inputs are pairwise distinct, and there are nine unary constraints, each of the form “this variable is equal to d ” for some $d \in D$. Instances of this constraint satisfaction problem have 81 variables, and the variables that the constraints apply to depend on the specific instance. Other examples of CSPs include HornSAT and 3Coloring: HornSAT is the constraint satisfaction problem where the task is to determine whether or not a set of Horn clauses (implications) admits an assignment of the variables satisfying each Horn clause, and 3Coloring is the constraint satisfaction problem whose task is to determine whether or not a given graph admits a 3-coloring.

The linear programming relaxation of an instance of a CSP is, informally, the set of all locally consistent probability distributions over its variables and constraint relations, and is defined by a collection of linear inequalities (a rigorous but terse definition is given in section 2). A CSP is solved by its linear programming relaxation iff there is a way to turn points in the linear programming relaxation of every instance of the CSP to a solution of that CSP, known as a rounding scheme.

While finding solutions to instances of a general CSP is NP-hard, solving the linear programming relaxation of any CSP only takes polynomial time, so there is an efficient way to find a solution to an instance of a CSP if it is solved by its linear programming relaxation. Usually this is not the case, but for these special CSPs, local probability distributions of solutions, called fractional solutions, can be converted to true solutions using the rounding scheme.

The CSPs solved by their linear programming relaxation have been characterized in [9]. They are precisely the ones for which the clone of operations preserving each relation defining the CSP contains symmetric operations of every arity.

Theorem 3 (Theorem 2 of [9]). *The CSP defined by a collection of relations Γ is solved by the linear programming relaxation if and only if the clone of operations that preserves each relation in Γ contains symmetric operations of every arity.*

An equivalent characterization of clones with symmetric operations of every arity appears in a recent article by Butti and Dalmau about solving CSPs with *distributed* algorithms [6]. In their setup, each agent has access to a single variable or to a single constraint, agents can communicate only when one owns a constraint involving the variable owned by the other, and the agents are *anonymous*, so there is no obvious way to elect a leader.

Theorem 4 (Theorem 6 of [6]). *The CSP defined by a collection of relations Γ can be solved in the distributed setting described above if and only if the clone of operations that preserves each relation in Γ contains symmetric operations of every arity. If it can be solved in the distributed setting at all, then it can be solved in this setting in polynomial time.*

For brevity, we call such clones *round*, as they can be used to construct a rounding scheme that turns fractional solutions into solutions of instances of CSPs. However, determining whether or not a given clone is round is difficult (possibly undecidable); that is, unless Conjecture 1 is true.

Recall that Conjecture 1 states that the existence of symmetric operations of arities $1, 2, \dots, |D|$ is a sufficient (and necessary, but this direction is obvious) condition for a clone to contain symmetric operations of every arity. Equivalently, Conjecture 1 asserts that given $|D|$ symmetric operations of arities $1, 2, \dots, |D|$, one can create a symmetric operation of any desired arity by composing the $|D|$ “base” operations in some way. If Conjecture 1 is true, then determining whether or not a clone is round becomes a finite case check!

Although Conjecture 1 remains an open problem in full generality, we prove Conjecture 1 for all clones over a domain of size at most 4. We achieve this by classifying all minimal idempotent clones over a domain of size at most 4 satisfying Conjecture 1’s hypothesis and proving that each one is round. It is our hope that this classification will help future researchers form stronger hypotheses, verify their truth for a large number of examples, and ultimately take steps closer towards a proof (or a disproof) of Conjecture 1.

1.2 Road Map

The remainder of our paper is organized as follows. In section 2, we go over definitions. In section 3, we summarize related results. In section 4, we prove results that simplify the enumeration of all minimal idempotent round clones over a domain of size at most four. In section 5, we enumerate all such clones. In section 6, we sketch a plausible line of attack on the general case of Conjecture 1. In the appendix, we describe connections between the linear programming relaxation of a CSP and certain weak consistency conditions, and we conjecture a precise connection between them.

2 Definitions

2.1 Constraint Satisfaction Problems

A *domain*, denoted by the capital letter D , is a set of values a variable can be assigned to. A k -ary relation R over a domain D is a subset of the k -fold Cartesian product $D^k := D \times \dots \times D$; k is known as the constraint's *arity*, denoted $\text{ar}(R)$. A tuple (a_1, \dots, a_k) *satisfies* a relation R if $(a_1, \dots, a_k) \in R$. A *constraint* is a pair consisting of a relation R and an $\text{ar}(R)$ -tuple of variables. A *constraint satisfaction problem* is a pair $P = (D, \Gamma)$ where D is its domain and Γ is a set of relations over D . An *instance* of a constraint satisfaction problem (D, Γ) is a pair $I = (X, T)$ where

- $X = \{x_1, \dots, x_{|X|}\}$ is a finite set of variables, and
- T is a set of constraints involving the variables in X , such that each constraint relation is an element of Γ . Formally, T is a set of pairs (\mathbf{x}, R) , where $R \in \Gamma$ and $\mathbf{x} \in X^{\text{ar}(R)}$ is the tuple of variables the relation R is applied to.

An assignment $x_1 = a_1, \dots, x_n = a_n$ of the variables to elements of D is a *solution* to that instance if, for each pair $(\mathbf{x}, R) \in T$, R is satisfied by the tuple \mathbf{x} after replacing each x_i in \mathbf{x} with a_i . The *linear programming relaxation* of a k -ary relation $R = \{r_1, \dots, r_{|R|}\}$ over a finite domain $D = \{d_1, \dots, d_{|D|}\}$ is the polyhedron in $(\mathbb{R}^{|D|})^k$ defined by the set of all points

$$\begin{aligned} & \left((v_1)_{d_1}, (v_1)_{d_2}, \dots, (v_1)_{d_{|D|}} \right), \\ & \left((v_2)_{d_1}, (v_2)_{d_2}, \dots, (v_2)_{d_{|D|}} \right), \\ & \quad \vdots \\ & \left((v_k)_{d_1}, (v_k)_{d_2}, \dots, (v_k)_{d_{|D|}} \right) \end{aligned}$$

for which there exist reals $p_{r_1}, \dots, p_{r_{|R|}}$ such that

$$\begin{aligned} 0 & \leq p_{r_1}, \dots, p_{r_{|R|}} \leq 1, \\ \sum_{r \in R} p_r & = 1, \end{aligned}$$

and

$$(v_i)_{d_j} = \sum_{r \in R | r_i = d_j} p_r$$

for all $1 \leq i \leq k$ and $1 \leq d \leq |D|$. The *linear programming relaxation* of an instance $I = (X, T)$ of a CSP over a domain D is the polyhedron in $(\mathbb{R}^{|D|})^{|X|}$ defined by the set of all points

$$\left(\begin{array}{c} ((x_1)_{d_1}, (x_1)_{d_2}, \dots, (x_1)_{d_{|D|}}), \\ ((x_2)_{d_1}, (x_2)_{d_2}, \dots, (x_2)_{d_{|D|}}), \\ \vdots \\ ((x_{|X|})_{d_1}, (x_{|X|})_{d_2}, \dots, (x_{|X|})_{d_{|D|}}) \end{array} \right)$$

such that for each pair $(\mathbf{x}, R) \in T$, \mathbf{x} lies in R 's linear programming relaxation when each variable in \mathbf{x} is replaced by its corresponding $|D|$ -tuple. A *fractional solution* of an instance of a CSP is a point inside its linear programming relaxation. We say that a CSP is *solved* by its linear programming relaxation if, for every instance I of the CSP, the existence of a fractional solution implies the existence of a solution.

Let $P = (D, \Gamma)$ be a constraint satisfaction problem, let $X = \{x_1, \dots, x_{|X|}\}$ be a set of variables, and let

$$I = (X, \{(\mathbf{x}_1, R_1), \dots, (\mathbf{x}_{|R|}, R_{|R|})\})$$

be an instance of P . A *step* $s = (k, (i, j))$ in I is defined to be a constraint relation (\mathbf{x}_k, R_k) and a pair of integers $1 \leq i, j \leq \text{ar}(R_k)$; we think of the step s as connecting the variable $(\mathbf{x}_k)_i$ to the variable $(\mathbf{x}_k)_j$. A *cycle* p in I is a finite sequence of steps

$$s_1, \dots, s_{|p|} = (k_1, (i_1, j_1)), \dots, (k_{|p|}, (i_{|p|}, j_{|p|}))$$

for which $(\mathbf{x}_{k_1})_{j_1} = (\mathbf{x}_{k_2})_{i_2}$, $(\mathbf{x}_{k_2})_{j_2} = (\mathbf{x}_{k_3})_{i_3}$, $(\mathbf{x}_{k_{|p|-1}})_{j_{|p|-1}} = (\mathbf{x}_{k_{|p|}})_{i_{|p|}}$, and $(\mathbf{x}_{k_{|p|}})_{j_{|p|}} = (\mathbf{x}_{k_1})_{i_1}$, where the subscript n on each $(\mathbf{x}_i)_n$ represents the variable in X corresponding to the n^{th} coordinate of \mathbf{x}_i . If B is a subset of D and $s = (k, (i, j))$ is a step in I , then we define the sum $B + s$ as

$$B + s := \{d \in D \mid \exists \mathbf{a} \in R_k \text{ s.t. } \mathbf{a}_i \in B \wedge \mathbf{a}_j = d\}$$

and the sum $B - s$ as

$$B - s := \{d \in D \mid \exists \mathbf{a} \in R_k \text{ s.t. } \mathbf{a}_j \in B \wedge \mathbf{a}_i = d\},$$

where the subscript \mathbf{a}_i denotes the i^{th} coordinate of the tuple \mathbf{a} . If $p = s_1, \dots, s_{|p|}$ is a cycle in I , then we define the sum $B + p$ as

$$B + p := B + s_1 + \dots + s_p$$

and the sum $B - p$ as

$$B - p := B - s_p - \dots - s_1.$$

2.2 Clones

Let D be some domain. An *operation* f is a function $f : D^k \rightarrow D$ for some positive integer k , known as its *arity*. An operation with arity 1 is *unary*, an operation with arity 2 is *binary*, and an operation with arity 3 is *ternary*. In general, an operation f with arity k is *k-ary*, and its arity is denoted $\text{ar}(f)$. The output of an operation f with inputs x_1, x_2, \dots, x_k is denoted

$$f(x_1, x_2, \dots, x_k).$$

Let $f : D^k \rightarrow D$ be an operation. We extend f to an operation on vectors in D^n by applying it coordinatewise, i.e.,

$$f \left(\begin{array}{c} [(a_1)_1] \\ [(a_1)_2] \\ \vdots \\ [(a_1)_n] \end{array}, \begin{array}{c} [(a_2)_1] \\ [(a_2)_2] \\ \vdots \\ [(a_2)_n] \end{array}, \dots, \begin{array}{c} [(a_k)_1] \\ [(a_k)_2] \\ \vdots \\ [(a_k)_n] \end{array} \right) := \begin{array}{c} [f((a_1)_1, (a_2)_1, \dots, (a_k)_1)] \\ [f((a_1)_2, (a_2)_2, \dots, (a_k)_2)] \\ \vdots \\ [f((a_1)_n, (a_2)_n, \dots, (a_k)_n)] \end{array}.$$

A k -ary operation $f : D^k \rightarrow D$ *preserves* a relation R if

$$f(r_1, \dots, r_k) \in R$$

for all $r_1, \dots, r_k \in R$ (note that the arity of f has nothing to do with the arity of R). The relation *generated* by f with generators $\mathbf{x}_1, \mathbf{x}_2, \dots$ is the smallest relation containing $\mathbf{x}_1, \mathbf{x}_2, \dots$ that f preserves, and is denoted

$$\text{Sg}_f(\mathbf{x}_1, \mathbf{x}_2, \dots).$$

For every $k \in \mathbb{Z}^+$ and integer $1 \leq i \leq k$, the *projection operation* $\pi_i^k : D^k \rightarrow D$ over a domain D is defined as

$$\pi_i^k(d_1, \dots, d_k) := d_i.$$

Note that π_1^1 is the identity operation over D . A *clone* over a domain D is a set \mathcal{O} of finite-arity operations that contains every projection operation over D and is closed under multiple composition; that is, if $f \in \mathcal{O}$ is an m -ary operation and $g_1, \dots, g_m \in \mathcal{O}$ are n -ary operations, then the operation

$$h(x_1, \dots, x_n) := f(g_1(x_1, \dots, x_n), \dots, g_m(x_1, \dots, x_n))$$

is also in \mathcal{O} . Note that \mathcal{O} is closed under any “natural” way of composing operations because every projection operation is in \mathcal{O} . A clone is *compatible* with a cyclic automorphism if there is a renaming defined by a cyclic permutation of the domain elements such that the renamed clone contains the same operations as the original clone.

A *subclone* \mathcal{O}' of a clone \mathcal{O} is a subset of \mathcal{O} that is a clone. The subclone is *proper* if $\mathcal{O}' \neq \mathcal{O}$. The clone *generated* by a set of operations $\{f_1, f_2, \dots\}$ is the smallest clone containing $\{f_1, f_2, \dots\}$ and is denoted

$$\langle f_1, f_2, \dots \rangle.$$

An operation f over a domain D is *idempotent* if

$$f(\underbrace{x, \dots, x}_{\text{ar}(f) \text{ } x\text{'s}}) = x$$

for all $x \in D$. A clone \mathcal{O} is *idempotent* if every operation in \mathcal{O} is idempotent. A k -ary operation f is *symmetric* if

$$f(x_1, \dots, x_k) = f(x_{\sigma(1)}, \dots, x_{\sigma(k)})$$

for all $x_1, \dots, x_k \in D$ and permutations $\sigma : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$. We call a clone *round* if it contains symmetric operations of every arity. We call a clone over a domain D *semi-round* if it contains symmetric operations of arities $1, 2, \dots, |D|$. The relation *generated* by a clone \mathcal{O} with generators $\mathbf{x}_1, \mathbf{x}_2, \dots$ is the smallest relation that every operation in \mathcal{O} preserves, and it is denoted

$$\text{Sg}_{\mathcal{O}}(\mathbf{x}_1, \mathbf{x}_2, \dots).$$

Let P be a property of a clone. We say a clone is *minimal* with respect to property P if it does not contain a proper subclone with property P . For example, a clone is *minimally round* if it does not contain a proper round subclone, and a clone is *minimally semi-round* if it does not contain a proper semi-round subclone. Lastly, a *minimal idempotent round clone* is an idempotent clone that is minimally round, and a *minimal idempotent semi-round clone* is an idempotent clone that is minimally semi-round.

2.3 Algebraic concepts

An *algebraic structure* $\mathbb{A} = (A, f_1, f_2, \dots)$, also known as an *algebra*, is a domain A , which we call the *underlying set* of \mathbb{A} , with some operations $f_i : A^{k_i} \rightarrow A$, which we call the *basic operations* of \mathbb{A} . Algebraic structures will always be written in blackboard bold. The sequence of arities k_1, k_2, \dots is called the *signature* of the algebra \mathbb{A} .

Given an algebraic structure $\mathbb{A} = (A, f_1, f_2, \dots)$, we define the *power* $\mathbb{A}^m = (A^m, f_1, f_2, \dots)$ to be an algebraic structure of the same signature, where each f_i acts coordinatewise on A^m . A *subalgebra* \mathbb{B} of an algebra \mathbb{A} , denoted as $\mathbb{B} \leq \mathbb{A}$, consists of a subset $B \subseteq A$ closed under the basic operations of \mathbb{A} , and basic operations which are restrictions of the basic operations of \mathbb{A} to B . If B is any subset of A^m , then we define the *subalgebra generated by B* , denoted $\text{Sg}_{\mathbb{A}^m}(B)$, to be the smallest subalgebra of \mathbb{A}^m which contains B .

The *clone* of an algebraic structure \mathbb{A} , written $\text{Clo}(\mathbb{A})$, is the clone generated by the basic operations of \mathbb{A} . When we study algebraic structures, we will mainly be interested in properties which only depend on their clones instead of their basic operations.

A *relation* $\mathbb{R} \leq \mathbb{A}^m$ on an algebra \mathbb{A} will always refer to a subalgebra of \mathbb{A}^m for some integer m , known as the *arity* of \mathbb{R} . Alternatively, a relation $\mathbb{R} \leq \mathbb{A}^m$ is a subset of A^m which is *compatible* with the basic operations of \mathbb{A} ; that is, for each basic operation f_i of arity k_i and for every choice of k_i tuples $\mathbf{a}^1, \dots, \mathbf{a}^{k_i} \in \mathbb{R}$, we have $f(\mathbf{a}^1, \dots, \mathbf{a}^{k_i}) \in \mathbb{R}$, where f acts coordinatewise. Whether or not a given set $R \subseteq A^m$ defines a relation compatible with the basic operations of \mathbb{A} only depends on the clone of \mathbb{A} .

A *congruence* θ on an algebraic structure \mathbb{A} is an equivalence relation $\theta \leq \mathbb{A}^2$ compatible with the basic operations of \mathbb{A} , and the *quotient* \mathbb{A}/θ is an algebraic structure with domain A/θ with the basic operations defined in the natural way.

If $\mathbb{R} \leq \mathbb{A}^m$ is a relation and I is a subset of the coordinates $\{1, \dots, m\}$, then we define the *existential projection* $\pi_I(\mathbb{R})$ as

$$\pi_I(\mathbb{R}) := \{\mathbf{x} \in A^I \mid \exists \mathbf{y} \in \mathbb{R} \text{ s.t. } \mathbf{y}_i = \mathbf{x}_i \ \forall i \in I\}.$$

For brevity, define $\pi_{i,j,\dots}(\mathbb{R})$ to be $\pi_{\{i,j,\dots\}}(\mathbb{R})$. A relation $\mathbb{R} \leq \mathbb{A}^m$ is *subdirect*, denoted $\mathbb{R} \leq_{sd} \mathbb{A}^m$, if the i^{th} projection $\pi_i(\mathbb{R})$ is equal to A for every integer $1 \leq i \leq m$.

If $\mathbb{R}, \mathbb{S} \leq \mathbb{A}^2$ are binary relations, then we define their *composition* as

$$\mathbb{R} \circ \mathbb{S} := \{(x, z) \in A^2 \mid \exists y \in A \text{ s.t. } (x, y) \in \mathbb{R} \wedge (y, z) \in \mathbb{S}\}.$$

We define the *reverse* of the binary relation \mathbb{R} , denoted \mathbb{R}^- , as

$$\mathbb{R}^- := \{(y, x) \in A^2 \mid (x, y) \in \mathbb{R}\}.$$

If $\mathbb{R} \leq_{sd} \mathbb{A}^2$, then we define the *linking congruence* of \mathbb{R} on the first coordinate to be the congruence

$$\bigcup_{n \geq 1} (\mathbb{R} \circ \mathbb{R}^-)^{\circ n}.$$

The linking congruence on the second coordinate is defined similarly, with \mathbb{R} and \mathbb{R}^- swapped. If $B \subseteq A$ is a set and $\mathbb{R} \leq \mathbb{A}^2$ is a binary relation, then we define the *sum* $B + \mathbb{R}$ as

$$B + \mathbb{R} := \{y \in A \mid \exists x \in B \text{ s.t. } (x, y) \in \mathbb{R}\}.$$

and we define the *difference* $B - \mathbb{R}$ as

$$B - \mathbb{R} := B + \mathbb{R}^- = \{x \in A \mid \exists y \in B \text{ s.t. } (x, y) \in \mathbb{R}\}.$$

If \mathcal{A} is a collection of algebraic structures which all have the same signature, then we define $P(\mathcal{A})$ to be the collection of all products of algebras in \mathcal{A} , $S(\mathcal{A})$ to be the collection of all subalgebras of algebras in \mathcal{A} , and $H(\mathcal{A})$ to be the collection of all homomorphic images of algebras in \mathcal{A} (which is the collection of all algebras which are isomorphic to the quotient of some algebra \mathbb{A} in \mathcal{A} by some congruence on A). If $\mathbb{B} \in HS(\mathbb{A})$, then we call \mathbb{B} a *subquotient* of \mathbb{A} .

A clone is *Taylor* if it contains idempotent operations that satisfy some functional equation that cannot be satisfied by projection operations. An algebra is called *Taylor* if its clone is Taylor. By Birkhoff's *HSP* theorem, an idempotent algebra is Taylor if and only if $HSP(\mathbb{A})$ does not contain a two-element algebra with each of its basic operations equal to a projection.

2.4 Miscellaneous

For sets A and B , the set $A + B$ is defined as

$$A + B := \{a + b \mid a \in A, b \in B\}.$$

The function $\text{sgn} : \mathbb{R} \rightarrow \mathbb{R}$ is defined as

$$\text{sgn}(x) := \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ 1 & \text{if } x > 0. \end{cases}$$

For a $(k - 1)$ -ary operation f , we define $c_f : D^k \rightarrow D^k$ as

$$c_f((x_1, \dots, x_k)) := (f(x_2, x_3, \dots, x_{k-1}, x_k), \\ f(x_1, x_3, \dots, x_{k-1}, x_k), \\ f(x_1, x_2, \dots, x_{k-1}, x_k), \\ \vdots, \\ f(x_1, x_2, x_3, \dots, x_{k-1})).$$

We also define

$$c_f(x_1, \dots, x_k) := c_f((x_1, \dots, x_k))$$

for convenience. Additionally, given a tuple $\mathbf{x} = (x_1, \dots, x_k)$ and a k -ary operation f , define

$$f(\mathbf{x}) := f(x_1, \dots, x_k).$$

We say an operation f over a domain D *acts like a height-1 semilattice* over a subset $D' \subseteq D$ of its domain if it is idempotent and

$$f(x_1, \dots, x_{\text{ar}(f)})$$

is the same fixed value $c \in D'$ over all non-constant tuples $(x_1, \dots, x_{\text{ar}(f)}) \in (D')^{\text{ar}(f)}$. We say a clone *acts like a height-1 semilattice* over a subset $D' \subseteq D$ of its domain if all its operations act like a height-1 semilattice over D' , and the constant c is the same across all operations.

We say a binary operation f over a domain D *acts linearly* over a subset $D' \subseteq D$ if $|D'|$ is odd and the domain elements of f , when restricted to the domain D' , can be renamed such that the new operation $f' : (\mathbb{Z}/|D'|\mathbb{Z})^2 \rightarrow \mathbb{Z}/|D'|\mathbb{Z}$ satisfies

$$f'(x, y) = \frac{x + y}{2} \pmod{|D'|}$$

for all $x, y \in \mathbb{Z}/|D'|\mathbb{Z}$.

A binary operation f over a domain D is a *semilattice operation* if there exists a poset on D such that f represents “join”; that is,

$$f(x, y) = x \wedge y$$

for all $x, y \in D$.

Lastly, for all $a, b, c \in \{-, 0, +\}$ we define f_{abc} to be the symmetric binary operation

$$\begin{array}{c|ccc} f_{abc} & - & 0 & + \\ \hline - & - & c & b \\ 0 & c & 0 & a \\ + & b & a & + \end{array},$$

and for $a, b, c, d \in \{0, 1, 2, 3\}$ we define f_{abcdef} to be the symmetric binary operation

f_{abcdef}	0	1	2	3
0	0	a	b	c
1	a	1	d	e
2	b	d	2	f
3	c	e	f	3

3 Related Results

Recall that our goal is to verify Conjecture 1 for all clones over a domain of size 4 or less; that is, our goal is to prove that all semi-round clones over a domain of size at most 4 are round.

The relevance of this problem to CSPs has been demonstrated in [9]. The authors prove that a CSP is solved by its linear programming relaxation if and only if the clone of all operations preserving the relations defining the CSP is round.

A related result about cyclic operations is proved in [1]. They prove that every Taylor algebra contains a cyclic operation of every prime arity greater than the size of its domain.

It turns out that the $|D|$ bound is tight when $|D|$ is prime.

Proposition 5. *Let D be a domain with prime cardinality. Then there exists a clone \mathcal{O} over D containing symmetric operations of arities $1, 2, \dots, |D| - 1$ that is not round.*

Proof. By construction, we can force \mathcal{O} to be compatible a cyclic automorphism. Let p be the cardinality of the domain, let $D = \mathbb{Z}/p\mathbb{Z}$, and define the k -ary operation f_k for $k \in \{1, 2, \dots, |D| - 1\}$ to be any symmetric operation satisfying

$$f_k(x_1 + c, \dots, x_k + c) = f_k(x_1, \dots, x_k) + c$$

for all $(x_1, \dots, x_k) \in D^k$ and $c \in D$. For instance, we can take

$$f_k(x_1, \dots, x_k) := \frac{x_1 + \dots + x_k}{k} \pmod{p}.$$

Then

$$\langle f_1, f_2, \dots, f_{p-1} \rangle$$

is compatible with a cyclic automorphism, so no p -ary symmetric operation f_p can exist, as

$$f_p(x_1, \dots, x_p)$$

cannot be preserved by the automorphism. □

Remark 6 (Lemma 4 of [7]). The theorem statement is false if $|D|$ is not required to be prime; the construction in [7] proves that a 4-ary symmetric operation $f_4(w, x, y, z)$ can be constructed from a binary symmetric operation $f_2(x, y)$ and a ternary symmetric operation $f_3(x, y, z)$:

$$\begin{aligned} f_4(w, x, y, z) = f_3(f_2(f_2(w, x), f_2(y, z)), \\ f_2(f_2(w, y), f_2(x, z)), \\ f_2(f_2(w, z), f_2(x, y))) \end{aligned}$$

is symmetric.

4 Preliminary Theorems

These results will help us classify the semi-round clones over a domain of size at most four, up to a renaming of the domain elements.

Theorem 7. *Let Σ be a set of functional equations. Then every clone \mathcal{O} containing operations that simultaneously satisfy each functional equation in Σ contains a minimal subclone containing operations that simultaneously satisfy each functional equation in Σ .*

Proof. For a set Σ_1 of functional equations, let property P_{Σ_1} of a clone \mathcal{O} denote the assertion that \mathcal{O} contains a set of operations simultaneously satisfying each equation in Σ_1 .

For any finite subset $\Sigma' \subseteq \Sigma$, there can only be finitely many combinations of operations in \mathcal{O} that satisfy Σ' , since the domain is finite. Therefore, the intersection of every chain

$$\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$$

of clones with property $P_{\Sigma'}$ also has property $P_{\Sigma'}$ for every finite subset $\Sigma' \subseteq \Sigma$. By the logical compactness theorem, \mathcal{O} has property P_{Σ} if \mathcal{O} has property $P_{\Sigma'}$ for every finite subset $\Sigma' \subseteq \Sigma$, which allows us to conclude that the intersection of every chain

$$\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$$

of clones with property P_{Σ} also has property P_{Σ} . Therefore, every chain in the poset of subclones of \mathcal{O} with property P_{Σ} , ordered by reverse inclusion, has an upper bound, so Zorn's lemma implies the existence of a minimal clone with property P_{Σ} . \square

Corollary 8. *Every round clone contains a minimal round subclone.*

Proof. Let Σ be the set of functional equations

$$\begin{aligned} f_2(x, y) &= f_2(y, x) \\ f_3(x, y, z) &= f_3(x, z, y) = f_3(y, x, z) = f_3(y, z, x) = f_3(z, x, y) = f_3(z, y, x) \\ &\vdots \end{aligned}$$

where each line asserts the existence of a symmetric operation of some arity. Then P_{Σ} is equivalent to roundness, so there exists a minimal round subclone by Theorem 7. \square

Corollary 9. *Every semi-round clone contains a minimal semi-round subclone.*

Proof. This is the same as the proof of Corollary 8, except the set of equations Σ asserts the existence of symmetric operations of arities $1, 2, \dots, |D|$, where D is the domain, instead of symmetric operations of every arity. \square

Theorem 10. *Suppose \mathcal{O} is a minimal counterexample to Conjecture 1 over the smallest possible domain. Then \mathcal{O} is idempotent.*

Proof. This is a corollary of a folklore result that seems to have first shown up in the context of CSPs in [5], but we reproduce it here. It suffices to show that the only unary operation in a minimal counterexample \mathcal{O} is π_1^1 . To do this, let the domain be D and let f_k be a k -ary symmetric operation in \mathcal{O} for each $k \in \{1, 2, \dots, |D|\}$. If there is a non-identity unary operation $u(x) \in \mathcal{O}$, then either some unary operation is not injective or all unary operations are permutations.

If all unary operations are permutations, then one can find idempotent symmetric operations of arities $1, 2, \dots, |D|$. To do this, define

$$u_k(x) = f_k(\underbrace{x, \dots, x}_{k \text{ times}}),$$

which must be a permutation. Then $u_k^{N_k}(x) = x$ for some positive integer N_k . Define

$$g_k(x_1, \dots, x_k) = f_k \left(u_k^{N_k-1}(x_1), \dots, u_k^{N_k-1}(x_k) \right)$$

and note that g_k is idempotent for each k . Finally,

$$\langle g_1, \dots, g_k \rangle \subsetneq \mathcal{O}$$

is an idempotent semi-round proper subclone of \mathcal{O} which is also a counterexample to the conjecture, contradicting the minimality of \mathcal{O} .

If there is a unary operation $u(x)$ that is not injective, then $u^{2N}(x) = u^N(x)$ for some positive integer N , and $u^N(a) = u^N(b) = a$ for some distinct $a, b \in D$. Define

$$g_k(x_1, \dots, x_k) = u^N \left(f_k \left(u^N(x_1), \dots, u^N(x_k) \right) \right)$$

and note that each g_k is a symmetric operation that acts on a and b the same way; that is, replacing a with b or b with a in g_k 's input will not change its output. Since

$$\langle g_1, \dots, g_k \rangle \subseteq \mathcal{O}$$

is effectively a clone on the domain $D \setminus \{b\}$, \mathcal{O} is not minimal with respect to domain size, contradiction. \square

Proposition 11. *Let f be a semilattice operation. Then $\langle f \rangle$ is round.*

Proof. Consider the poset determined by f . By induction, the symmetric k -ary operation f_k that returns the greatest upper bound of its k inputs is in the clone.

Indeed, $f_1 = \pi_1^1$ and $f_2 = f$. Now

$$f_{k+1}(x_1, \dots, x_{k+1}) = f_2(f_k(x_1, \dots, x_k), x_{k+1})$$

by induction. \square

Theorem 12. *Let \mathcal{O} be an idempotent round clone over a domain D , and suppose some binary operation in \mathcal{O} acts like a height-1 semilattice over some subset $D' \subseteq D$. Then \mathcal{O} contains a round subclone that acts like a height-1 semilattice over D' .*

Proof. Let t_2 be the given binary operation, and for each $k \in \mathbb{Z}^+$, let f_k denote a k -ary symmetric operation in \mathcal{O} . Then the binary operation

$$g_2(x, y) = f_2(t_2(x, y), t_2(y, x))$$

is symmetric and acts like a height-1 semilattice over D' . By induction, the k -ary operation

$$g_k(x_1, \dots, x_k) = f_k \left(c_{g_{k-1}}^2(x_1, \dots, x_k) \right)$$

is symmetric and acts like a height-1 semilattice over D' . (Recall that $c_{g_{k-1}}$ is the function that takes a k -tuple and returns the k -tuple whose i^{th} entry is g_{k-1} applied on the $k-1$ variables not in the i^{th} coordinate, for all $1 \leq i \leq k$.) Then

$$\langle g_1, g_2, \dots \rangle$$

is a round subclone of \mathcal{O} that acts like a height-1 semilattice over D' , as desired. \square

Corollary 13. *Let \mathcal{O} be an idempotent semi-round clone over a domain D , and suppose some binary operation in \mathcal{O} acts like a height-1 semilattice over some subset $D' \subseteq D$. Then \mathcal{O} contains a semi-round subclone that acts like a height-1 semilattice over D' .*

Proof. This is the essentially the same as the proof of Theorem 12, except the induction stops after constructing the $|D|$ -ary symmetric operation. \square

Theorem 14. *Let \mathcal{O} be an idempotent clone over a domain D with a binary symmetric operation f_2 that acts linearly on a subset $D' \subseteq D$ with odd prime cardinality, and suppose \mathcal{O} contains a $|D'|$ -ary symmetric operation $f_{|D'|}$. Then one can find symmetric operations $g_1, g_2, \dots \in \mathcal{O}$ and a constant c such that*

$$g_k(x_1, \dots, x_k) = \begin{cases} x & \text{if } \{x_1, x_2, \dots, x_k\} = \{x\} \\ c & \text{otherwise} \end{cases}$$

for all $k \in \mathbb{Z}^+$ and tuples $(x_1, x_2, \dots, x_k) \in (D')^k$.

Proof. We prove this by induction on k . Define $c := f_{|D'|}(d_1, \dots, d_{|D'|})$ where $D' = \{d_1, \dots, d_{|D'|}\}$, and for $k = 1$ note $g_1 = \pi_1^1$. For the $k = 2$ construction, assume $D' = \{0, 1, \dots, p-1\}$ for some odd prime p . Since $f_2(x, y) = \frac{x+y}{2} \pmod{p}$ over D' , one can construct any operation that acts like

$$\frac{a}{2^b}x + \frac{2^b - a}{2^b}y \pmod{p}$$

over D' by composing f_2 with itself for any $a, b \in \mathbb{Z}^+$ with $0 \leq a \leq 2^b$; in particular, by choosing $b = p-1$, the operation

$$ax - (a-1)y \pmod{p}$$

can be constructed using only f_2 for any $a \in \mathbb{Z}$, by Fermat's Little Theorem. Finally, define

$$g_2(x, y) := f_k(x, 2x - y, 3x - 2y, \dots, px - (p-1)y) \pmod{p};$$

this works because

$$\{x, 2x - y, 3x - 2y, \dots, px - (p-1)y\} = \{0, 1, \dots, p-1\} \pmod{p}$$

whenever $x - y \neq 0$, by the primality of p . To construct g_{k+1} given g_k , define

$$g_{k+1}(x_1, x_2, \dots, x_{k+1}) := g_2(g_k(x_1, x_2, \dots, x_k), g_k(f_2(x_1, x_{k+1}), f_2(x_2, x_{k+1}), \dots, f_2(x_k, x_{k+1}))).$$

This works by casework on whether or not $x_1 = x_2 = \dots = x_k$; if they are not all equal, then both arguments of g_2 are c , and if they are all equal, then both arguments of g_2 are distinct elements of D' . \square

Lemma 15. *Suppose a semi-round clone \mathcal{O} over a domain D with $|D| \leq 4$ contains a binary symmetric operation that acts linearly on a three-element set $D' \subseteq D$. Then \mathcal{O} is not minimally semi-round.*

Proof. Let c be the constant guaranteed by Theorem 14. If $c \in D'$ then we are done by Corollary 13. Thus we can assume $D = \{0, 1, 2, 3\}$, $D' = \{0, 1, 2\}$, and $c = 3$.

The binary symmetric operation that acts linearly over D' must be of the form

f_{21a0bc}	0	1	2	3
0	0	2	1	a
1	2	1	0	b
2	1	0	2	c
3	a	b	c	3

and g_2 must be of the form

f_{33d3ef}	0	1	2	3
0	0	3	3	d
1	3	1	3	e
2	3	3	2	f
3	d	e	f	3

If at least two of d, e, f are equal to 3, then f_{33d3ef} acts like a height-1 semilattice over a subset $D' \subset D$ with size 3, so we can apply Corollary 13 to show that \mathcal{O} is not minimally semi-round. The set D' contains 3 and the corresponding domain elements for the two of d, e, f that are equal to 3.

If there is at most one 3 among $d, e,$ and f , we can assume there is at most one 3 among $a, b,$ and c by Corollary 13 on a two-element subset of D ; otherwise \mathcal{O} is not minimally semi-round. If $a \in \{1, 2\}$ then one can check that

$$\begin{bmatrix} a \\ a \end{bmatrix} \in \text{Sg}_{\mathcal{O}}\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \end{bmatrix}\right).$$

Thus, some binary operation acts like a height-1 semilattice over $\{1, 2\}$, so \mathcal{O} is not minimally semi-round. Similarly, if $b \in \{0, 2\}$ or $c \in \{0, 1\}$ then \mathcal{O} is not minimally semi-round.

If exactly one of a, b and c is equal to 3, then without loss of generality assume $c = 3$, so the binary operation that acts linearly over D' is forced to be f_{210013} :

f_{210013}	0	1	2	3
0	0	2	1	0
1	2	1	0	1
2	1	0	2	3
3	0	1	3	3

Since we can assume $a \notin \{1, 2\}$ and $b \notin \{0, 2\}$, by Corollary 13 either \mathcal{O} is not minimally semi-round, or the binary symmetric operation f_{330313} guaranteed by Theorem 14 is in \mathcal{O} :

f_{330313}	0	1	2	3
0	0	3	3	0
1	3	1	3	1
2	3	3	2	3
3	0	1	3	3

Then one can check that

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} \in \text{Sg}_{\mathcal{O}}\left(\begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 2 \\ 1 \end{bmatrix}\right)$$

so \mathcal{O} is not minimally semi-round by Corollary 13.

If $3 \notin \{a, b, c\}$, then the binary operation that acts linearly over D' is forced to be f_{210012} :

f_{210012}	0	1	2	3
0	0	2	1	0
1	2	1	0	1
2	1	0	2	2
3	0	1	2	3

By Corollary 13 either \mathcal{O} is not minimally semi-round, or the binary symmetric operation f_{330312} guaranteed by Theorem 14 is in \mathcal{O} :

f_{330312}	0	1	2	3
0	0	3	3	0
1	3	1	3	1
2	3	3	2	2
3	0	1	2	3

Then

$$f_{330312}(f_{330312}(x, y), f_{330312}(f_{330312}(x, z), f_{330312}(y, z)))$$

is a ternary symmetric operation, so

$$\langle f_{330312} \rangle \subsetneq \mathcal{O}$$

is a semi-round clone by Remark 6, and one can check that it does not contain a symmetric binary operation other than f_{330312} . \square

5 Classification

In this section, we present a complete catalogue of every minimal idempotent semi-round clone over a domain of size at most 4, up to a renaming of the domain elements. Additionally, we prove that every semi-round clone over a domain of size 4 is also round. Non-idempotent clones are not considered, as Theorem 10 states that a minimal counterexample to the conjecture is idempotent, and non-minimal round clones are not considered, as Corollary 9 guarantees the existence of a minimally semi-round subclone of every semi-round clone.

5.1 Domain of Size 1

There is only one clone over a domain of size 1, which is both semi-round and round.

5.2 Domain of Size 2

There is only one minimal semi-round clone over a domain of size 2, up to a renaming of the domain elements. To prove this, let $D = \{0, 1\}$ and let $f_2(x, y)$ be the clone's binary symmetric operation. We can, without loss of generality, assume that $f_2(x, y) = xy$:

$$\begin{array}{c|cc} f_2 & 0 & 1 \\ \hline 0 & 0 & 0 \\ 1 & 0 & 1 \end{array}.$$

Then $\langle f_2 \rangle$ is round by Proposition 11 because f_2 is a semilattice operation, so it is the unique minimal semi-round clone over $\{0, 1\}$, up to a renaming of the domain elements.

5.3 Domain of Size 3

Assume $D = \{-1, 0, 1\}$, which we abbreviate as $\{-, 0, +\}$. By a finite case check, any minimal clone must contain, up to a renaming of the domain elements, one of following symmetric binary operations:

$$\begin{array}{c} \begin{array}{c|ccc} f_{000} & - & 0 & + \\ \hline - & - & 0 & 0 \\ 0 & 0 & 0 & 0 \\ + & 0 & 0 & + \end{array} & \begin{array}{c|ccc} f_{++0} & - & 0 & + \\ \hline - & - & 0 & + \\ 0 & 0 & 0 & + \\ + & + & + & + \end{array} & \begin{array}{c|ccc} f_{+00} & - & 0 & + \\ \hline - & - & 0 & 0 \\ 0 & 0 & 0 & + \\ + & 0 & + & + \end{array} \\ \\ \begin{array}{c|ccc} f_{00+} & - & 0 & + \\ \hline - & - & + & 0 \\ 0 & + & 0 & 0 \\ + & 0 & 0 & + \end{array} & \begin{array}{c|ccc} f_{+0-} & - & 0 & + \\ \hline - & - & - & 0 \\ 0 & - & 0 & + \\ + & 0 & + & + \end{array} & \begin{array}{c|ccc} f_{+-0} & - & 0 & + \\ \hline - & - & 0 & - \\ 0 & 0 & 0 & + \\ + & - & + & + \end{array} \\ \\ \begin{array}{c|ccc} f_{-0+} & - & 0 & + \\ \hline - & - & + & 0 \\ 0 & + & 0 & - \\ + & 0 & - & + \end{array} \end{array}.$$

To determine the minimal round clones, we casework on the symmetric binary operation.

- $\langle f_{000} \rangle$ is minimally round by Proposition 11, since f_{000} is a semilattice operation.
- $\langle f_{++0} \rangle$ is minimally round by Proposition 11, since f_{++0} is a semilattice operation.

- $\langle f_{++0} \rangle \subseteq \langle f_{+00} \rangle$, so f_{+00} does not need to be considered, since

$$f_{+00}(f_{+00}(x, f_{+00}(x, y)), f_{+00}(y, f_{+00}(x, y))) = f_{++0}(x, y).$$

- $\langle f_{++0} \rangle \subseteq \langle f_{00+} \rangle$, so f_{00+} does not need to be considered, since

$$f_{00+}(f_{00+}(x, f_{00+}(x, y)), f_{00+}(y, f_{00+}(x, y))) = f_{++0}(x, y).$$

- $\langle f_{+0-} \rangle$ is minimally semi-round and minimally round, since all operations of the form

$$f_k(x_1, \dots, x_k) = \text{sgn}(x_1 + \dots + x_k)$$

are in $\langle f_{-+0} \rangle$. For $k = 1$ take $f_1 = \pi_1^1$, and for $k = 2$ take $f_2 = f_{-+0}$. To construct f_{k+1} inductively, the identity

$$\begin{aligned} f_{k+1}(x_1, \dots, x_{k+1}) &= f_2(f_k(f_k(x_2, x_3, \dots, x_k, x_{k+1}), \\ &\quad f_k(x_1, x_3, \dots, x_k, x_{k+1}), \\ &\quad f_k(x_1, x_2, \dots, x_k, x_{k+1}), \\ &\quad \vdots, \\ &\quad f_k(x_1, x_2, x_3, \dots, x_{k+1})), f_k(x_1, \dots, x_k)) \end{aligned}$$

will suffice.

- $\langle f_{+-0} \rangle$ is not semi-round because the existence of a symmetric ternary operation is forbidden by the automorphism sending 0 to +, + to -, and - to 0. Hence, we will determine all minimal semi-round clones of the form $\langle f_{+-0}, g \rangle$, where g is some symmetric idempotent ternary operation.

By Corollary 13 there must be a symmetric ternary operation in \mathcal{O} of the form

$$f_3(x, y, z) = \begin{cases} f_{+-0}(d_1, d_2) & \text{if } \exists (d_1, d_2 \in D) \{x, y, z\} = \{d_1, d_2\} \\ c & \text{if } \{x, y, z\} = D \end{cases}$$

for some $c \in D$; without loss of generality assume $c = 0$. Now, $\langle f_{+-0}, f_3 \rangle$ is round because each operation of the form

$$f_k(x_1, \dots, x_k) := \begin{cases} f_{+-0}(d_1, d_2) & \text{if } \exists (d_1, d_2 \in D) \{x_1, \dots, x_k\} = \{d_1, d_2\} \\ 0 & \text{otherwise} \end{cases}$$

is in $\langle f_{+-0}, f_3 \rangle$. Note that the base cases $k \in \{1, 2, 3\}$ are true, so it suffices to construct f_k inductively. Define $\mathbf{x} = (x_1, \dots, x_{k+1})$, and define $\mathbf{y} = c_{f_k}^N(\mathbf{x})$, $\mathbf{y}' = c_{f_k}^N(f_2(\mathbf{x}, \mathbf{y}))$, and $\mathbf{y}'' = c_{f_k}^N(f_2(\mathbf{x}, \mathbf{y}'))$ for a sufficiently large integer N . Then

$$f_{k+1}(\mathbf{x}) = f_3(\mathbf{y}, \mathbf{y}', \mathbf{y}'')_1,$$

where the subscript denotes the first element of the tuple; this works because $\mathbf{y} = c_{f_k}^N(\mathbf{x})$ is a constant tuple for sufficiently large N , so $\{\mathbf{y}, \mathbf{y}', \mathbf{y}''\} = \{(-, \dots, -), (0, \dots, 0), (+, \dots, +)\}$ or $\mathbf{y} = \mathbf{y}' = \mathbf{y}''$.

- No clone containing f_{-0+} is minimally semi-round by Lemma 15, so this case does not need to be considered.

5.4 Domain of Size 4

The characterization of all minimal idempotent semi-round clones over a domain of size 4 will be done through casework on the symmetric binary operation it contains. There are $4^{\binom{4}{2}} = 4096$ such operations but they fall into 192 distinct equivalence classes under isomorphism; representative elements are listed below.

$f_{000000}, f_{000001}, f_{000002}, f_{000011}, f_{000012}, f_{000013}, f_{000021}, f_{000023}, f_{000033}, f_{000111},$
 $f_{000112}, f_{000121}, f_{000122}, f_{000123}, f_{000132}, f_{000321}, f_{001000}, f_{001001}, f_{001002}, f_{001003},$
 $f_{001010}, f_{001011}, f_{001012}, f_{001013}, f_{001020}, f_{001021}, f_{001022}, f_{001023}, f_{001030}, f_{001031},$
 $f_{001032}, f_{001033}, f_{001100}, f_{001101}, f_{001102}, f_{001103}, f_{001110}, f_{001111}, f_{001112}, f_{001113},$
 $f_{001120}, f_{001121}, f_{001122}, f_{001123}, f_{001130}, f_{001131}, f_{001132}, f_{001133}, f_{001200}, f_{001201},$
 $f_{001202}, f_{001203}, f_{001210}, f_{001211}, f_{001212}, f_{001213}, f_{001220}, f_{001221}, f_{001222}, f_{001223},$
 $f_{001230}, f_{001231}, f_{001232}, f_{001233}, f_{001300}, f_{001301}, f_{001302}, f_{001303}, f_{001310}, f_{001311},$
 $f_{001312}, f_{001313}, f_{001320}, f_{001321}, f_{001322}, f_{001323}, f_{001330}, f_{001331}, f_{001332}, f_{003000},$
 $f_{003001}, f_{003002}, f_{003011}, f_{003012}, f_{003013}, f_{003021}, f_{003100}, f_{003101}, f_{003102}, f_{003110},$
 $f_{003112}, f_{003113}, f_{003120}, f_{003121}, f_{003122}, f_{003300}, f_{003301}, f_{003302}, f_{003311}, f_{003312},$
 $f_{003321}, f_{011000}, f_{011001}, f_{011002}, f_{011010}, f_{011011}, f_{011012}, f_{011013}, f_{011020}, f_{011021},$
 $f_{011022}, f_{011023}, f_{011030}, f_{011031}, f_{011032}, f_{011120}, f_{011121}, f_{011122}, f_{011123}, f_{011130},$
 $f_{011131}, f_{011132}, f_{011120}, f_{011223}, f_{011230}, f_{011231}, f_{011320}, f_{011321}, f_{011322}, f_{012000},$
 $f_{012001}, f_{012002}, f_{012003}, f_{012010}, f_{012011}, f_{012012}, f_{012013}, f_{012020}, f_{012021}, f_{012022},$
 $f_{012030}, f_{012031}, f_{012032}, f_{012100}, f_{012101}, f_{012102}, f_{012103}, f_{012120}, f_{012121}, f_{012122},$
 $f_{012130}, f_{012131}, f_{012132}, f_{012200}, f_{012203}, f_{012210}, f_{012213}, f_{012230}, f_{012300}, f_{012301},$
 $f_{012302}, f_{012310}, f_{012311}, f_{012313}, f_{012320}, f_{012321}, f_{012330}, f_{013002}, f_{013010}, f_{013011},$
 $f_{013012}, f_{013021}, f_{013022}, f_{013102}, f_{013310}, f_{013321}, f_{022101}, f_{022301}, f_{022321}, f_{023321},$
 $f_{032000}, f_{032001}, f_{032020}, f_{032021}, f_{032030}, f_{032230}, f_{032320}, f_{032321}, f_{211000}, f_{211020},$
 $f_{211300}, f_{211301}.$

Many of these operations do not need to be considered because the clones they generate contain other binary operations. For example, only operations in the image of repeated composition of the map

$$f(a, b) \mapsto f(f(a, f(a, b)), f(b, f(a, b)))$$

need to be considered, which results in the following 37 operations:

$f_{000000}, f_{000002}, f_{000012}, f_{000013}, f_{000033}, f_{000111}, f_{000112}, f_{000123}, f_{000132}, f_{000321},$
 $f_{001030}, f_{001031}, f_{001032}, f_{001033}, f_{001130}, f_{001132}, f_{001133}, f_{001230}, f_{001231}, f_{001232},$
 $f_{001233}, f_{003012}, f_{003013}, f_{003112}, f_{003113}, f_{003312}, f_{003321}, f_{011231}, f_{011321}, f_{013310},$
 $f_{022101}, f_{023321}, f_{032230}, f_{032320}, f_{032321}, f_{211000}, f_{211020}.$

Under the map

$$f(a, b) \mapsto f(f(a, f(a, f(a, b))), f(b, f(b, f(a, b))))$$

the operations $f_{001030}, f_{001130}, f_{001230}, f_{011321}, f_{013310}$ can be removed from the list. Similarly, under the map

$$f(a, b) \mapsto f(f(a, f(b, f(a, b))), f(b, f(a, f(a, b))))$$

the operations f_{211000} and f_{022101} can be removed, and f_{211020} can be removed by considering the map

$$f(a, b) \mapsto f(a, f(a, f(b, f(b, f(a, b)))))$$

Lastly, the operations f_{000321} , f_{003321} , f_{023321} , f_{032230} , f_{032320} , and f_{032321} all act linearly over a three-element subset of their domain, so by Lemma 15 they don't need to be considered. Therefore, only the following 23 operations need to be considered for analysis:

$$f_{000000}, f_{000002}, f_{000012}, f_{000013}, f_{000033}, f_{000111}, f_{000112}, f_{000123}, f_{000132}, f_{001031},$$

$$f_{001032}, f_{001033}, f_{001132}, f_{001133}, f_{001231}, f_{001232}, f_{001233}, f_{003012}, f_{003013}, f_{003112},$$

$$f_{003113}, f_{003312}, f_{011231}.$$

Sixteen of these operations already generate minimal round clones. The following five operations generate round clones by Proposition 11 because they are semilattices:

f_{000000}	0	1	2	3	f_{000002}	0	1	2	3
0	0	0	0	0	0	0	0	0	0
1	0	1	0	0	1	0	1	0	0
2	0	0	2	0	2	0	0	2	2
3	0	0	0	3	3	0	0	2	3
f_{000012}	0	1	2	3	f_{000111}	0	1	2	3
0	0	0	0	0	0	0	0	0	0
1	0	1	0	1	1	0	1	1	1
2	0	0	2	2	2	0	1	2	1
3	0	1	2	3	3	0	1	1	3
f_{000112}	0	1	2	3					
0	0	0	0	0					
1	0	1	1	1					
2	0	1	2	2					
3	0	1	2	3					

The following eight operations also generate round clones:

f_{000013}	0	1	2	3	f_{000033}	0	1	2	3
0	0	0	0	0	0	0	0	0	0
1	0	1	0	1	1	0	1	0	3
2	0	0	2	3	2	0	0	2	3
3	0	1	3	3	3	0	3	3	3
f_{000123}	0	1	2	3	f_{001031}	0	1	2	3
0	0	0	0	0	0	0	0	0	1
1	0	1	1	2	1	0	1	0	3
2	0	1	2	3	2	0	0	2	1
3	0	2	3	3	3	1	3	1	3
f_{001133}	0	1	2	3	f_{001231}	0	1	2	3
0	0	0	0	1	0	0	0	0	1
1	0	1	1	3	1	0	1	2	3
2	0	1	2	3	2	0	2	2	1
3	1	3	3	3	3	1	3	1	3
f_{001233}	0	1	2	3	f_{011231}	0	1	2	3
0	0	0	0	1	0	0	0	1	1
1	0	1	2	3	1	0	1	2	3
2	0	2	2	3	2	1	2	2	1
3	1	3	3	3	3	1	3	1	3

For each of these operations f , the $(k+1)$ -ary operation

$$\begin{aligned}
f_{k+1}(x_1, \dots, x_{k+1}) := & f_2(f_k(f_k(x_2, x_3, \dots, x_k, x_{k+1}), \\
& f_k(x_1, x_3, \dots, x_k, x_{k+1}), \\
& f_k(x_1, x_2, \dots, x_k, x_{k+1}), \\
& \vdots, \\
& f_k(x_1, x_2, x_3, \dots, x_{k+1})), f_k(x_1, \dots, x_k)),
\end{aligned}$$

where $f_2 = f$, is symmetric; this is similar to the f_{+0-} case from the domain of size 3 enumeration. Lastly, the following three operations generate round clones:

f_{001032}	0	1	2	3	f_{001033}	0	1	2	3
0	0	0	0	1	0	0	0	0	1
1	0	1	0	3	1	0	1	0	3
2	0	0	2	2	2	0	0	2	3
3	1	3	2	3	3	1	3	3	3
f_{001232}	0	1	2	3					
0	0	0	0	1					
1	0	1	2	3					
2	0	2	2	2					
3	1	3	2	3					

For each of these operations, a $(k+1)$ -ary symmetric operation f_{k+1} can be constructed through the following induction, where

$$(y_1, \dots, y_{k+1}) = c_{f_k}^N(x_1, \dots, x_{k+1})$$

for a sufficiently large integer N .

$$\begin{aligned}
f_{k+1}(x_1, \dots, x_{k+1}) := & f_2(f_k(f_k(y_2, y_3, \dots, y_k, y_{k+1}), \\
& f_k(y_1, y_3, \dots, y_k, y_{k+1}), \\
& f_k(y_1, y_2, \dots, y_k, y_{k+1}), \\
& \vdots, \\
& f_k(y_1, y_2, y_3, \dots, y_{k+1})), f_k(y_1, \dots, y_k)).
\end{aligned}$$

This works because (y_1, \dots, y_{k+1}) always lies in $(D')^{k+1}$ for some three-element subset $D' \subsetneq D$ for sufficiently large N :

- For $\langle f_{001032} \rangle$, repeatedly applying c_{f_k} to the tuple (x_1, \dots, x_{k+1}) will always result in an element of $\{0, 1, 3\}^{k+1}$, unless $(x_1, \dots, x_{k+1}) \in \{2, 3\}^{k+1}$.
- For $\langle f_{001033} \rangle$, repeatedly applying c_{f_k} to the tuple (x_1, \dots, x_{k+1}) will always result in an element of $\{0, 1, 3\}^{k+1}$, unless $(x_1, \dots, x_{k+1}) = (2, \dots, 2)$.
- For $\langle f_{001232} \rangle$, repeatedly applying c_{f_k} to the tuple (x_1, \dots, x_{k+1}) will always result in an element of $\{0, 1, 2\}^{k+1}$, unless $(x_1, \dots, x_{k+1}) \in \{1, 3\}^{k+1}$.

The following five operations do not generate clones with a symmetric ternary operation, but they generate round clones when symmetric ternary operations are added:

f_{000132}	0	1	2	3	f_{003012}	0	1	2	3
0	0	0	0	0	0	0	0	0	3
1	0	1	1	3	1	0	1	0	1
2	0	1	2	2	2	0	0	2	2
3	0	3	2	3	3	3	1	2	3

f_{003013}	0	1	2	3	f_{003112}	0	1	2	3
0	0	0	0	3	0	0	0	0	3
1	0	1	0	1	1	0	1	1	1
2	0	0	2	3	2	0	1	2	2
3	3	1	3	3	3	3	1	2	3

f_{003113}	0	1	2	3
0	0	0	0	3
1	0	1	1	1
2	0	1	2	3
3	3	1	3	3

We casework on each one. For the remainder of this paragraph, let f_2 be the binary symmetric operation and let f_3 be the ternary symmetric operation. The induction used to prove that each case yields a round clone is similar to the f_{+-0} case from the domain of size 3 enumeration. To construct the k -ary operation f_{k+1} , define $\mathbf{x} := (x_1, \dots, x_{k+1})$, and define $\mathbf{y} := c_{f_k}^N(x)$, $\mathbf{y}' := c_{f_k}^N(f_2(\mathbf{x}, \mathbf{y}))$, and $\mathbf{y}'' := c_{f_k}^N(f_2(\mathbf{x}, \mathbf{y}'))$ for a sufficiently large integer N . Then

$$f_{k+1}(\mathbf{x}) := f_3(\mathbf{y}, \mathbf{y}', \mathbf{y}'')_1,$$

where the subscript denotes the first element of the tuple, is symmetric; this works for each case because $\mathbf{y} = c_{f_k}^N(\mathbf{x})$ is a constant tuple for sufficiently large N , so either $\mathbf{y} = \mathbf{y}' = \mathbf{y}''$ or $\{\mathbf{y}, \mathbf{y}', \mathbf{y}''\} = \{(d_1, \dots, d_1), (d_2, \dots, d_2), (d_3, \dots, d_3)\}$, where $\{d_1, d_2, d_3\}$ is chosen such that f_2 , when restricted to the domain $\{d_1, d_2, d_3\} \subset D$, can be renamed to f_{+-0} .

The function θ that takes a binary operation f_2 and a ternary operation f_3 as input and outputs a ternary operation is defined as

$$\theta(f_2, f_3)(x, y, z) := c_{f_2}^N(f_2(c_{f_2}^N(x, y, z), f_3(c_{f_2}^N(x, y, z))))_1$$

for a sufficiently large positive integer N . In each of the following cases, θ returns a symmetric ternary operation that modifies only one or two outputs of f_3 . The casework on the binary symmetric operation is below.

- Suppose a clone is generated by f_{000132} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_c(x, y, z) = \begin{cases} 0 & 0 \in \{x, y, z\} \\ 1 & \{x, y, z\} \in \{\{1\}, \{1, 2\}\} \\ 2 & \{x, y, z\} \in \{\{2\}, \{2, 3\}\} \\ 3 & \{x, y, z\} \in \{\{3\}, \{1, 3\}\} \\ c & \{x, y, z\} = \{1, 2, 3\} \end{cases}$$

for some $c \in D$. Since $\theta(f_{000132}, g_c)$ maps

$$g_1 \mapsto g_3 \mapsto g_2 \mapsto g_1,$$

this case gives two distinct minimal round clones.

- Suppose a clone is generated by f_{003012} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_{c,d}(x, y, z) = \begin{cases} 0 & \{x, y, z\} \in \{\{0\}, \{0, 1\}, \{0, 2\}, \{1, 2\}, \{0, 1, 2\}, \\ & \quad \{1, 2, 3\}\} \\ 1 & \{x, y, z\} \in \{\{1\}, \{1, 3\}\} \\ 2 & \{x, y, z\} \in \{\{2\}, \{2, 3\}\} \\ 3 & \{x, y, z\} \in \{\{3\}, \{0, 3\}\} \\ c & \{x, y, z\} = \{0, 1, 3\} \\ d & \{x, y, z\} = \{0, 2, 3\} \end{cases}$$

for some $c, d \in D$. Since $\theta(f_{003012}, g_c)$ maps

$$\begin{aligned} g_{0,0} &\mapsto g_{3,3} \mapsto g_{1,2} \mapsto g_{0,0} \\ g_{0,2} &\mapsto g_{3,0} \mapsto g_{1,3} \mapsto g_{0,2} \\ g_{0,3} &\mapsto g_{3,2} \mapsto g_{1,0} \mapsto g_{0,3} \end{aligned}$$

and maps each of $g_{0,1}, g_{1,1}, g_{2,0}, g_{2,1}, g_{2,2}, g_{2,3}, g_{3,1}$ to one of the above three cycles, this case gives three distinct minimal round clones.

- Suppose a clone is generated by f_{003013} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_c(x, y, z) = \begin{cases} 0 & \{x, y, z\} \in \{\{0\}, \{0, 1\}, \{0, 2\}, \{1, 2\}, \{0, 1, 2\}\} \\ 1 & \{x, y, z\} \in \{\{1\}, \{1, 3\}\} \\ 2 & \{x, y, z\} = \{2\} \\ 3 & \{x, y, z\} \in \{\{3\}, \{0, 3\}, \{2, 3\}, \{0, 2, 3\}\} \\ c & \{x, y, z\} \in \{\{0, 1, 3\}, \{1, 2, 3\}\} \end{cases}$$

for some $c \in D$. Since $\theta(f_{003013}, g_c)$ maps

$$g_0 \mapsto g_3 \mapsto g_1 \mapsto g_0$$

and maps g_2 to the above cycle, this case only gives one minimal round clone.

- Suppose a clone is generated by f_{003112} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_{c,d}(x, y, z) = \begin{cases} 0 & \{x, y, z\} \in \{\{0\}, \{0, 1\}, \{0, 2\}, \{0, 1, 2\}\} \\ 1 & \{x, y, z\} \in \{\{1\}, \{1, 2\}, \{1, 3\}, \{1, 2, 3\}\} \\ 2 & \{x, y, z\} \in \{\{2\}, \{2, 3\}\} \\ 3 & \{x, y, z\} \in \{\{3\}, \{0, 3\}\} \\ c & \{x, y, z\} = \{0, 1, 3\} \\ d & \{x, y, z\} = \{0, 2, 3\} \end{cases}$$

for some $c, d \in D$. Since $\theta(f_{003112}, g_c)$ maps

$$\begin{aligned} g_{0,0} &\mapsto g_{3,3} \mapsto g_{1,2} \mapsto g_{0,0} \\ g_{0,2} &\mapsto g_{3,0} \mapsto g_{1,3} \mapsto g_{0,2} \\ g_{0,3} &\mapsto g_{3,2} \mapsto g_{1,0} \mapsto g_{0,3} \end{aligned}$$

and maps each of $g_{0,1}, g_{1,1}, g_{2,0}, g_{2,1}, g_{2,2}, g_{2,3}, g_{3,1}$ to one of the above three cycles, this case gives three distinct minimal round clones.

- Suppose a clone is generated by f_{003113} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_c(x, y, z) = \begin{cases} 0 & \{x, y, z\} \in \{\{0\}, \{0, 1\}, \{0, 2\}, \{0, 1, 2\}\} \\ 1 & \{x, y, z\} \in \{\{1\}, \{1, 2\}, \{1, 3\}, \{1, 2, 3\}\} \\ 2 & \{x, y, z\} = \{2\} \\ 3 & \{x, y, z\} \in \{\{3\}, \{0, 3\}, \{2, 3\}, \{0, 2, 3\}\} \\ c & \{x, y, z\} = \{0, 1, 3\} \end{cases}$$

for some $c \in D$. Since $\theta(f_{003013}, g_c)$ maps

$$g_0 \mapsto g_3 \mapsto g_1 \mapsto g_0$$

and g_2 to the above cycle, this case gives two distinct minimal round clones.

One can prove that each of the clones enumerated above are distinct by computing relations; for each pair of distinct clones \mathcal{O}_1 and \mathcal{O}_2 in the above list, one can find a relation that is preserved by \mathcal{O}_1 but not by \mathcal{O}_2 .

The following two operations also don't generate clones with a symmetric ternary operation, but they generate round clones when symmetric ternary operations are added:

f_{001132}	0	1	2	3	f_{003312}	0	1	2	3
0	0	0	0	1	0	0	0	0	3
1	0	1	1	3	1	0	1	3	1
2	0	1	2	2	2	0	3	2	2
3	1	3	2	3	3	3	1	2	3

For the remainder of this paragraph, let f_2 be the binary symmetric operation and let f_3 be the ternary symmetric operation. To construct the k -ary operation f_{k+1} for each case, define:

$$\begin{aligned} \mathbf{x} &:= (x_1, \dots, x_{k+1}) \\ \mathbf{y} &:= c_{f_k}^N \left(f_2 \left(c_{f_k}^N(\mathbf{x}), c_{f_k}^{N+1}(\mathbf{x}) \right) \right) \\ \mathbf{y}' &:= c_{f_k}^N \left(f_2 \left(c_{f_k}^N(f_2(\mathbf{x}, \mathbf{y})), c_{f_k}^{N+1}(f_2(\mathbf{x}, \mathbf{y})) \right) \right) \\ \mathbf{y}'' &:= c_{f_k}^N \left(f_2 \left(c_{f_k}^N(f_2(\mathbf{x}, \mathbf{y}')), c_{f_k}^{N+1}(f_2(\mathbf{x}, \mathbf{y}')) \right) \right) \end{aligned}$$

for a sufficiently large integer N . Then

$$f_{k+1}(\mathbf{x}) := f_3(\mathbf{y}, \mathbf{y}', \mathbf{y}'')_1,$$

where the subscript denotes the first element of the tuple, is symmetric; this works for each case because $\mathbf{y} = c_{f_k}^N \left(f_2 \left(c_{f_k}^N(\mathbf{x}), c_{f_k}^{N+1}(\mathbf{x}) \right) \right)$ is a constant tuple for sufficiently large N , so either $\mathbf{y} = \mathbf{y}' = \mathbf{y}''$ or $\{\mathbf{y}, \mathbf{y}', \mathbf{y}''\} = \{(d_1, \dots, d_1), (d_2, \dots, d_2), (d_3, \dots, d_3)\}$, where f_2 restricted to the domain $\{d_1, d_2, d_3\} \subset D$ can be renamed to f_{+-0} .

The function Θ that takes a binary operation f_2 and a ternary operation f_3 as input and outputs a ternary operation is defined as

$$\Theta(f_2, f_3)(x, y, z) := c_{f_2}^N \left(f_2 \left(f_2(c_{f_2}^N(x, y, z), c_{f_2}^{N+1}(x, y, z)), \right. \right. \\ \left. \left. f_3 \left(f_2(c_{f_2}^N(x, y, z), c_{f_2}^{N+1}(x, y, z)) \right) \right) \right)_1$$

for a sufficiently large positive integer N . In each of the following cases, Θ returns a symmetric ternary operation that modifies only one or two outputs of f_3 .

- Suppose a clone is generated by f_{001132} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_c(x, y, z) = \begin{cases} 0 & \{x, y, z\} \in \{\{0, 2\}, \{0, 1, 2\}, \{0, 2, 3\}\} \\ 1 & \{x, y, z\} = \{1, 2\} \\ 2 & \{x, y, z\} \in \{\{2\}, \{2, 3\}\} \\ \sigma^{-1}(\text{sgn}(\sigma(x) + \sigma(y) + \sigma(z))) & \{x, y, z\} \subseteq \{0, 1, 3\} \\ c & \{x, y, z\} = \{1, 2, 3\} \end{cases}$$

for some $c \in D$, where $\sigma(0) = -1$, $\sigma(1) = 0$, and $\sigma(3) = 1$. This case only gives one minimal round clone, since $\Theta(f_{001132}, g_c)$ maps

$$g_1 \mapsto g_3 \mapsto g_2 \mapsto g_1.$$

To prove that $\langle f_{001132}, g_0 \rangle \subseteq \langle f_{001132}, g_1 \rangle$, let $\mathbf{x} := (x, y, z)$ and $\mathbf{y} := f_2(c_{f_{001132}}^N(\mathbf{x}), c_{f_{001132}}^{N+1}(\mathbf{x}))$, which will be constant unless $\{x, y, z\} = \{1, 2, 3\}$. Additionally, let $\mathbf{y}' := c_{f_{001132}}(\mathbf{y})$ and $\mathbf{y}'' := c_{f_{001132}}(\mathbf{y}')$ be the cyclic rotations of \mathbf{y} . Then

$$c_{f_{003312}}^N(f_2(f_2(f_2((g_0(\mathbf{y}), g_0(\mathbf{y})), \mathbf{y}), \mathbf{y}'), \mathbf{y}'', \mathbf{y}''))_1 = g_1(\mathbf{x})$$

where the subscript denotes the first element of the tuple, as desired.

- Suppose a clone is generated by f_{003312} and a symmetric ternary operation. Using Corollary 13, one can force the existence of a symmetric ternary operation of the form

$$g_{c,d}(x, y, z) = \begin{cases} 0 & \{x, y, z\} \in \{\{0\}, \{0, 1\}, \{0, 2\}\} \\ 3 & \{x, y, z\} \in \{\{0, 3\}, \{0, 1, 2\}\} \\ \sigma^{-1}(\text{sgn}(\sigma(x) + \sigma(y) + \sigma(z))) & \{x, y, z\} \subseteq \{1, 2, 3\} \\ c & \{x, y, z\} = \{0, 1, 3\} \\ d & \{x, y, z\} = \{0, 2, 3\} \end{cases}$$

for some $c, d \in D$, where $\sigma(1) = -1$, $\sigma(2) = 1$, and $\sigma(3) = 0$. This case gives three minimal round clones, since $\Theta(f_{003312}, g_{c,d})$ maps

$$g_{0,0} \mapsto g_{3,3} \mapsto g_{1,2} \mapsto g_{0,0}$$

$$g_{0,2} \mapsto g_{3,0} \mapsto g_{1,3} \mapsto g_{0,2}$$

$$g_{0,3} \mapsto g_{3,2} \mapsto g_{1,0} \mapsto g_{0,3}$$

and also eventually maps all others to the above cycles, since $g_{1,1} \mapsto g_{0,1} \mapsto g_{3,0}$, $g_{2,2} \mapsto g_{2,0} \mapsto g_{0,3}$, $g_{2,1} \mapsto g_{1,2}$, $g_{2,3} \mapsto g_{3,2}$, and $g_{3,1} \mapsto g_{1,3}$.

Since we have exhausted all cases, we have established the proof of Theorem 2.

5.5 Domain of Size 5

With computer assistance, it has been shown that every idempotent clone over a domain of size 5 that contains symmetric operations of arities 1, 2, 3, 4, and 5 contains symmetric operations of arities up to 20. The code used to verify this is available on Github at <https://github.com/The-Turtle/PRIMES>.

6 Future Work

If we want to make progress on larger domains, we need a way to determine whether or not a clone has a symmetric operation of a given arity without explicitly generating one.

Definition 16. Let \mathbb{A} be an algebra with underlying set A . For any tuple $\mathbf{a} = (a_1, \dots, a_k) \in A^k$, define the *symmetric relation* on \mathbf{a} to be the set

$$\text{Sym}(\mathbf{a}) := \text{Sg}_{\mathbb{A}^{k!}} \left(\begin{bmatrix} a_{\sigma_1(1)} \\ \vdots \\ a_{\sigma_{k!}(1)} \end{bmatrix}, \dots, \begin{bmatrix} a_{\sigma_1(k)} \\ \vdots \\ a_{\sigma_{k!}(k)} \end{bmatrix} \right),$$

where $\sigma_1, \sigma_2, \dots, \sigma_{k!}$ are the $k!$ permutations of the tuple $(1, \dots, k)$.

Proposition 17. Let \mathbb{A} be an algebra with underlying set A , and suppose that for every $j \leq k$ and every tuple $\mathbf{a} \in A^j$, the symmetric relation $\text{Sym}(\mathbf{a})$ contains a constant tuple. Then \mathbb{A} has a symmetric operation of every arity less than or equal to k .

Proof. We prove this by induction on k ; for the base case $k = 1$, take $f_1 = \pi_1^1$. By the inductive hypothesis, there are symmetric operations f_1, f_2, \dots, f_{k-1} of every arity strictly less than k . Now suppose that f is a k -ary operation such that the set $T \subseteq A^k$ of tuples for which f acts symmetrically on is maximal. We claim that T must equal A^k ; to prove this, it suffices to show that if $t \in A^k \setminus T$ is a tuple which f does not act symmetrically on, then there is a k -ary operation g which acts symmetrically on $T \cup \{t\}$.

We will first construct, for each $j < k$, an operation g_j which acts symmetrically on T and which is unchanged by every permutation of its first j variables. We start by taking $g_1 = f$, and then we inductively define g_j as

$$\begin{aligned} g_j(x_1, \dots, x_k) &:= f_j(g_{j-1}(x_1, x_2, \dots, x_{j-1}, x_j, \quad x_{j+1}, \dots, x_k), \\ &\quad g_{j-1}(x_2, x_3, \dots, x_j, \quad x_1, \quad x_{j+1}, \dots, x_k), \\ &\quad \vdots, \\ &\quad g_{j-1}(x_j, x_1, \dots, x_{j-2}, x_{j-1}, x_{j+1}, \dots, x_k)). \end{aligned}$$

Finally, let \mathbf{a} be the tuple

$$\mathbf{a} := (g_{k-1}(t_1, \dots, t_k), g_{k-1}(t_2, \dots, t_k, t_1), \dots, g_{k-1}(t_k, t_1, \dots, t_{k-1})).$$

By assumption, $\text{Sym}(\mathbf{a})$ contains a constant tuple, so there must be some k -ary operation $h \in \text{Clo}(\mathbb{A})$ which acts symmetrically on \mathbf{a} . Then we define g by

$$g(x_1, \dots, x_k) := h(g_{k-1}(x_1, \dots, x_k), g_{k-1}(x_2, \dots, x_k, x_1), \dots, g_{k-1}(x_k, x_1, \dots, x_{k-1})). \quad \square$$

The relation $\text{Sym}(\mathbf{a})$ has a useful special property.

Proposition 18. Let \mathbb{A} be an algebra with underlying set A . For any tuple $\mathbf{a} \in A^k$ and any pair of permutations (σ, τ) on $\{1, \dots, k\}$, let $\mathbb{P} \leq \mathbb{A}^2$ be the binary relation $\pi_{i_\sigma, i_\tau}(\text{Sym}(\mathbf{a}))$, where i_σ and i_τ are the indices of σ and τ as defined in Definition 16. Then for any subset $B \subseteq A$, we have

$$B + \mathbb{P} = B \implies B - \mathbb{P} = B.$$

Proof. Define

$$\mathbb{P}^{\circ n} := \underbrace{\mathbb{P} \circ \dots \circ \mathbb{P}}_{n \text{ } \mathbb{P}'\text{s}}.$$

Then we have $\mathbb{P}^- \subseteq \mathbb{P}^{\circ(k!-1)}$, since $\mathbb{P}^{\circ(k!-1)}$ contains the generators of \mathbb{P}^- , so $B - \mathbb{P} \subseteq B + \mathbb{P}^{\circ(k!-1)} = B$. Similarly, $B = B + \mathbb{P} \subseteq B - \mathbb{P}^{\circ k!-1} \subseteq B$. Hence $B - \mathbb{P}$ must in fact equal B . \square

Definition 19. Let \mathbb{A} be an algebra with underlying set A . Say that a relation $\mathbb{R} \leq \mathbb{A}^m$ is *reversible* if it satisfies the following two properties:

- for all i, j we have $\pi_i(\mathbb{R}) = \pi_j(\mathbb{R})$, and
- for every sequence $p = ((i_1, j_1), \dots, (i_k, j_k))$ of pairs of coordinates of \mathbb{R} , if we define the binary relation $\mathbb{P}_p \leq \mathbb{A}^2$ by

$$\mathbb{P}_p := \pi_{i_1, j_1}(\mathbb{R}) \circ \dots \circ \pi_{i_k, j_k}(\mathbb{R}),$$

then for every $B \subseteq A$, we have

$$B + \mathbb{P}_p = B \implies B - \mathbb{P}_p = B.$$

Proposition 20. *For every algebra \mathbb{A} with underlying set A and every tuple $\mathbf{a} \in A^n$, the relation $\text{Sym}(\mathbf{a})$ is reversible.*

Proof. Since the marginal distributions of each coordinate of the uniform distribution on the set of tuples in $\text{Sym}(\mathbf{a})$ are equal, this follows from the implication (e) \implies (a) of Proposition 23 below. \square

We have the following strong refinement of our main conjecture.

Conjecture 21. *Suppose that \mathbb{A} is a finite idempotent algebra, such that for every subquotient $\mathbb{B} \in HS(\mathbb{A})$ there is some element $b \in \mathbb{B}$ which is fixed by every automorphism of \mathbb{B} . Then every reversible relation $\mathbb{R} \leq \mathbb{A}^n$ contains a constant tuple.*

The condition involving arbitrary compositions of two-variable projections of the relation \mathbb{R} in the definition of reversibility is necessary, as demonstrated by the following example.

Example 22. Let $\mathbb{A} = (\{-, 0, +\}, \text{sgn}(x + y))$ and let $\mathbb{R} \leq \mathbb{A}^5$ be the relation

$$\mathbb{R} := \{(x_1, x_2, x_3, x_4, x_5) \in \mathbb{A}^5 \mid x_1 + x_2 + x_3 \geq 1 \wedge x_4 = -x_5\}.$$

Then every binary projection of \mathbb{R} is reversible, but the relation \mathbb{R} is *not* reversible: we have

$$\{-\} + \pi_{1,2}(\mathbb{R}) + \pi_{4,5}(\mathbb{R}) = \{-\}$$

but

$$\{-\} - \pi_{4,5}(\mathbb{R}) - \pi_{1,2}(\mathbb{R}) = \{-, 0, +\}.$$

Since \mathbb{R} does not contain any of the constant tuples $(-, \dots, -)$, $(0, \dots, 0)$, or $(+, \dots, +)$, we need the stronger condition about arbitrary compositions of two-variable projections.

For binary relations, the concept of reversibility simplifies.

Proposition 23. *If $\mathbb{R} \leq_{sd} \mathbb{A}^2$ is a binary subdirect relation on a finite algebra \mathbb{A} with underlying set A , then the following are equivalent.*

(a) *For every $B \subseteq A$, we have*

$$B + \mathbb{R} = B \implies B - \mathbb{R} = B.$$

(b) *If we consider the ordered pairs of \mathbb{R} as the edges of a directed graph G with vertex set A , then every weakly connected component of G is also strongly connected.*

(c) *If we consider the ordered pairs of \mathbb{R} as the edges of a directed graph G with vertex set A , then every directed edge of G is contained in a directed cycle of G .*

(d) *There is some $n \geq 1$ such that $\mathbb{R}^- \subseteq \mathbb{R}^{on}$.*

(e) *There is a positive probability distribution with support \mathbb{R} such that the marginal distributions on the first and second coordinates agree.*

(f) *The binary relation \mathbb{R} is reversible; that is, every binary relation which can be written as a composition of copies of \mathbb{R} and \mathbb{R}^- satisfies (a).*

Proof. (a) \implies (b): define a quasiorder \preceq on A by $a \preceq b$ if there is any $k \geq 0$ such that $(a, b) \in \mathbb{R}^{\circ k}$. For any $a \in A$, there is a \preceq -maximal element $b \in A$ such that $a \preceq b$, by the finiteness of A . Let B be the set of all b' such that $b \preceq b'$, then the \preceq -maximality of b implies that B is a strongly connected component of \mathbb{R} and that $B + \mathbb{R} = B$. Then (a) implies that we have $B - \mathbb{R} = B$, so we have

$$a \in \{b\} - \mathbb{R}^{\circ k} \subseteq B - \mathbb{R}^{\circ k} = B,$$

and similarly any element in the weakly connected component containing a is also contained in B .

(b) \implies (c) is obvious. For (c) \implies (d), pick for each directed edge of \mathbb{R} a directed cycle containing it, and choose n such that $n+1$ is a common multiple of the lengths of all of these directed cycles. (d) \implies (a) and (f) \implies (a) are also obvious.

To prove that (c) \implies (e), find a collection \mathcal{C} of directed cycles of \mathbb{R} that contains every edge of \mathbb{R} at least once. Define a probability distribution on \mathbb{R} by the following two step process: first pick a uniformly random cycle $C \in \mathcal{C}$, then pick a uniformly random edge $(x, y) \in C$.

For (e) \implies (a), let $p_{(a,b)} > 0$ be the probability assigned to a given element $(a, b) \in \mathbb{R}$ (and set $p_{(a,b)} = 0$ for $(a, b) \notin \mathbb{R}$), and let

$$p_b := \sum_{a \in A} p_{(a,b)} = \sum_{c \in A} p_{(b,c)}$$

be the marginal probability of seeing b on either the first or second coordinate. For any subset $B \subseteq A$, define $p(B)$ by

$$p(B) := \sum_{b \in B} p_b.$$

Then we have

$$p(B + \mathbb{R}) = \sum_{b \in B + \mathbb{R}} p_b = \sum_{b \in B + \mathbb{R}} \sum_{a \in \{b\} - \mathbb{R}} p_{(a,b)} \geq \sum_{b \in B + \mathbb{R}} \sum_{a \in B} p_{(a,b)} = \sum_{a \in B} p_a = p(B),$$

with equality only when every element $b \in B + \mathbb{R}$ has $\{b\} - \mathbb{R} \subseteq B$. If $B + \mathbb{R} = B$, then we must have equality above, so $B - \mathbb{R} = B + \mathbb{R} - \mathbb{R} = B$.

Given the equivalence between (a) and (e), (e) \implies (f) follows from the fact that if \mathbb{R} and \mathbb{S} are any pair of binary relations such that there are positive probability distributions p and q supported on \mathbb{R} and \mathbb{S} , respectively, such that the marginal of p on the second coordinate equals the marginal of q on the first coordinate, then there is a positive probability distribution “ $p \circ q$ ” supported on $\mathbb{R} \circ \mathbb{S}$ such that the marginals of p and $p \circ q$ on the first coordinate are equal, and the marginals of $p \circ q$ and q on the second coordinate are equal.

The equivalence of (f) can also be shown by proving (d) \implies (f). Let \mathbb{R}' be a composition of i copies of \mathbb{R} and j copies of \mathbb{R}^- , in some order; it suffices to show that \mathbb{R}' satisfies (d). If $i > j$, then $\mathbb{R}' \supseteq \mathbb{R}^{\circ(i-j)}$, so if $\mathbb{R}^- \subseteq \mathbb{R}^{\circ n}$, then

$$(\mathbb{R}^-)^{\circ(i-j)} \subseteq \mathbb{R}^{\circ n(i-j)} \subseteq \mathbb{R}'^{\circ n},$$

and we can finish since \mathbb{R}' and \mathbb{R}'^- are each contained in some composition of $\mathbb{R}^{\circ(i-j)}$ and $(\mathbb{R}^-)^{\circ(i-j)}$. The case $i < j$ is similar, so we are left with the case $i = j$.

To deal with the case $i = j$, the case where \mathbb{R}' is a composition of an equal number of copies of \mathbb{R} and \mathbb{R}^- in some order, code the sequence of copies of \mathbb{R} and \mathbb{R}^- as a sequence of i copies of $+$ and i copies of $-$. Let a and $-b$ be the largest value and smallest value, respectively, of the partial sums of the sequence of $+$'s and $-$'s. Then it's easy to see that \mathbb{R}' contains the relations

$$\mathbb{R}_{\pm a} := \mathbb{R}^{\circ a} \circ (\mathbb{R}^-)^{\circ a}$$

and

$$\mathbb{R}_{\mp b} := (\mathbb{R}^-)^{\circ b} \circ \mathbb{R}^{\circ b}.$$

Thus, both \mathbb{R}' and $(\mathbb{R}')^-$ are contained in some composition of copies of $\mathbb{R}_{\pm a}$ and $\mathbb{R}_{\mp b}$, as desired. \square

Theorem 24. *Conjecture 21 is true for the algebra $\mathbb{A} = (\{-, 0, +\}, \text{sgn}(x + y))$.*

Proof. Let \mathbb{A} have underlying set A , and let $\mathbb{R} \leq \mathbb{A}^n$ be a reversible relation. If $\pi_i(\mathbb{R}) \neq \mathbb{A}$, then $\pi_i(\mathbb{R})$ is a semilattice - we leave this case to the reader. We are left with the case $\pi_i(\mathbb{R}) = A$ for all i ; that is, the case where \mathbb{R} is subdirect.

A brute force enumeration shows that every binary subdirect relation on \mathbb{A} is one of the seven relations

$$\begin{aligned} & \{(x, y) \in A^2 \mid x = y\}, \{(x, y) \in A^2 \mid x = -y\}, \{(x, y) \in A^2 \mid x \leq y\}, \\ & \{(x, y) \in A^2 \mid x \geq y\}, \{(x, y) \in A^2 \mid x + y \geq 0\}, \{(x, y) \in A^2 \mid x + y \leq 0\}, A^2. \end{aligned}$$

In particular, each binary subdirect relation $\mathbb{S} \leq_{sd} \mathbb{A}^2$ is completely determined by the intersection $\mathbb{S} \cap \{-, +\}^2$; in fact, the composition of any pair of binary subdirect relations on \mathbb{A} is also determined by the composition of their restrictions to $\{-, +\}$.

Among these seven relations, the two binary relations $\{(x, y) \in A^2 \mid x \leq y\}$ and $\{(x, y) \in A^2 \mid x \geq y\}$ are *not* reversible. Since

$$\{(x, y) \in A^2 \mid x = -y\} \circ \{(y, z) \in A^2 \mid y + z \geq 0\} = \{(x, z) \in A^2 \mid x \leq z\}$$

and

$$\{(x, y) \in A^2 \mid x + y \leq 0\} \circ \{(y, z) \in A^2 \mid y + z \geq 0\} = \{(x, z) \mid x \leq z\},$$

we see that every reversible subdirect arity- k relation \mathbb{R} either

- (a) has $\pi_{i,j}(\mathbb{R}) \in \{\{(x, y) \in A^2 \mid x = y\}, \{(x, y) \in A^2 \mid x = -y\}, A^2\}$ for all integers $1 \leq i, j \leq k$,
- (b) has $\pi_{i,j}(\mathbb{R}) \in \{\{(x, y) \in A^2 \mid x = y\}, \{(x, y) \in A^2 \mid x + y \geq 0\}, A^2\}$ for all integers $1 \leq i, j \leq k$, or
- (c) has $\pi_{i,j}(\mathbb{R}) \in \{\{(x, y) \in A^2 \mid x = y\}, \{(x, y) \in A^2 \mid x + y \leq 0\}, A^2\}$ for all integers $1 \leq i, j \leq k$.

We will show that in case (a), we have $(0, \dots, 0) \in \mathbb{R}$, in case (b) we have $(+, \dots, +) \in \mathbb{R}$, and in case (c) we have $(-, \dots, -) \in \mathbb{R}$. By symmetry, we only have to consider cases (a) and (b). Case (b) follows from the following claim.

Claim: For any relation $\mathbb{S} \leq \mathbb{A}^n$ such that $(+, +) \in \pi_{i,j}(\mathbb{S})$ for all i, j , we have $(+, \dots, +) \in \mathbb{S}$.

Proof: We will prove, by induction on $|I|$ that for every subset $I \subseteq \{1, 2, \dots, n\}$ there is a tuple $s_I \in \mathbb{S}$ such that its i^{th} coordinate is $+$ for all $i \in I$. The base case $|I| \leq 2$ is our assumption on \mathbb{S} . For the inductive step $|I| \geq 3$, let i, j , and k be any three distinct elements of I . Then we define s_I inductively by

$$s_I := \text{sgn}(s_{I \setminus \{i\}}, s_{I \setminus \{j\}}, s_{I \setminus \{k\}}),$$

using the fact that the three-variable operation $\text{sgn}(x + y + z)$ is in the clone generated by the two-variable operation $\text{sgn}(x + y)$, as proven in the $\langle f_{+0-} \rangle$ case of section 5.3. \blacksquare

Case (a) also follows from the claim. To see this, find a maximal subset $I \subseteq \{1, 2, \dots, n\}$ such that no pair of indices $i, j \in I$ has $\pi_{i,j}(\mathbb{R}) = \{(x, y) \in A^2 \mid x = -y\}$. Then we can use the claim to show that the tuple s given by

$$\pi_i(s) = \begin{cases} + & i \in I \\ - & i \notin I, \end{cases}$$

is in \mathbb{R} . By symmetry, $-s \in \mathbb{R}$ as well. Therefore $\text{sgn}((s) + (-s)) = (0, \dots, 0) \in \mathbb{R}$, so we are done. \square

Using some stronger background theory, we can confirm that Conjecture 21 is true for binary relations.

Theorem 25. *Conjecture 21 holds for binary relations: if every subquotient of a finite idempotent algebra \mathbb{A} has an element fixed by its automorphism group, then every binary reversible relation $\mathbb{R} \leq \mathbb{A}^2$ contains a constant tuple.*

Proof. For idempotent algebras, the assumption implies that \mathbb{A} is Taylor by Proposition 4.14 of [4]; in fact, a more general form of this result is proved in Proposition 2.1 of [8].

Let A and B be the underlying sets of \mathbb{A} and \mathbb{B} , respectively. Assume without loss of generality that \mathbb{R} is subdirect; that is, $\pi_1(\mathbb{R}) = \pi_2(\mathbb{R}) = A$. Let θ be the limit of the linking congruence of the binary relation $\mathbb{R}^{\circ m}$ when m gets large. An alternative way to describe θ is as follows: consider \mathbb{R} to be the edges of a directed graph on \mathbb{A} , and consider two vertices to be equivalent if there is an undirected path connecting them such that the total number of forward edges along the path equals the total number of backward edges along the path. Then \mathbb{R}/θ is the graph of an automorphism of \mathbb{A}/θ , so by assumption there is a congruence class B of θ which is fixed by this automorphism. Since \mathbb{A} is idempotent, \mathbb{B} is a subalgebra of \mathbb{A} , and since B is fixed by this automorphism of \mathbb{A}/θ , $B + \mathbb{R} = B$. The fact that B is a congruence class of θ is equivalent to the restriction of \mathbb{R} to B defining a directed graph of “algebraic length 1,” so we can apply the Loop Lemma of [1] to conclude that \mathbb{R} contains a constant tuple (b, b) with $b \in B$. \square

7 Acknowledgements

We would like to thank the MIT-PRIMES program — including Dr. Tanya Khovanova, Dr. Alexander Vitanov, Dr. Slava Gerovitch, and Prof. Pavel Etingof — for giving us the resources to make this research possible.

References

- [1] Libor Barto and Marcin Kozik. Absorbing subalgebras, cyclic terms, and the constraint satisfaction problem. *Log. Methods Comput. Sci.*, 8(1):1:07, 27, 2012.
- [2] Libor Barto and Marcin Kozik. Robust satisfiability of constraint satisfaction problems. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing, STOC '12*, pages 931–940, New York, NY, USA, 2012. ACM.
- [3] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):Art. 3, 19, 2014.
- [4] Andrei Bulatov and Peter Jeavons. Algebraic structures in combinatorial problems. 2001.
- [5] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.
- [6] Silvia Butti and Victor Dalmau. The complexity of the distributed constraint satisfaction problem. *arXiv preprint arXiv:2007.13594*, 2020.
- [7] Catarina Carvalho and Andrei Krokhin. On algebras with many symmetric operations. *Internat. J. Algebra Comput.*, 26(5):1019–1031, 2016.
- [8] Ralph Freese and Matthew A Valeriote. On the complexity of some maltsev conditions. *International Journal of Algebra and Computation*, 19(01):41–77, 2009.
- [9] Gabor Kun, Ryan O’Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 484–495. ACM, New York, 2012.

A Reversible instances and linear programming

The linear programming relaxation of a CSP is closely related to certain weak local consistency checking procedures. The most basic form of local consistency is known as arc-consistency.

Definition 26. An instance I of a CSP with domain A is called *arc-consistent* if there is a way to associate to each variable x of I a subset $A_x \subseteq A$ such that for every constraint relation R of the instance which involves the variable x , the set of possible values of x which are compatible with the relation R is exactly A_x .

Originally it was believed that a CSP was solved by the basic linear programming relaxation if and only if every arc-consistent instance had a solution – this claim appeared in [9], but the proof was flawed; the CSP defined by the three element algebra $(\{-, 0, +\}, \text{sgn}(x + y))$ is a counterexample.

A stronger form of local consistency was introduced in [3].

Definition 27. An instance I of a CSP is called a *weak Prague instance* if it satisfies the following three conditions.

(P1) The instance I is arc-consistent.

(P2) For every variable x , every set $B \subseteq A_x$, and every cycle p from x to x ,

$$B + p = B \implies B - p = B.$$

(P3) For every variable x , every set $B \subseteq A_x$, and every pair of cycles p, q from x to x ,

$$B + p + q = B \implies B + p = B.$$

An alternative form of condition (P2) is given in [2].

Proposition 28 (Barto, Kozik [2]). *If an instance satisfies condition (P1), then (P2) is equivalent to the following condition.*

(P2*) For all variables x , sets $B \subseteq A_x$, and cycles p from x to x with first step s_1 such that $B + p = B$,

$$B + s_1 - s_1 = B;$$

that is, B is a union of linked components of s_1 .

Conditions (P1) and (P2) are closely related to the basic linear programming relaxation of a CSP, while condition (P3) is closely related to the basic *semidefinite* programming relaxation of a CSP (see [2]).

Theorem 29. *If I is an instance of a CSP such that the basic linear programming relaxation of I has a solution assigning probability vectors p_R to each constraint R of I and probability vectors p_x to each variable x , then the instance I' obtained by restricting each constraint relation of I to the support of the corresponding probability distribution p_R , and similarly for the variable domains, satisfies conditions (P1) and (P2).*

Proof. Assume for simplicity that $I = I'$; that is, all of the probability vectors have full support. The compatibility of the probability vectors p_R with the probability vectors on the variable domains ensures that I is arc-consistent, so (P1) is satisfied. For (P2), it is easier to check condition (P2*) from Proposition 28. For each set $B \subseteq A_x$, we attach a probability $P(B)$ given by

$$P(B) := \sum_{a \in B} p_{x,a}.$$

Now consider any step p_1 from a variable x to an adjacent variable y within a constraint with corresponding relation R . Let $S \subseteq A_x \times A_y$ be the binary projection of the corresponding relation R onto x and y , and let p_S be the corresponding marginal distribution of p_R . Then we have

$$P(B + S) = \sum_{b \in B+S} p_{y,b} \geq \sum_{b \in B+S} \sum_{a \in B} p_{S,(a,b)} = \sum_{a \in B} p_{x,a} = P(B),$$

with equality when $B + S - S = B$. Thus if $B + p = B$, then we have

$$P(B) \leq P(B + p_1) \leq P(B + p) = P(B),$$

so $P(B + p_1) = P(B)$, and thus we have $B + p_1 - p_1 = B$. □

In fact, Theorem 29 has a converse when we restrict our attention to a single cycle at a time. The proof is a straightforward generalization of the implication (a) \implies (e) from Proposition 23.

Theorem 30. *If I is an instance of a CSP such that the associated hypergraph of variables and relations consists of a single cycle, then I has properties (P1) and (P2) if and only if the basic linear relaxation of I has a solution such that for each constraint R of I , the support of the corresponding probability distribution p_R is exactly equal to the relation corresponding to R .*

The connections between the basic linear programming relaxation and conditions (P1) and (P2) make the following conjecture natural.

Conjecture 31. *A CSP defined by relations Γ is solved by its linear programming relaxation if and only if every instance I of the CSP which satisfies conditions (P1) and (P2) has a solution.*

Conjecture 21 is a slight strengthening of a special case of Conjecture 31, where we restrict to the case of CSPs which have just a single variable and a single relation.