Systems of Equations over Finite Semigroups and the #CSP Dichotomy Conjecture

Ondřej Klíma$^{\star 1}$, Benoît Larose$^{\star\star 2}$, and Pascal Tesson$^{\star\star\star 3}$

[1] Department of Mathematics, Masaryk University
klima@math.muni.cz
[2] Department of Mathematics and Statistics, Concordia University
larose@mathstat.concordia.ca
[3] Département d'Informatique et de Génie Logiciel, Université Laval
pascal.tesson@ift.ulaval.ca

Abstract. We study the complexity of counting the number of solutions to a system of equations over a fixed finite semigroup. We show that this problem is always either in FP or #P-complete and describe the borderline precisely. We use these results to convey some intuition about the conjectured dichotomy for the complexity of counting the number of solutions in constraint satisfaction problems.

1 Introduction

Constraint satisfaction problems (or CSPs) are a natural way to formalize a number of computational problems arising from combinatorial optimization, artificial intelligence and database theory. Informally, an instance of CSP consists of a domain, a list of variables and a set of constraints relating the values of the different variables. One then has to decide if the constraints can be simultaneously satisfied. Considerable attention has been given to the case where the constraints are constructed using a finite set of relations Γ and it has been conjectured that for any such Γ the problem CSP(Γ) is either in P or NP-complete [11]. Over the boolean domain Schaefer's classical result [20] states that CSP(Γ) is indeed always in P or NP-complete. More recently, deep results of Bulatov have established a similar dichotomy over the three-element domain [2].

It is similarly believed that the corresponding counting problem #CSP(Γ) is always either tractable (in FP) or #P-complete [3]. This dichotomy is known to hold over the boolean domain [9]. The dichotomy conjectures for CSP and #CSP have been the subject of intense research over the last fifteen years and the algebraic approach uncovered in [14, 13] and extended to the counting problem in [3] has underlied the considerable progress made towards these classification

* Research supported by the Ministry of Education of the Czech Republic.

** Research supported in part by NSERC and CRM.

*** Research supported in part by NSERC and the Alexander von Humboldt Fellowship programme (while the author was at the University of Tübingen).

results. It was shown that the tractability of both CSP(Γ) and #CSP(Γ) depends on the algebraic properties of the set of operations which preserve the relations of Γ. There are now very broad classes of Γs for which the tractability or NP/#P-completeness of CSP(Γ) or #CSP(Γ) can be guaranteed through this algebraic approach. Most notably, Dalmau's recent result on the tractability of CSP(Γ) for any Γ closed under a generalized majority-minority operation provides one of the largest class of tractable CSP and, on the hardness side, any Γ which is not closed under a *Taylor term* is such that CSP(Γ) is NP-complete [8]. This approach also underlies the dichotomy results of Bulatov for CSP over the three-element domain and for the list-homomorphism problem.

Despite this remarkable progress, the resolution of the dichotomy conjecture for CSP still seems a few years away. For #CSP on the other hand, there are good reasons to be more optimistic: the deep results of [3, 7] provide an algebraic criterion for the tractability of #CSP(Γ) which is known to be necessary and is very close to being shown as sufficient [1]. In their seminal paper, Bulatov and Dalmau proved that #CSP(Γ) is #P-complete if Γ is not preserved by any Mal'tsev operation and conjectured that the problem is tractable otherwise [3]. That hypothesis was later refuted by Bulatov and Grohe who completely classified the complexity of #CSP(Γ) when Γ consists of two equivalence relations [7].

In order to illuminate the current status of this conjecture, we build on the work of Nordh and Jonsson [19] and study the problem #EQN$^*_S$ of counting the number of solutions to a system of equations over a fixed finite semigroup S. We show that for any S, the problem #EQN$^*_S$ is either in FP or #P-complete and precisely describe the class of semigroups involved in the result. Such a result is to be expected if one believes that a #CSP dichotomy holds and we show how our classification precisely matches the conjectured criterion for tractability of #CSP. Our results also provide simple examples illustrating the very delicate nature of the dividing line between hard and easy cases of #CSP.

Systems of equations over finite semigroups have already been used as interesting case studies for the complexity of constraint satisfaction problems. In [15] it is shown that for any set of relations Γ, there exists a finite semigroup S_Γ such that CSP(Γ) is polynomial time equivalent to the problem EQN$^*_{S_\Gamma}$ of testing if a system over S_Γ has a solution and so proving a dichotomy for this class of problems is equivalent to proving the CSP dichotomy conjecture. If we consider only the problem of solving systems of equations over finite monoids, then the problem is either in P or NP-complete and this result led to the identification of a new class of tractable CSPs [10]. Nordh also considered the problem of testing if two systems are equivalent or isomorphic [18].

We review in Section 2 the basics of the algebraic approach to CSPs and discuss the current conjectures about the #CSP dichotomy. In Section 3, we give the relevant semigroup theoretic notions and rely on the deep results of [7, 3] to show our dichotomy result. Finally, we consider in Section 4 a more elementary point of view on the same results and thus provide some intuition on the conjectured classification of #CSP. Due to space restrictions, a number of proofs have been moved to the Appendix.

2 Universal Algebra and Constraint Satisfaction Problems

Let D be a finite domain and Γ be a finite set of relations over D. The constraint satisfaction problem over Γ, denoted $\mathrm{CSP}(\Gamma)$ is the following decision problem. The input consists of a list of variables $x_1, \ldots, x_n$ and constraints that are pairs (S_i, R_i) where R_i is a k_i-ary relation in Γ and S_i, the scope of the constraint, is an ordered list of k_i variables. We ask whether there exists an assignment of values in D to the variables such that every constraint is satisfied. The related counting problem $\#\mathrm{CSP}(\Gamma)$ consists of counting the number of such assignments. Throughout the paper, Γ denotes a *constraint language*, i.e. a finite set of relations over some domain D.

The algebraic approach mentioned in our introduction considers the closure properties of Γ. An *operation* f on D is simply a function $f : D^t \to D$. We naturally extend f so that it takes as inputs t k-tuples $\overline{a_1}, \ldots, \overline{a_t}$ of values in D by defining $f(\overline{a_1}, \ldots, \overline{a_t}) = (f(a_{11}, \ldots, a_{t1}), \ldots, f(a_{1k}, \ldots, a_{tk}))$. We say that a k-ary relation R over D is *closed under* f, or that f is a *polymorphism of* R if for any t k-tuples of R, say $\overline{a_1}, \ldots, \overline{a_t}$, we also have $f(\overline{a_1}, \ldots, \overline{a_t}) \in R$. Pictorially,

$$\begin{array}{rlccclr}
 & (& a_{11}, & \ldots, & a_{1k} &) & \in R \\
 & (& \vdots, & \vdots & \vdots &) & \in R \\
 & (& a_{t1}, & \ldots, & a_{tk} &) & \in R \\
\cline{2-6}
\Longrightarrow & (& f(a_{11}, \ldots, a_{t1}), & \ldots, & f(a_{1k}, \ldots, a_{tk}) &) & \in R
\end{array}$$

In other words, if each of the t rows represents a tuple in R then we can apply f on each of the k columns and again obtain a tuple in R.

By extension we say that Γ is closed under f or that f is a polymorphism of Γ if every relation of Γ is closed under f, and denote as $\mathrm{Pol}(\Gamma)$ the set of all such finitary operations f. The fundamental link to the complexity of counting CSPs is the following theorem whose counterpart for the decision problem was proved in [13].

Theorem 1 ([3]). *If Γ_1, Γ_2 are sets of relations over D such that $\mathrm{Pol}(\Gamma_1) \subseteq \mathrm{Pol}(\Gamma_2)$ then $\#\mathrm{CSP}(\Gamma_2)$ is polynomial-time Turing reducible to $\#\mathrm{CSP}(\Gamma_1)$.*

A ternary operation m over D is a *Mal'tsev term* if it satisfies the identities $m(x, y, y) = m(y, y, x) = x$. Bulatov and Dalmau showed that if $\mathrm{Pol}(\Gamma)$ contains a Mal'tsev term then $\mathrm{CSP}(\Gamma)$ is tractable [5]. A very broad criterion for #P-completeness of $\#\mathrm{CSP}(\Gamma)$ can also be given in terms of these operations.

Theorem 2 ([3]). *If Γ is a constraint language such that $\mathrm{Pol}(\Gamma)$ contains no Mal'tsev term, then $\#\mathrm{CSP}(\Gamma)$ is #P-complete.*

In their original conference paper [3], Bulatov and Dalmau conjectured that the presence of a Mal'tsev term in $\mathrm{Pol}(\Gamma)$ was in fact sufficient for the tractability of $\#\mathrm{CSP}(\Gamma)$. That conjecture was disproved by later work of Bulatov and

Grohe [7, 5]: the algorithm that guarantees the tractability of CSP(Γ) when Pol(Γ) contains a Mal'tsev term cannot quite be adapted to solve #CSP(Γ) efficiently. It can be salvaged in one important special case discussed below.

An *algebra* $\mathbb{D}$ over a domain D is a pair $\langle D; F\rangle$ where F is a set of operations over D, called the *fundamental operations* of $\mathbb{D}$. For an algebra $\mathbb{D}$, we denote as Inv($\mathbb{D}$) the set of relations over D which are preserved by all its fundamental operations. Let $\langle \Gamma\rangle$ denote the set of relations[1] Inv(Pol(Γ)). We say that the constraint language Γ is #-tractable (resp. #P-complete) if #CSP(Γ) is in FP (resp. #P-complete). By extension we say that the algebra $\mathbb{D}$ is #-tractable if every finite $\Lambda \subseteq$ Inv($\mathbb{D}$) is #-tractable and say that $\mathbb{D}$ is #P-complete if there exists a finite subset $\Lambda \subseteq$ Inv($\mathbb{D}$) such that Λ is #P-complete. It follows from Theorem 1 that Γ is #-tractable (resp. #P-complete) iff its associated algebra $\langle D; \mathrm{Pol}(\Gamma)\rangle$ is #-tractable (resp. #P-complete).

It will also be convenient to consider standard algebraic constructions: given an algebra $\mathbb{D}$, we fix some indexing of its fundamental operations, and can then consider *subalgebras*, *homomorphic images* and *products* of algebras (see [17] or [5]). Bulatov and Dalmau have shown that if an algebra is #-tractable then so is every finite algebra obtained from it by these constructions; and conversely, if a power or subalgebra or homomorphic image of an algebra $\mathbb{D}$ is #P-complete then so is $\mathbb{D}$. A *congruence* of an algebra is an equivalence relation on its universe which is invariant under the fundamental operations.

Γ is said to be *uniform* if the following holds: for every binary relation $\theta \in \langle\Gamma\rangle$ such that there exists a subset E of D such that θ is an equivalence relation on E, the blocks of θ all have the same size. Equivalently, Γ is uniform if its associated algebra $\mathbb{D}$ is uniform, i.e. if θ is a congruence of a subalgebra of $\mathbb{D}$ then its blocks all have the same size.

Theorem 3 ([3]). *A uniform algebra containing a Mal'tsev term is #-tractable.*

We will give in the next section examples of uniform and non-uniform constraint languages related to systems of equations over Abelian groups. The sufficient condition for tractability provided by this theorem is not necessary. An algebra is #P-complete if it contains no Mal'tsev term and #-tractable if it is uniform and contains a Mal'tsev term but the dividing line between easy and hard cases of #CSP lies in the small gap between these two criteria.

Bulatov and Grohe considered the complexity of the $\#CSP(\Gamma)$ problem for the special case in which Γ consists of two equivalence relations α, β. For any such α, β, we can construct an integer matrix $M_{\alpha,\beta}$ with rows labeled by the α-classes, columns labeled by β-classes and integer entries given by the size of the intersection of the corresponding α and β classes. Although their result is more general, we cite a weaker theorem that is sufficient for our purposes and really represents the core of their arguments.

Theorem 4 ([3]). *If $M_{\alpha\beta}$ is positive and has rank strictly larger than* 1*, then $\#CSP(\alpha, \beta)$ is $\#P$ complete.*

[1] Alternatively, $\langle\Gamma\rangle$ is the set of relations expressible through primitive positive formulas over Γ and the equality relation [5].

Corollary 5. *If Γ is a set of relations and α, β are equivalence relations in $\langle \Gamma \rangle$ with $M_{\alpha,\beta}$ positive of rank strictly larger than* 1*, then* $\#\mathrm{CSP}(\Gamma)$ *is #P-complete.*

It is conjectured that the above corollary provides the frontier between the tractable and #P-complete cases of #CSP. More precisely, $\#\mathrm{CSP}(\Gamma)$ should be tractable if, for every homomorphic image $\mathbb{B}$ of a subalgebra of a finite power of the algebra associated to Γ, and every pair of congruences α and β of $\mathbb{B}$, the matrix $M_{\alpha,\beta}$ has rank 1 if it is positive [1]. Note that by the last result and the remarks preceding Theorem 3 the condition is necessary.

3 Systems of Equations and Dual Algebras

To study the complexity of $\#\mathrm{EQN}^*_S$, we reuse some of the simple but useful observations of [19, 16, 15]. The first concerns the complexity of solving systems over the direct product of two semigroups.

Lemma 6. *Let S and T be finite semigroups such that $\#\mathrm{EQN}^*_T$ is in FP. Then $\#\mathrm{EQN}^*_{S\times T}$, $\#\mathrm{EQN}^*_{S\times S}$ and $\#\mathrm{EQN}^*_S$ are polynomial time Turing equivalent.*

A proof of this simple fact is given in the appendix. Given a system over S, we can introduce for each $s \in S$ a new variable x_s and the equation $x_s = s$ without affecting the number of solutions to the system. Moreover an equation $y_1y_2y_3 = z_1z_2$ can be replaced by the set of equations $y_1y_2 = y'$, $y'y_3 = z'$ and $z_1z_2 = z'$ where y' and z' new dummy variables, again without affecting the number of solutions. We thus assume that our systems consist only of equations of the form $xy = z$, $x = y$ or $x = c$ where x, y, z are variables and $c \in S$ is a constant. Therefore, the problem $\#\mathrm{EQN}^*_S$ can be viewed as a #CSP with domain S and constraint language Γ_S consisting of $|S| + 2$ relations: the $|S|$ singleton unary relations, the equality relation and the ternary relation $\cdot_S = \{(x, y, z) : xy = z\}$. As we explained in the previous section, the complexity $\#\mathrm{EQN}^*_S$ is completely determined by $\mathrm{Pol}(\Gamma_S)$ and we wish to analyze the structure of that set.

We say that an operation $f : S^k \to t$ *commutes* with f if for any $s_1, \ldots, s_k$, $t_1, \ldots, t_k \in S$ it holds that $f(s_1t_1, \ldots, s_kt_k) = f(s_1, \ldots, s_k)f(t_1, \ldots, t_k)$. We further say that f is *idempotent* if $f(x, \ldots, x) = x$. For a semigroup S, we denote as $\mathcal{D}(S)$ the *dual algebra* of S, i.e. the algebra over S containing all operations that commute with $\cdot_S$.

Lemma 7 ([16, 19]). *Let Γ_S be the constraint language defined by equations over the semigroup S, then an operation $f : S^k \to S$ is a polymorphism of Γ_S iff f is idempotent and commutes with f.*

A proof is in the appendix. Combined with Theorems 2 and 3 we get:

Lemma 8. *If $\mathcal{D}(S)$ does not contain a Malt'sev term then $\#\mathrm{EQN}^*_S$ is #P-complete.*

If $\mathcal{D}(S)$ is uniform and contains a Malt'sev term then $\#\mathrm{EQN}^*_S$ is solvable in polynomial time.

As we will see there are semigroups fitting neither of these criteria but, as a first step, we want to identify the classes of semigroups corresponding to these two cases and this requires the introduction of some notions of semigroup theory. Recall that an element e of a semigroup S is *idempotent* if $e^2 = e$: in a finite semigroup, there exists an integer ω (which will have this meaning throughout the paper) such that x^ω is idempotent for all $x \in S$.

We say that S is a *left-zero* (resp. *right-zero*) semigroup if it satisfies $ab = b$ (resp. $ab = a$). In particular, all elements of such semigroups are idempotent. We further say that S is a *rectangular band* if it is the direct product of a right-zero and a left-zero semigroup or, equivalently, if it satisfies $xyz = xz$ and $x^2 = x$.

For a semigroup S, let S^1 denote the monoid obtained from S by adjoining an identity element if no such element exists in S. A semigroup is called *simple* if for any two elements $a, b \in S$, we have $S^1aS^1 = S^1bS^1$. An equivalent requirement is that for any a, b there exist $x, y \in S^1$ such that $xay = b$. In particular, groups and rectangular bands are simple semigroups. It can be easily shown that a semigroup is simple iff for any two idempotents $e, f \in S$ we have $(efe)^\omega = e$.

A semigroup is said to be *orthodox* if the product of two idempotents of S is itself an idempotent and a simple semigroup is orthodox iff it is the direct product of a group and a rectangular band [12].

We will say that a semigroup S is an *inflated simple semigroup* if it consists of a simple subsemigroup T and elements $g_1, \ldots, g_n$ such that for all g_i there exists not necessarily distinct elements $t_1, \ldots, t_n \in T$ satisfying $t_i s = g_i s$ and $st_i = sg_i$ for all $s \in S$. We say that g_i is a *ghost* of t_i. The terminology of course stresses the fact that the actions defined by left and right multiplication of t_i and g_i are indistinguishable. For an element $t \in T$, we denote as $g(t)$ the set of ghosts of t (including t itself) and call $w(t) = |g(t)|$ the *weight* of t in S. We say that S is a *uniform inflation* of T if each $t \in T$ has the same weight.

Lemma 9. *A finite semigroup S is an inflated simple orthodox semigroup with only Abelian subgroups if and only if it satisfies $wxyz = wyxz$ and $xy^\omega z = xz$.*

We prove this lemma in the appendix. In the sequel, we denote as $\mathbf{V}$ the class of semigroups which, as in the statement of the lemma, are the inflation of a direct product of an Abelian group and a rectangular band and for $S \in \mathbf{V}$ we will denote as $c(S)$ the maximal simple subsemigroup of S. In this case, $c(S)$ is always the direct product $A \times L \times R$ of an Abelian group, a left-zero band and a right-zero band. The class $\mathbf{V}$ is tightly connected with dual Malt'sev terms.

Theorem 10.

(a) Let S be a finite semigroup: the dual algebra $\mathcal{D}(S)$ contains a Malt'sev term iff S is in $\mathbf{V}$.

(b) Furthermore, if $S \in \mathbf{V}$ then $\mathcal{D}(S)$ is uniform and contains a Malt'sev term if and only if S is a uniform inflation of $c(S)$.

Proof. Item (a) was established in [19] for simple semigroups in $\mathbf{V}$. For $S \in \mathbf{V}$, we thus know that there exists a Mal'tsev term in the dual algebra of $c(S)$ and our proof, given in the appendix works by extending that operation so S.

The proof of (b) is more technical and is given in the appendix.

□

Theorem 11. *Let S be a semigroup obtained as the inflation of an abelian group A. Let $\mathbb{D}$ denote the dual algebra of S. If S is not a uniform inflation of A then there exist congruences α and β of the algebra $\mathbb{D}^3$ such that the matrix $M_{\alpha\beta}$ is positive of rank strictly greater than 1.*

Corollary 12. *If S is the direct product of a uniformly inflated Abelian group, an inflated left-zero semigroup and an inflated right-zero semigroup then $\#\mathrm{EQN}^*_S$ is tractable. Otherwise, $\#\mathrm{EQN}^*_S$ is #P-complete.*

Proof. We will explicitly prove the upper bound in the next section. The hardness result for non-uniformly inflated Abelian groups is provided by Theorem 11 and results of Section 2. Finally, we will show in the next section that $\#\mathrm{EQN}^*_S$ is #P-complete if S is not the direct product of an inflated Abelian group, an inflated right-zero semigroup and an inflated left-zero semigroup.

□

4 Elementary Arguments for the Complexity of $\#\mathbf{EQN}^*_S$

Suppose that we want to count the number of solutions to a system over an inflated simple semigroup. As noted earlier, we can assume that every equation is of the form $x = y$, $xy = z$ or $z = c$ where c is a constant. We can remove all equations of the form $z = c$ by replacing every occurrence of z by the constant c. In the resulting system, if any solution exists, there exists one in which every variable is set to a value in $c(S)$ because any variable x set to a ghost value s can just as well be set to $s^{\omega+1}$. It is tempting to think that in fact x can be set to any ghost of s but this is not quite the case: if x occurs in an equation of the form $yz = x$ then x can only take values in $c(S)$. We will say that such variables are *regular*. Any solution to the system in which every variable is set to a value in $c(S)$ thus corresponds to a whole set of solutions in which every non-regular variable can be set to any corresponding ghost value. We can formalize these ideas as follows. We say that a solution $\boldsymbol{a} = (a_1, \ldots, a_n)$ is *regular* if all a_i lie in $c(S)$ and define the weight of $\boldsymbol{a}$ as $w(\boldsymbol{a}) = \prod_{x_i \text{ non regular}} w(a_i)$ and it is easy to see that the number of solutions to the system is the sum of all $w(\boldsymbol{a})$ where $\boldsymbol{a}$ is a regular solution.

By Theorem 10 a semigroup S has a dual Mal'tsev term and Γ_S is uniform if and only if S is a product of cyclic groups, right-zero semigroups and left-zero semigroups, the whole of which is uniformly inflated. In this case, tractability of $\#\mathrm{EQN}^*_S$ is guaranteed by Theorem 8. Let us briefly sketch a polynomial-time algorithm in this case. Since $c(S)$ is uniformly inflated, every regular solution $\boldsymbol{a}$ has weight k^t where k is the number of ghosts of any element and t is the number of non-regular variables in the system. It therefore suffices to exhibit a polynomial-time algorithm to count the number of solutions to a system over the direct product of an Abelian group and a rectangular band. By Lemma 6, we can argue separately for the two cases. For completeness, we sketch in the Appendix a proof of the next lemma which can also be obtained through Theorem 3.

Lemma 13. *Let S be an Abelian group. Then $\#\mathrm{EQN}^*_S$ is in FP.*

One can also count in polynomial time the number of solutions to a system of equations over a rectangular band and a stronger result in fact holds.

Lemma 14. *Let S be an inflation of a right-zero band or of a left-zero band. Then $\#\mathrm{EQN}^*_S$ is in FP.*

Proof. We are only interested in summing up the weights of regular solutions. We first identify the regular variables and replace any constant by its representative in $c(S)$ (note that we can do this without harm once the equations of the form $x = c$ have been removed). The resulting system can be viewed as a system over the right-zero semigroup $c(S)$ and we want to understand the structure of the set of solutions. Every equation of the form $xy = z$ is in fact equivalent to $y = z$. Thus if the system has a solution, it is simply defining an equivalence relation on the set of variables and constants. Formally, the system partitions the set of variables and constants into classes $Y_{c_1}, \dots, Y_{c_{|c(S)|}}, X_1, \dots, X_m$ where the constant c_i lies in Y_{c_i}. We have $\boldsymbol{a}$ a solution, iff all variables in Y_{c_i} are set to c_i and all variables in X_i have the same value a_i. We will abuse notation and denote as $|X_i|$ the number of non-regular variables in the set X_i. Now the weight of $\boldsymbol{a}$ is simply

$$w(\boldsymbol{a}) = \prod |g(c_i)|^{|Y_i|} |g(a_i)|^{|X_i|}.$$

The sum of all these weights is thus

$$(\prod |g(c_i)|^{|Y_i|})(\prod_{i=1}^{m} \sum_{s \in c(S)} |g(s)|^{|X_i|}).$$

□

Note that if S is a non-uniform inflation of a right-zero band then, by Theorem 10, the dual algebra $\mathcal{D}(S)$ is non-uniform. Thus, the above lemma provides examples of constraint languages Γ such that $\mathrm{Pol}(\Gamma)$ is non-uniform but $\#\mathrm{CSP}(\Gamma)$ is nonetheless tractable.

Corollary 15. *If S is the direct product of a uniformly inflated Abelian group, an inflated right-zero semigroup and an inflated left-zero semigroup then $\#\mathrm{EQN}^*_S$ is in FP.*

We now want to show that if S is not in this class then $\#\mathrm{EQN}^*_S$ is #P-complete and we provide one explicit proofs of hardness by reducing from the #P-complete problem Permanent. Recall that for an $n \times n$ matrix A, the permanent of A is $Perm(A) = \sum_{\sigma \in S_n} \prod_i a_{i\sigma(i)}$. Valiant proved that the problem of computing the permanent of a matrix over $\mathbb{F}_2$ is #P-complete [21].

Let us start with the simple example of the three-element semigroup C_2' which consists of the two-element group C_2 (with the operation written additively) to which we add a ghost for the element 1. In other words, C_2' has elements $\{0, 1, 1'\}$ and the operation is specified by $0 + 1 = 1 + 0 = 0 + 1' = 1' + 0 = 1$ and $1 + 1' = 1' + 1 = 1 + 1 = 1' + 1' = 0 + 0 = 0$.

Theorem 16. *$\#\mathrm{EQN}_{C_2'}$ is #P-complete under Turing reductions.*

Proof. We first prove that the computation of the permanent reduces to the following problem. Given a system $\mathcal{E}$ of equations over the group C_2 with m variables and an integer $0 \leq i \leq m$, determine the number of solutions to $\mathcal{E}$ that contain i 1's and $m - i$ 0's. We call this problem N_{C_2}.

Let $A = (a_{ij})_{1 \leq i,j \leq n}$ be the matrix over $\mathbb{F}_2$ whose permanent we wish to compute. We construct a system $\mathcal{E}$ of equations over C_2 with n^2 variables y_{ij} for $1 \leq i, j \leq n$. There are $2n$ equations in $\mathcal{E}$ corresponding to the n rows and n columns of A.

Specifically, we have for each i an equation $\sum_{j=1}^{n} a_{ij} y_{ij} = 1$ and for each j an equation $\sum_{i=1}^{n} a_{ij} y_{ij} = 1$. We claim that $Perm(A)$ is exactly the number of solutions of $\mathcal{E}$ that contain n 1's. Indeed, for each permutation $\sigma \in S_n$ such that $\prod_i a_{i\sigma(i)} = 1$, the assignment to the y_{ij}'s that sets $y_{ij} = 1$ if $j = \sigma(i)$ and $y_{ij} = 0$ otherwise is an assignment with n 1's. Moreover it is a solution of the system because $a_{ij} y_{ij} = 1$ iff $j = \sigma(i)$ and in particular the sum of these products over one row or one column is exactly one. Conversely, every solution to $\mathcal{E}$ with n of the y_{ij} set to 1 must be such that there exists a permutation σ with the property that $y_{ij} = 1$ iff $j = \sigma(i)$ for otherwise at least one row or one column has all its y_{ij}'s set to 0 and an equation is left unsatisfied. Furthermore it must be that $a_{i\sigma(i)} = 1$ for otherwise, again, one of the sums $\sum_{j=1}^{n} a_{ij} y_{ij}$ or $\sum_{i=1}^{n} a_{ij} y_{ij}$ is 0.

Hence there is a one-to-one correspondence between solutions of $\mathcal{E}$ with n 1's and permutations $\sigma \in S_n$ such that $\prod_i a_{i\sigma(i)} = 1$ and so $Perm$ reduces to N_{C_2}.

To complete the proof we show a reduction from the problem N_{C_2} to the problem $\#\mathrm{EQN}_{C_2'}$.

Let $\mathcal{E}$ be a system of equations over C_2 and suppose for convenience that the n variables in $\mathcal{E}$ are $x_{11}, \ldots, x_{n1}$. We construct a system $\mathcal{E}'$ of equations over the super-semigroup C_2' from the system $\mathcal{E}$ as follows. For each $2 \leq i \leq n$ we introduce $n - 1$ new variables x_{ij} for $1 \leq j \leq n$ and add the equations $x_{i1} = x_{i2} = \ldots = x_{in}$. Furthermore we replace any occurrence of a variable y by $3y$.

Note that any solution to $\mathcal{E}$ containing i 1's and $(n - i)$ 0's gives rise to 2^{ni} solutions in $\mathcal{E}'$. Indeed if $x_{i1} = 1$ in the solution to $\mathcal{E}$, then x_{i1} can be either 1 or the ghost of 1. The same happens for each copy x_{ij}. Thus if N_i denotes the number of solutions of weight i in $\mathcal{E}$ then the number of solutions in $\mathcal{E}'$ is $\sum 2^{ni} N_i$. Since $N_i \leq \binom{n}{i} \leq 2^n$ we know that the ith block of n bits in the sum $\sum 2^{ni} N_i$ is precisely N_i.

□

This example shows the subtlety of the dividing line between tractable and intractable cases of #CSP. Here, the constraint language $\mathit{\Gamma}_{C_2'}$ is closed under a

Mal'tsev operation but $\Gamma_{C_2'}$ is not uniform. Indeed, one can easily verify that the function $m(x, y, z)$ which is x if $y = z$, z if $x = y$ and $x + y + z$ otherwise is a Mal'tsev polymorphism but the equivalence relation $0 + x = 0 + y$ has two equivalence classes $\{0\}$ and $\{1, 1'\}$ of different sizes. In the extended abstract [3] and the accompanying technical report [4], Bulatov and Dalmau claim that a constraint language over a domain of size three is tractable iff it admits a Mal'tsev polymorphism. Their argument was in fact flawed and the claim was retracted in the full version of the paper [6].

As we mentioned in Section 2, equivalence relations in $\langle\Gamma\rangle$ play a crucial role in the complexity of #CSP(Γ). For a semigroup $S \in \mathbf{V}$, there are a number of very natural equivalence relations defined through equations over S. We know that the simple subsemigroup $c(S)$ can be decomposed as the direct product of an Abelian group A, a right-zero semigroup R and a left-zero semigroup L. Correspondingly, we write an element of this subgroup as (a, r, l). Note that since R and L are right and left-zero, the multiplication in $c(S)$ is given by $(a_1, r_1, l_1)(a_2, r_2, l_2) = (a_1a_2, r_2, l_1)$. Let e denote the element $(1_A, r_0, l_0)$ where 1_A denotes the identity element the group A and r_0, l_0 are arbitrarily chosen elements of respectively R and L. Note that e is idempotent. Consider the binary relations α_A, α_L and α_R defined as

1. $\alpha_A = \{(x, y) : exe = eye\}$;
2. $\alpha_R = \{(x, y) : ex^\omega = ey^\omega\}$;
3. $\alpha_L = \{(x, y) : x^\omega e = y^\omega e\}$;

Clearly, all three are equivalence relations. Furthermore ghosts of a same element are equivalent under all three relations. Thus, in each case, an equivalence class is completely determined by its elements in the simple semigroup $c(S)$. For an element $x = (a, r, l)$ of $c(S)$, we have $x^\omega e = (1_A, r, l)(1_A, r_0, l_0) = (1_A, r, l_0)$ and so $x, y \in c(S)$ are α_R equivalent iff their L coordinate is the same. Similarly, x, y are α_R equivalent if their R coordinate is the same and are α_A equivalent if they agree on their group coordinate. The intersection of these three equivalence relations is the equivalence relation $\{(x, y) : x^{\omega+1} = y^{\omega+1}\}$ which equates two elements which are ghosts of a common element of $c(S)$.

Consider the two equivalence relations $\alpha = \alpha_R$ and $\beta = \alpha_A \cap \alpha_L$ and the corresponding matrix $M_{\alpha\beta}$ (as described before Theorem 4). The entries of this matrix correspond to the cardinality of the intersection of an α and a β class. Each such intersection contains precisely a unique element of $c(S)$ and all its ghosts. If the matrix thus formed has rank greater than 1, we know that $\#\mathrm{EQN}^*_S$ is #P-complete by Corollary 5. Otherwise, a folklore fact about positive integer matrices of rank 1 guarantees that $M_{\alpha,\beta}$ is the product of a row vector ρ and a column vector κ which are both positive integer. This allows us to show:

Lemma 17. *Let S be an inflation of $A \times L \times R$. If $M_{\alpha\beta} = \rho\kappa$ has rank 1, then S is isomorphic to the direct product of an inflation of R and an inflation of $A \times L$.*

Proof. Let $T = A \times L$. Note that two semigroups $S_1, S_2 \in \mathbf{V}$ are isomorphic iff there is an isomorphism ϕ between the simple semigroups $c(S_1)$ and $c(S_2)$ such that for all x, the ghost classes of x and $\phi(x)$ have the same size.

The matrix $M_{\alpha,\beta}$ has dimension $|R| \times |T|$ and we view rows and columns as being labeled with elements r of R and elements t of T respectively. In S, the number of ghosts of the regular element (r, t) is given by the (r, t) entry of the matrix which is $\rho[r]\kappa[t]$.

Consider the inflation R' of R in which the element r has $\rho[r]$ ghosts and similarly let T' be the inflation of T specified by the column vector κ. It is easy to verify that $R' \times T'$ is indeed isomorphic to S since the number of ghosts of the regular element (r, t) in $R' \times T'$ is the product $\rho[r]\kappa[t]$ of the number of ghosts of r in R' and the number of ghosts of t in T'.

□

One can interpret the result as follows: if $M_{\alpha,\beta}$ has rank 1, then we can "peel off" the inflated right-zero band out of S. Since the problem $\#\mathrm{EQN}^*_{R'}$ is tractable for any inflation of a right-zero band (Lemma 14), we get by Lemma 6 that the complexity of $\#\mathrm{EQN}^*_S$ is exactly that of $\#\mathrm{EQN}^*_{T'}$ where T' is the inflation of the product of $L \times A$ given in the previous lemma. We can of course repeat the above argument and now consider over T' the two equivalence relations $\alpha = \alpha_L$ and $\beta = \alpha_A$ and build the matrix $M_{\alpha,\beta}$. This matrix is positive and if its rank is not 1 then $\#\mathrm{EQN}^*_{T'}$ is #P-complete by Theorem 4. Otherwise T' is the direct product of an inflation L' of L and an inflation A' of A. Since $\#\mathrm{EQN}^*_{L'}$ is tractable, the problem $\#\mathrm{EQN}^*_S$ reduces to the problem $\#\mathrm{EQN}^*_{A'}$. This argument completes the proof of Corollary 12.

We argued that $\#\mathrm{EQN}^*_{A'}$ is #P-complete if it is not a uniform inflation of A by using the sophisticated machinery of [5, 7]. An alternative route can also be pursued: it is possible to generalize the argument of Theorem 16 to show that any non-uniform inflation of an Abelian group of prime-power order leads to a #P-complete problem. The proof is a tedious case analysis and involves tricks reminiscent of the *thickening* of [7]. With this result in hand, one can continue to apply the above reasoning to successively peel off from A' uniform inflations of Abelian groups of a given prime power order and thus progressively factor out the tractable components of the $\#\mathrm{EQN}^*_S$ problem. One of three things must happen: in the first case, this process decomposes S as the direct product of semigroups for which the counting problem is tractable and so $\#\mathrm{EQN}^*_S$ is tractable. In the second case, we hit a positive matrix $M_{\alpha,\beta}$ of rank at least 2, in which case the problem is #P-complete by Theorem 4. In the last case we isolate in S a non-uniform inflation of an Abelian group of prime power order and the problem is #P-complete. This approach seems more transparent but, because of the technicality of the arguments, we chose the presentation given in Section 3.

References

1. A. Bulatov. Personnal communication.

2. A. Bulatov. A dichotomy theorem for constraints on a three-element set. In *Proc. of 43rd Foundations of Comp. Sci. (FOCS'02)*, pages 649–658, 2002.
3. A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. In *44th IEEE Symp. on Foundations of Comp¸. Sci. (FOCS'03)*, pages 562–571, 2003.
4. A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. Technical report, PRG-RR-03-13, Oxford University, 2003.
5. A. Bulatov and V. Dalmau. Malt'sev constraints are tractable. 2006. To appear.
6. A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. 2006.
7. A. Bulatov and M. Grohe. The complexity of partition functions. In *Proc. 31st Int. Coll. on Automata Languages and Programming (ICALP'04)*, pages 294–306, 2004.
8. A. Bulatov, A. Krokhin, and P. Jeavons. Constraint satisfaction problems and finite algebras. In *Proceedings 27th International Colloquium on Automata, Languages and Programming—ICALP'00*, volume 1853 of *Lecture Notes in Computer Science*, pages 272–282, 2000.
9. N. Creignou and M. Hermann. Complexity of generalized satisfiability counting problems. *Inf. Comput.*, 125(1):1–12, 1996.
10. V. Dalmau, R. Gavaldà, P. Tesson, and D. Thérien. Tractable clones of polynomials over semigroups. In *Principles and Practice of Constraint Programming—CP'05*, pages 196–210, 2005.
11. T. Feder and M. Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. on Computing*, 28(1):57–104, 1999.
12. J. Howie. *Fundamentals of Semigroup Theory*. Claredon Press, Oxford, 1995.
13. P. Jeavons. On the algebraic structure of combinatorial problems. *Theor. Comput. Sci.*, 200(1-2):185–204, 1998.
14. P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997.
15. O. Klíma, P. Tesson, and D. Thérien. Dichotomies in the complexity of solving systems of equations over finite semigroups. *Theory of Computing Systems*, 2006. To appear.
16. B. Larose and L. Zádori. Taylor terms, constraint satisfaction and the complexity of polynomial equations over finite algebras. *Int. J. on Algebra and Computation*, 2006. To appear, 17 pages.
17. R. McKenzie, G. McNulty, and W. Taylor. *Algebras, Lattices and Varieties.* Wadsworth and Brooks/Cole, 1987.
18. G. Nordh. The complexity of equivalence and isomorphism of systems of equations over finite groups. In *Proc. Math. Found. Comp. Sci. (MFCS'04)*, pages 380–391, 2004.
19. G. Nordh and P. Jonsson. The complexity of counting solutions to systems of equations over finite semigroups. In *Proc. 10th Conf. Computing and Combinatorics (COCOON'04)*, pages 370–379, 2004.
20. T. J. Schaefer. The complexity of satisfiability problems. In *Proc. 10th ACM STOC*, pages 216–226, 1978.
21. L. G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979.

Appendix

We provide proofs for a number of technical results which were omitted due to space constraints.

Lemma 9. *A finite semigroup S is an inflated simple orthodox semigroup that contains only Abelian subgroups if and only if it satisfies $wxyz = wyxz$ and $xy^\omega z = xz$.*

Proof. To show left to right implication, note that it suffices to show that these identities hold for the maximal simple subsemigroup of S because any ghost element can be replaced by its representative in this subsemigroup without affecting the product. Since this simple subsemigroup is orthodox, it is in fact the direct product of an Abelian group with a rectangular band and both terms of the product clearly satisfy the two identities.

For the right to left implication, since S satisfies $wxyz = wyxz$, all its subgroups are Abelian. Consider the subsemigroup of $c(S)$ consisting of elements which can be written as es for some $e, s \in S$ and e idempotent. One can easily verify that $c(S)$ is the maximal simple subsemigroup of S. For any two idempotents $u, v \in c(S)$ we have $uvuv = uv$ by the second identity so $c(S)$ is an orthodox simple subsemigroup. For any element x in $S - c(S)$ we have $x^{\omega+1}$ in $c(S)$ and $x^{\omega+1}z = xz$ as well as $zx^{\omega+1} = zx$ by this same identity. Thus, x is a ghost of the element $x^{\omega+1}$ of the simple subsemigroup.

□

Lemma 6. *Let S and T be finite semigroups such that $\#\mathrm{EQN}^*_T$ is in FP. Then $\#\mathrm{EQN}^*_{S\times T}$, $\#\mathrm{EQN}^*_{S\times S}$ and $\#\mathrm{EQN}^*_S$ are polynomial time Turing equivalent.*

Proof. For a system of equations $\mathcal{E}$, let $\#\mathcal{E}$ denote the number of solutions of $\mathcal{E}$. A system $\mathcal{E}_{S\times T}$ over the direct product $S \times T$ can be viewed as a pair of independent systems $\mathcal{E}_S$ and $\mathcal{E}_T$ over S and T respectively and $\#(\mathcal{E}_{S\times T})$ is simply the product of $\#(\mathcal{E}_S)$ and $\#(\mathcal{E}_T)$. So if $\#(\mathcal{E}_T)$ is computable in polynomial time we have $\#\mathrm{EQN}^*_{S\times T} \leq_T \#\mathrm{EQN}^*_S$ and $\#\mathrm{EQN}^*_{S\times S} \leq_T \#\mathrm{EQN}^*_S$.

Suppose now that $\mathcal{E}_S$ is a system over S. We obtain a system $\mathcal{E}_{S\times S}$ over $S \times S$ by replacing any constant $s \in S$ by (s, s). Clearly $\#(\mathcal{E}_{S\times S}) = \#(\mathcal{E}_S)^2$ so $\#\mathrm{EQN}^*_S \leq_T \#\mathrm{EQN}^*_{S\times S}$. A similar argument shows that if $\#\mathrm{EQN}^*_T$ is in FP then $\#\mathrm{EQN}^*_S \leq_T \#\mathrm{EQN}^*_{S\times T}$.

□

Lemma 7. *Let Γ_S be the constraint language defined by equations over the semigroup S, then an operation $f : S^k \to S$ is a polymorphism of Γ_S iff f is idempotent and commutes with f.*

Proof. We give a brief sketch, a detailed proof appears in [16]. If an operation f preserves all the singleton relations $R_s = \{s\}$ for $s \in S$ then it must be idempotent since $f(s, \ldots, s)$ has to belong to R_s and is thus s. Conversely, any idempotent operation clearly preserves all singleton relations. Similarly, f preserves the relation defined by $xy = z$ iff for any $x_1y_1 = z_1, \ldots, x_ky_k = z_k$ we have $f(x_1, \ldots, x_k)f(y_1, \ldots, y_k) = f(z_1, \ldots, z_k) = f(x_1y_1, \ldots, x_ky_k)$. So f preserves $xy = z$ iff it commutes with the operation of S.

□

Theorem 10.
(a) Let S be a finite semigroup: the dual algebra $\mathcal{D}(S)$ contains a Malt'sev term iff S is in $\mathbf{V}$.
(b) Furthermore, if $S \in \mathbf{V}$ then $\mathcal{D}(S)$ is uniform and contains a Malt'sev term if and only if S is a uniform inflation of $c(S)$.

Proof. **Part (a)**

Let S be an arbitrary semigroup of $\mathbf{V}$. The simple semigroup $c(S)$ is a direct product $A \times R \times L$ of an Abelian group, a right-zero band and a left-zero band. We first show that the dual algebra of $c(S)$ contains a Mal'tsev term [19]. In fact, it is sufficient to show that the duals of A, R and L all contain Mal'tsev terms and obtain a Mal'tsev term for $A \times R \times L$ as the product of these individual functions. For A, let $m_A(x, y, z) = xy^{-1}z$. This operation is clearly Mal'tsev since $m_A(x, y, y) = xy^{-1}y = x$ and similarly $m_A(y, y, x) = x$. Furthermore m_A commutes with a since $m_A(xx', yy', zz') = xx'(yy')^{-1}zz'$ which by commutativity is $xy^{-1}zx'y'^{-1}z' = m_A(x, y, z)m_A(x', y', z')$.

To show that R also has a dual Mal'tsev term, we first arbitrarily choose a group structure over the set R and denote its multiplication as $\circ$. Now, let $m_R(x, y, z) = x \circ y^{-1} \circ z$ where the inverse is taken with respect to $\circ$. Again, the operation is clearly Mal'tsev. To show that it commutes with the right-zero band R, we have $m_R(xx', yy', zz') = m_R(x', y', z')$ since the band is right-zero. Similarly $m_R(x, y, z)m_R(x', y', z') = m_R(x', y', z')$ and so m_R commutes with the band R. The case of L is handled symmetrically.

Let m be the Malt'sev term in the dual algebra of $c(S) = A \times R \times L$ which is defined by applying m_A, m_R, m_L separately to the A, R, L coordinates. We define $m' : S^3 \to S$ by extending m to the larger domain. We set

$$m'(x, y, z) = \begin{cases} x & \text{if } y = z \\ z & \text{if } x = y \\ m(x^{\omega+1}, y^{\omega+1}, z^{\omega+1}) & \text{otherwise.} \end{cases}$$

By definition, m' is a Malt'sev term and coincides with m on elements of $c(S)$ since $x^{\omega+1} = x$ for $x \in c(S)$. Finally, we want to show that $m'(x_1x_2, y_1y_2, z_1z_2) = m'(x_1, y_1, z_1)m'(x_2, y_2, z_2)$. We already know that this holds if all x_i, y_i, z_i lie in $c(S)$ and the key is to observe that for all $s \in S$ we have $m'(x_1, y_1, z_1)s = m'(x_1^{\omega+1}, y_1^{\omega+1}, z_1^{\omega+1}s$. This follows directly if $x_1 \neq y_1$ and $y_1 \neq z_1$. If however

$x_1 = y_1$ then $m'(x_1, y_1, z_1) = z_1$ but since z_1 is a ghost of the element $z_1^{\omega+1}$ we have

$$m'(x_1, y_1, z_1)s = z_1 s = z_1^{\omega+1}s = m'(x_1^{\omega+1}, y_1^{\omega+1}, z_1^{\omega+1})s$$

as required. Thus,

$$\begin{aligned} m'(x_1, y_1, z_1)m'(x_2, y_2, z_2) &= m'(x_1^{\omega+1}, y_1^{\omega+1}, z_1^{\omega+1})m'(x_2^{\omega+1}, y_2^{\omega+1}, z_2^{\omega+1}) \\ &= m'(x_1^{\omega+1}x_2^{\omega+1}, y_1^{\omega+1}y_2^{\omega+1}, z_1^{\omega+1}z_2^{\omega+1}) \\ &= m'(x_1x_2, y_1y_2, z_1z_2). \end{aligned}$$

Conversely, we want to show that no Malt'sev operation commutes with the operation of a semigroup outside of $\mathbf{V}$. This fact is in fact proved implicitly by [19], which show that if S commutes with a Malt'sev operation then it must satisfy the two identities $wxyz = wyxz$ and $xy^\omega z = xz$.

Part (b) For the left to right implication of (b), let $\mathbb{D} = \mathcal{D}(S)$ and let $c(S) = L \times R \times A$. The operations of $\mathbb{D}$ must preserve the relation $\rho = \{(x, y) : x^{\omega+1} = y^{\omega+1}\}$. It is also immediate that ρ is a congruence of $\mathbb{D}$: each of its blocks corresponds to the 'ghosts' of a single element in $L \times R \times A$. So all ghost-classes have the same size.

For the converse, let $T = c(S) = L \times R \times A$. For every x in S, let x' denote its ρ-representative in T, i.e. $x' = x^{\omega+1}$.

Claim 0. The dual algebra $\mathcal{D}(T)$ is uniform.

Proof of Claim 0. We must prove that for any subalgebra X of $\mathcal{D}(T)$, every congruence of X has blocks of equal size. Since we know this result holds for groups, it will suffice to find an Abelian group structure on T such that the operation $M(x, y, z) = x - y + z$ is a term of $\mathcal{D}(T)$. In fact, choose *any* Abelian group structures on L and R and use the group structure of A, and take as your group structure on T the product of these. It is a simple exercise to verify that $M(x, y, z) = x - y + z$ actually commutes with the product on T.

Claim 1. An idempotent operation ϕ is a term of D if and only if it preserves T, and if f denotes its restriction to T then f commutes with the product in T and $\phi(x_1, \ldots, x_n)$ is in the ρ-block of $f(x_1', \ldots, x_n')$ for all $x_i \in D$.

Proof of Claim 1. Clearly every term of D preserves $T = \{x \in D : x^{\omega+1} = x\}$ and is idempotent. Since ϕ commutes with the product of S so does its restriction f. Now

$$\begin{aligned} f(x_1', \ldots, x_n')^{\omega+1} &= \phi(x_1', \ldots, x_n')^{\omega+1} \\ &= \phi(x_1^{\omega+1}, \ldots, x_n^{\omega+1})^{\omega+1} \\ &= \phi(x_1, \ldots, x_n)^{\omega+1} \end{aligned}$$

hence f and ϕ satisfy the desired condition. Conversely

$$\begin{aligned} \phi(x_1y_1, \ldots, x_ny_n) &= \phi((x_1y_1)', \ldots, (x_ny_n)') \\ &= f((x_1y_1)', \ldots, (x_ny_n)') \\ &= f(x_1'y_1', \ldots, x_n'y_n') \\ &= f(x_1', \ldots, x_n')f(y_1', \ldots, y_n') \end{aligned}$$

this last equality holding because f commutes with the product in T; finally notice that this last element must equal

$$\phi(x_1, \dots, x_n)\phi(y_1, \dots, y_n)$$

because T contains only one element in each ρ block and products always lie in T.

By Claims 0 and 1, we may construct a term ϕ of D starting from a term operation f of $\mathcal{D}(T)$ by defining $\phi(x, \dots, x) = x$ for all x and for each $x_i \in S$ not all equal, defining $\phi(x_1, \dots, x_n)$ to be *any* element in the ρ block of $f(x'_1, \dots, x'_n)$.

Claim 2. A subset X of S is a subalgebra of D if and only if there exists a subalgebra Y of $\mathcal{D}(T)$ such that (i) $X = Y$ or (ii) $X = \{x \in S : x^{\omega+1} \in Y\}$.

Proof of Claim 2. If X is a subalgebra of D then $Y = X \cap T$ is a subalgebra of D contained in T, and hence is a subalgebra of $\mathcal{D}(T)$; we must show that every x ρ-equivalent to a y in Y is in X, (unless $Y = X$, which we'll assume is not the case.) Indeed, let $f = id$ on T and define ϕ to be, on each ρ-block, any permutation of the elements of the block (other than the element in T). This shows that if one ghost is in X then all the others in its ρ block are there as well. Now suppose that the ρ block of y_1 is in X, and that y_2 is in Y. Take $f = M$ a Mal'tsev operation on T (as in Claim 0), and define ϕ with it such that $\phi(y_1, z_1, y_2) = z_2$ where z_i is any ghost of y_i. This shows that z_2 is in X and we're done.

By Claim 2, if D has equal sized blocks, then the same holds for any of its subalgebras. One may show, using arguments similar to those in Claim 0, to see that every subalgebra of $\mathcal{D}(T)$ is isomorphic of the form $\mathcal{D}(L' \times R' \times A')$ where L' is a left-zero band, R' is a right-zero band and A' is an Abelian group.

Hence the next claim will be sufficient to prove our result:

Claim 3. An equivalence relation θ is a congruence of D if and only if there exists a congruence α of $\mathcal{D}(T)$ such that (x, y) is in θ if and only if $(x', y') \in \alpha$.

Proof of Claim 3. Let θ be a congruence of D. Obviously the restriction $\theta|_T$ is a congruence of the subalgebra $D(T)$: call it α. Now let u be a ghost of x and let v a ghost of y (where x and y are in T.) We show that $u\theta v$ if and only if $x\alpha y$. Pick ANY element z of S outside T. Let f be the first binary projection $f(s, t) = s$ on T, and extend it to a term ϕ of D such that $\phi(u, z) = x$ and $\phi(v, z) = y$. Then if u is ρ related to v it forces x to be ρ related to y. For the converse the construction is identical, this time choosing $\phi(x, z) = u$ and $\phi(y, z) = v$.

□

Theorem 11. *Let S be a semigroup obtained as the inflation of an Abelian group A. Let $\mathbb{D}$ denote the dual algebra of S. If S is not a uniform inflation of A then there exist congruences α and β of the algebra $\mathbb{D}^3$ such that the matrix $M_{\alpha\beta}$ is positive of rank strictly greater than 1.*

Proof. Let ρ denote the partition of the semigroup S into blocks of ghosts; we know this is a congruence of the algebra $\mathbb{D}$. Recall that $\mathbb{D}$ admits a Malt'sev operation, and hence so does $\mathbb{D}^3$; it follows that any reflexive binary relation on D^3 which is invariant under the terms of $\mathbb{D}^3$ is a congruence. Define the relations α and β as follows: α consists of all pairs $((a,b,c),(a',b',c'))$ such that $a+b+c = a'+b'+c'$ (here the $+$ denotes the operation of S.) This is clearly reflexive and invariant so it is a congruence. Define β as the set of all pairs $((a,b,c),(a',b',c'))$ such that $b \,\rho\, b'$ and $c \,\rho\, c'$. Again it is clear this is a congruence of $\mathbb{D}^3$.

Let $g_1,\dots,g_k$ denote the elements of A and let a_i denote the number of ghosts of g_i.

Consider a block U of α and a block V of β: there exists a unique $g_i \in A$ such that $(a,b,c) \in U$ if and only if $a+b+c = g_i$ and there exists a unique pair $(g_j, g_l) \in A^2$ such that $(a,b,c) \in V$ if and only if $b \,\rho\, g_j$ and $c \,\rho\, g_l$. Hence a triple (a,b,c) lies in $U \cap V$ if and only if $b \,\rho\, g_j$ and $c \,\rho\, g_l$ and $a \,\rho\, (g_i - g_j - g_l)$; hence the number of elements of S in $U \cap V$ is $a_i a_j a_t$ where $g_t = g_i - g_j - g_l$. In particular, the matrix $M_{\alpha\beta}$ is positive.

If the number of ghosts is unbalanced, then there exists a member of A that has a number of ghosts different from the number of ghosts of 0; suppose without loss of generality that $g_1 = 0$ and $a_1 \neq a_2$. Then we consider the rows of the matrix that correspond to the β-block associated to (g_1, g_1) and to (g_1, g_2): these are

$$\begin{array}{cccc} a_1^2 a_1 & a_1^2 a_2 & \dots & a_1^2 a_k \\ a_1 a_2 a_{t_1} & a_1 a_2 a_{t_2} & \dots & a_1 a_2 a_{t_k} \end{array}$$

where $g_{t_i} = g_i - g_1 - g_2 = g_i - g_2$ for all i.

If the matrix had rank 1 then both these rows would be integer multiples of a vector $(c_1,\dots,c_n)$, i.e. we'd have equality of the ratios of all entries, and hence:

$$\frac{a_1}{a_{t_1}} = \frac{a_2}{a_{t_2}} = \dots = \frac{a_k}{a_{t_k}}.$$

Since $g_{t_2} = g_2 - g_2 = 0 = g_1$ we have $a_{t_2} = a_1$, so the common ratio a_2/a_1 is *not* equal to 1. Let a_i be the maximum value of $a_1,\dots,a_k$; then we must have that $a_i/a_{t_i} > 1$; but there exists some j such that $a_{t_j} = a_i$ and thus $a_j/a_{t_j} < 1$, a contradiction.

□

Lemma 13. *Let S be an Abelian group. Then $\#\mathrm{EQN}^*_S$ is in FP.*

Proof. By the remarks preceding the statement of the lemma, it is sufficient to show that $\#\mathrm{EQN}^*_S$ is in FP if S is an Abelian group. Moreover, we can assume that S is a cyclic group of prime power order since any Abelian group is a direct product of such groups.

If S is a cyclic group of prime order p, then $\#\mathrm{EQN}^*_S$ is the problem of counting the number of solutions to a system of linear equations over the field $\mathbb{Z}_p$. We can in polynomial time diagonalize such a system to obtain the dimension d of the solution space and the number of solutions is then p^d.

For $m > 1$ we diagonalize such a system in a similar way. If there is a variable x which has in some equation a coefficient relatively prime to p, then the value of the variable x is fully determined by other variables from this equation. So, we can put this equation aside and eliminate x from the other equations. We continue this process and obtain one of two cases. First, if we reduce the system to the empty system with d free variables then the original system has $(p^m)^d$ solutions. Second, if all coefficients of the system with d variables are divisible by p, then we can divide all equalities by p and obtain a system in the cyclic group of order p^{m-1}. If the resulting system has N solutions (which we can efficiently compute by our induction hypothesis) then the original system has $N \cdot p^d$ solutions.

□