Dichotomies for classes of homomorphism problems involving unary functions

Tomás Feder[*],
268 Waverley St., Palo Alto, CA 94301, U.S.A.

Florent Madelaine[†],
Dépt. math-info BP 7151 15, Université de La Réunion,
av. René Cassin, 97715 Saint Denis, France.

Iain A. Stewart[‡],
Department of Computer Science, University of Durham,
Science Labs, South Road, Durham DH1 3LE, U.K.

August 24, 2003

Abstract

We study non-uniform constraint satisfaction problems where the underlying signature contains constant and function symbols as well as relation symbols. Amongst our results are the following. We establish a dichotomy result for the class of non-uniform constraint satisfaction problems over the signature consisting of one unary function symbol by showing that every such problem is either complete for **L**, via very restricted logical reductions, or trivial (depending upon whether the template function has a fixed point or not). We show that the class of non-uniform constraint satisfaction problems whose templates are structures over the signature λ_2 consisting of two unary function symbols reflects the full computational significance of the class of non-uniform constraint satisfaction problems over relational structures. We prove a dichotomy result for the class of non-uniform constraint satisfaction problems where the template is a λ_2-structure with the property that

[*]e-mail: `tomas@theory.stanford.edu`

[†]e-mail: `florent.madelaine@univ-reunion.fr`. Much of this work was undertaken whilst at the University of Leicester.

[‡]e-mail: `i.a.stewart@durham.ac.uk`. Much of this work was undertaken whilst at the University of Leicester.

the two unary functions involved are the reverse of one another, in that every such problem is either solvable in polynomial-time or **NP**-complete. Finally, we extend some of our results to the situation where instances of non-uniform constraint satisfaction problems come equipped with lists of elements of the template structure which restrict the set of allowable homomorphisms.

1 Introduction

Constraint satisfaction problems consist of finding assignments of values to variables subject to constraints on the values which can be simultaneously assigned to certain specified subsets of variables [21]. They are of great importance in computer science and artificial intelligence, and have strong links with database theory, combinatorics and universal algebra. For example, the general constraint satisfaction problem is nothing other than the conjunctive-query containment problem from database theory and the homomorphism problem from combinatorics [2]; and there is a strong link between the tractability of constraint satisfaction problems and the study of the closure of relations under certain operations in universal algebra [14]. This diversity has meant that the study of these constraint satisfaction problems has progressed on a number of different fronts and according to different motivations.

Our formulation of constraint satisfaction involves the existence of a homomorphism from one finite structure to another, and in this paper we are concerned with the computational complexity of constraint satisfaction problems when the structures involved are restricted. The *general constraint satisfaction problem*, CSP, has: as its instances pairs of finite structures $(\mathcal{A}, \mathcal{B})$ over the same signature; and as its yes-instances instances $(\mathcal{A}, \mathcal{B})$ for which there is a homomorphism from $\mathcal{A}$ to $\mathcal{B}$. The general constraint satisfaction problem is trivially in **NP** (non-deterministic polynomial-time) and is easily shown to be **NP**-complete (via polynomial-time many-one reductions); and it is usual to restrict the problem so that all finite structures come from some specific class of finite structures or, further, so that the second component, the template, of any instance is some fixed finite structure. The former problems are called *uniform* constraint satisfaction problems, as the two structures in an instance can be arbitrarily drawn from the given class of structures, whilst the latter problems are called *non-uniform* constraint satisfaction problems, as the second structure in an instance must be a given fixed structure (rather than thinking of instances of non-uniform constraint satisfaction problems as pairs of finite structures $(\mathcal{A}, \mathcal{T})$, with $\mathcal{T}$

fixed, we simply think of them as finite structures $\mathcal{A}$, with yes-instances those instances $\mathcal{A}$ for which there exists a homomorphism to $\mathcal{T}$). The computational complexity of these restricted problems is then studied with the ultimate goal being a classification as to the conditions under which a (uniform or non-uniform) constraint satisfaction problem has a given computational complexity.

There are two outstanding and well-known results which illustrate this attempt at classification. The first was established by Schaefer [19] who completely classified the complexity of a non-uniform constraint satisfaction problem when the template is a finite structure whose domain consists of two elements, *i.e.*, the template is a Boolean structure. He showed that if the template belongs to one of six specific classes of Boolean structures then the non-uniform constraint satisfaction problem is solvable in polynomial-time, otherwise it is **NP**-complete. Note the dichotomy here: a non-uniform constraint satisfaction problem with a Boolean template is either in **P** (deterministic polynomial-time) or is **NP**-complete (recall that, in general, if $\mathbf{P} \neq \mathbf{NP}$ then there is an infinite collection of distinct classes of polynomial-time equivalent problems between **P** and **NP**: see [10], for example). The second result is due to Hell and Nešetřil [11] who showed that if all structures involved are finite undirected graphs (without self-loops) then the non-uniform constraint satisfaction problem is solvable in polynomial-time if the template is bipartite, otherwise it is **NP**-complete (again, note the dichotomy).

It is with such dichotomy results that we are concerned in this paper. Apart from those established by Hell and Nešetřil and by Schaefer, there are many others, of one flavour or another: see [1, 3, 4, 5, 7, 8, 9, 17, 18] for a selection. Let us discuss some of these results further as they have a direct bearing on our contributions and provide much of our motivation.

Perhaps the focal paper as regards the classification of non-uniform constraint satisfaction problems is [9] for it was in that paper that Feder and Vardi (inspired by the results of Hell and Nešetřil and of Schaefer) posed the question of whether it might be the case that *every* non-uniform constraint satisfaction problem is solvable in polynomial-time or **NP**-complete. In more detail, Feder and Vardi essentially: considered the relationship between the class of non-uniform constraint satisfaction problems and a logic, MMSNP, embodying essential characteristics of constraint satisfaction; considered sub-classes of non-uniform constraint satisfaction problems which reflect the full computational significance of the whole class; and studied requirements for tractability as regards non-uniform constraint satisfaction problems. It is the second of these investigations which is particularly relevant to our studies; and key is Feder and Vardi's result that the problem of deciding whether the class of non-uniform constraint satisfaction problems over relational structures has a dichotomy is equivalent to that of deciding whether the class of non-uniform constraint satisfaction problems involving digraphs has a dichotomy (we later make precise our definitions of classes such as these and of our notion of equivalence: indeed, definitions of all concepts mentioned here are defined in full later on). That is, the (sub-)class of non-uniform constraint satisfaction problems consisting of the class of non-uniform constraint satisfaction problems involving digraphs reflects the full computational significance of the class of non-uniform constraint satisfaction problems over relational structures (in fact, Feder and Vardi also provided other even more restricted classes of non-uniform constraint satisfaction problems involving digraphs which reflect the full computational significance of the class of relational non-uniform constraint satisfaction problems, *e.g.*, the class where the template digraph is necessarily balanced).

Whilst Feder and Vardi's question as to whether the class of non-uniform constraint satisfaction problems exhibits a dichotomy is still unanswered, significant partial positive results have been proven, notably, though not always, for classes where the template is a digraph of a particular type (we refer the reader to Feder and Vardi's equivalence result mentioned above). For example, in [5] it was proven that the class of non-uniform constraint satisfaction problems where the template is a digraph in the form of an oriented cycle has a dichotomy. Furthermore, the search for a classification of non-uniform constraint satisfaction problems has resulted in the study of a number of related concepts. For example, if every instance $\mathcal{A}$ of some non-uniform constraint satisfaction problem comes complete with a list U_u of elements of the template $\mathcal{T}$, for every element $u \in |\mathcal{A}|$, then the problem of deciding whether there is a homomorphism φ from $\mathcal{A}$ to $\mathcal{T}$ with the additional property that for every element $u \in |\mathcal{A}|$, $\varphi(u) \in U_u$, is known as a list homomorphism problem. It has been proven that every list homomorphism problem where the template is a reflexive undirected graph (that is, every vertex has a self-loop) is solvable in polynomial-time, if the template is an interval graph, and **NP**-complete otherwise [6]. If the template is an irreflexive undirected graph (that is, no vertex has a self-loop) then the list homomorphism problem is solvable in polynomial-time, if the complement of the template is a circular arc graph of clique cover number 2, and **NP**-complete otherwise [7]. Finally, if the template is a general undirected graph (that is, a vertex may or may not have a self-loop) then it is still the case that the list homomorphism problem is either solvable in polynomial-time or **NP**-complete [8].

All of the results mentioned above concern non-uniform constraint satisfaction problems over relational structures. Our study in this paper is essentially of non-uniform constraint satisfaction problems where the underlying signature contains constant and function symbols as well as relation symbols; and our results can be summarized as follows. After presenting basic definitions and concepts in Section 2, in Section 3 we consider the class of non-uniform constraint satisfaction problems over a signature consisting of one unary function symbol. We show a dichotomy in that every such problem is either complete for **L** (deterministic log-space), via very restricted logical reductions, or trivial (depending upon whether the template function has a fixed point or not). In Sections 4 and 5, we consider the class of non-uniform constraint satisfaction problems over a signature λ_2 consisting of two unary function symbols, as well as over signatures consisting of a mix of constant, function and relation symbols. Amongst other results, we show that the class of non-uniform constraint satisfaction problems whose templates are λ_2-structures reflects the full computational significance of the class of non-uniform constraint satisfaction problems over relational structures; and we establish similar results when the template is even further restricted. In Section 6, we prove a dichotomy result for the class of non-uniform constraint satisfaction problems where the template is a λ_2-structure with the property that the two unary functions involved are the reverse of one another. Finally, in Section 7 we extend some of our results to the situation where instances of non-uniform constraint satisfaction problems come equipped with lists. In particular: we show that any list homomorphism problem where the template consists of a single unary function is solvable in polynomial-time; and we prove that there is a dichotomy result for the class of list homomorphism problems where the template consists of an arbitrary number of unary functions (the complexity of such a list homomorphism problem is determined by whether the template admits a choice majority function, similarly to as in [6, 7, 8]).

2 Basic definitions

A *signature* consists of a finite collection of constant symbols, function symbols and relation symbols, and each function and relation symbol has an associated arity. A *finite structure* $\mathcal{A}$ over the signature σ, or σ*-structure*, consists of a finite set $|\mathcal{A}|$, the *domain* or *universe*, together with a constant $c^{\mathcal{A}}$ (resp. function $f^{\mathcal{A}}$, relation $R^{\mathcal{A}}$) for every constant symbol c (resp. function symbol f, relation symbol R) of σ, with functions and relations being of the appropriate arity (we often dispense with superscripts in the names of our constants, functions and relations when it is clear as to which structure we are dealing with). The *size* of a structure $\mathcal{A}$ is the size of the domain and is denoted $|\mathcal{A}|$ also. A structure is *relational* if it is over a signature consisting entirely of relation symbols (which we refer to as a *relational signature*).

A *homomorphism* $\varphi : \mathcal{A} \to \mathcal{B}$ from a σ-structure $\mathcal{A}$ to a σ-structure $\mathcal{B}$ is a map $\varphi : |\mathcal{A}| \to |\mathcal{B}|$ such that:

- any constant of $\mathcal{A}$ is mapped to the corresponding constant of $\mathcal{B}$;
- if f is a function symbol of arity a then

$$f^{\mathcal{A}}(u_1, u_2, \ldots, u_a) = v \Rightarrow f^{\mathcal{B}}(\varphi(u_1), \varphi(u_2), \ldots, \varphi(u_a)) = \varphi(v), \quad \text{for all } u_1, u_2, \ldots, u_a, v \in |\mathcal{A}|;$$

- if R is a relation symbol of arity b then

$$R^{\mathcal{A}}(u_1, u_2, \ldots, u_b) \text{ holds } \Rightarrow R^{\mathcal{B}}(\varphi(u_1), \varphi(u_2), \ldots, \varphi(u_b)) \text{ holds}, \quad \text{for all } u_1, u_2, \ldots, u_b \in |\mathcal{A}|.$$

If there exists a homomorphism from $\mathcal{A}$ to $\mathcal{B}$ then we write $\mathcal{A} \to \mathcal{B}$.

Let $\mathbf{C}$ be a class of finite structures. The *uniform constraint satisfaction problem* $\mathrm{CSP}_{\mathbf{C}}$ has: as its instances pairs $(\mathcal{A}, \mathcal{B})$ of structures from $\mathbf{C}$ over the same signature; and as its yes-instances those instances $(\mathcal{A}, \mathcal{B})$ for which there exists a homomorphism from $\mathcal{A}$ to $\mathcal{B}$. The *size* of an instance $(\mathcal{A}, \mathcal{B})$ is $|\mathcal{A}| + |\mathcal{B}|$. If all structures in $\mathbf{C}$ are over the same signature and $\mathcal{T} \in \mathbf{C}$ then the *non-uniform constraint satisfaction problem* $\mathrm{CSP}_{\mathbf{C}}(\mathcal{T})$ has: as its instances structures $\mathcal{A} \in \mathbf{C}$; and as its yes-instances those instances $\mathcal{A}$ for which there exists a homomorphism from $\mathcal{A}$ to $\mathcal{T}$. The *size* of an instance $\mathcal{A}$ is $|\mathcal{A}|$. If $\mathbf{C}$ is the class of all σ-structures and $\mathcal{T}$ is a σ-structure then we abbreviate $\mathrm{CSP}_{\mathbf{C}}(\mathcal{T})$ by $\mathrm{CSP}(\mathcal{T})$. We should add that the individual tractability of an infinite collection of non-uniform constraint satisfaction problems $\{\mathrm{CSP}_{\mathbf{C}}(\mathcal{T}) : \mathcal{T} \in \mathbf{C}\}$ does not automatically yield the tractability of the uniform constraint satisfaction problem $\mathrm{CSP}_{\mathbf{C}}$; for it may be the case that the size of the template, whilst a constant in a non-uniform problem, might play an exponential role in some time bound (see [15] for an examination of this issue).

We shall be involved with problems solvable in $\mathbf{L}$ and complete for this complexity class. As regards completeness, the notion of reduction we work with comes from finite model theory and is the quantifier-free projection.

Before giving a definition of a quantifier-free projection, we present an example of a quantifier-free projection from one problem to another. As it turns out, we will need this actual reduction later on. The reader is referred to, for example, [12, 13, 20] for more on quantifier-free projections and other logical reductions (especially [13]), and their relevance as low-resource reductions: we only sketch the issues here.

Let the signature σ_{2++} consist of the binary relation symbol E and the two constant symbols c and d. We can think of a σ_{2++}-structure as a digraph, possibly with self-loops, with two designated vertices (which may be identical). The problem $\mathrm{DTC}_{0,1}$ has: as its instances the class of σ_{2++}-structures which, when considered as digraphs with self-loops, have the property that every vertex has out-degree *at most* 1; and as its yes-instances those instances with the property that there is a path in the digraph from the vertex c to the vertex d. The problem DTC_1 has: as its instances the class of σ_{2++}-structures which, when considered as digraphs with self-loops, have the property that every vertex has degree *exactly* 1; and as its yes-instances those instances with the property that there is a path in the digraph from the vertex c to the vertex d.

We shall derive four quantifier-free formulae over the signature σ_{2++} and we shall use our formulae to describe, given an instance $\mathcal{A}$ of $\mathrm{DTC}_{0,1}$, an instance $\rho(\mathcal{A})$ of DTC_1: the first formula will define the vertex set of $\rho(\mathcal{A})$; the second formula will describe the edge relation of our instance; and the third and fourth formulae will describe the source and target vertices.

The domain of $\rho(\mathcal{A})$ is $|\mathcal{A}|^2$. We assume that, regardless of the signature, we always have a binary relation *succ* at our disposal that is always interpreted as a *successor relation* on the domain of any structure, *i.e.*, as a relation of the form:

$$\{(i_j, i_{j+1}) : j = 0, 1, \ldots, n-1\},$$

when the domain of a structure of size n is $\{i_0, i_1, \ldots, i_{n-1}\}$, and also two constant symbols, 0 and max, that are always interpreted as the least and greatest elements, respectively, of the successor relation *succ* (more of this successor relation later). Let us suppose for simplicity that the elements of $|\mathcal{A}|$ are $\{0, 1, \ldots, n-1\}$ and abbreviate '$succ(u, w)$' by '$w = u + 1$'. The vertices of $\{(u, w) : w = 0, 1, \ldots, n-1\}$ will form a path $(u, 0), (u, 1), \ldots, (u, n-1)$ in $\rho(\mathcal{A})$, with a self-loop at $(u, n-1)$, except that:

- if (u, w) is an edge of $E^{\mathcal{A}}$, where $u \neq w$, then there is no edge $((u, w), (u, w+1))$ in $\rho(\mathcal{A})$ nor self-loop $((u, n-1), (u, n-1))$, if

$w = n-1$, but there is an edge $((u,w),(w,0))$ in $\rho(\mathcal{A})$; and

- if (u,u) is an edge of $E^{\mathcal{A}}$ then there is no edge $((u,u),(u,u+1))$ in $\rho(\mathcal{A})$ but there is a self-loop $((u,u),(u,u))$.

The source vertex of $\rho(\mathcal{A})$ is the vertex $(c^{\mathcal{A}},0)$ and the target vertex is $(d^{\mathcal{A}},0)$. It is easy to see that an instance $\mathcal{A}$ of $\mathrm{DTC}_{0,1}$ is a yes-instance if, and only if, the instance $\rho(\mathcal{A})$ is a yes-instance of DTC_1 (as whenever $u \neq w$, there is an edge (u,w) in $E^{\mathcal{A}}$ if, and only if, there is a path from vertex $(u,0)$ to vertex $(w,0)$ in $\rho(\mathcal{A})$).

The formulae ψ_0, ψ_E, ψ_c and ψ_d describing the above construction are as follows:

$$\begin{array}{rcl}
\psi_0(x_1,x_2) & \equiv & x_1 = x_1 \\
\psi_E(x_1,x_2,y_1,y_2) & \equiv & (x_1 = y_1 \wedge y_2 = x_2 + 1 \wedge \neg E(x_1,x_2)) \\
 & & \vee(x_1 = y_1 \wedge x_2 = y_2 = max \wedge \neg E(x_1, max)) \\
 & & \vee(x_1 \neq x_2 \wedge y_1 = x_2 \wedge y_2 = 0 \wedge E(x_1,x_2)) \\
 & & \vee(x_1 = x_2 \wedge x_1 = y_1 \wedge x_2 = y_2 \wedge E(x_1,x_2)) \\
\psi_c(x_1,x_2) & \equiv & x_1 = c \wedge x_2 = 0 \\
\psi_d(x_1,x_2) & \equiv & x_1 = d \wedge x_2 = 0
\end{array}$$

The formula $\psi_0(x_1,x_2)$ tells us that the vertex set of $\rho(\mathcal{A})$ is the whole of $|\mathcal{A}|^2$ (it might have restricted the vertex set to be some appropriately defined subset of $|\mathcal{A}|^2$ but in this case didn't); and ψ_E, ψ_c and ψ_d describe the edge relation, the source vertex and the target vertex of $\rho(\mathcal{A})$, respectively.

So, we can say that DTC_1 is a *quantifier-free first-order translation* of $\mathrm{DTC}_{0,1}$ (as the defining formulae are quantifier-free first-order); but we can actually say more. Note that the above formula ψ_E is of the following form:

$$\bigvee\{(\alpha_i \wedge \beta_i) : i = 1,2,\ldots,k\},$$

for some $k \geq 1$, where:

- each α_i is a conjunction of atoms and negated atoms not involving any relation or function symbols of the underlying signature (σ_{2++} in the illustration above);

- the α_i's are *mutually exclusive*, *i.e.*, for any valuation on the variables (and constants) of any α_i and α_j, where $i \neq j$, it is not the case that both α_i and α_j hold;

- each β_i is an atom or a negated atom (over the underlying signature).

Indeed, the formulae ψ_c and ψ_d are trivially of this form too; and, furthermore, ψ_0 is a quantifier-free first-order formula not involving any relation or function symbols of the underlying signature. Hence, there is a *quantifier-free projection* from the problem $\mathrm{DTC}_{0,1}$ to the problem DTC_1. It was proven in [20] that $\mathrm{DTC}_{0,1}$ is complete for $\mathbf{L}$ via quantifier-free projections; and consequently DTC_1 is also complete for $\mathbf{L}$ via quantifier-free projections.

Quantifier-free projections are so called because the defining formulae are quantifier-free first-order and any 'bit' of a target instance, *e.g.*, edge of $\rho(\mathcal{A})$, above, depends only upon at most one 'bit' of the source structure, *e.g.*, edge of $\mathcal{A}$, above. They are extremely restricted reductions between problems and can easily be translated into other restricted circuit-based or model-based reductions, *e.g.*, logtime-uniform NC^1-reductions, used in complexity theory (see [12]). The (built-in) successor relation and the two associated constants give us an ordering of our data which often enables us to model machine-based computations where all data (such as input strings and instantaneous descriptions) is ordered.

We have one final remark: in our example above, we used quantifier-free first-order formulae to describe an edge relation and two constants. We can equally well use such formulae to describe functions by treating an m-ary function f as an $(m+1)$-ary relation R_f where for any elements $u_1, u_2, \ldots, u_m$, there exists exactly one w such that $R_f(u_1, u_2, \ldots, u_m, w)$ holds (constants, *i.e.*, 0-ary functions, are described in this way too).

3 One unary function

In this section we show that the class of non-uniform constraint satisfaction problems with template a unary function has a dichotomy: such a problem is either $\mathbf{L}$-complete (via quantifier-free projections) or trivial.

Let λ_1 be the signature consisting of one unary function symbol f. The decision problem *Hom-Alg*$_1$ is the problem $\mathrm{CSP}_{\mathbf{C}}$, where $\mathbf{C}$ is the class of all λ_1-structures (we introduce the notation *Hom-Alg*, here and in what follows, to accentuate the algebraic aspect of function symbols).

Let $\mathcal{A}$ be a λ_1-structure and let $\sigma_2 = \langle E \rangle$, where E is a binary relation symbol (so, σ_2-structures are digraphs, possibly with self-loops). The *digraph* of $\mathcal{A}$ is the σ_2-structure $\dot{\mathcal{A}} = \langle |\mathcal{A}|, E \rangle$, where $E(u,w)$ holds if, and only if, $f(u) = w$ (note that it may be the case that $E(u,u)$ holds in $\dot{\mathcal{A}}$). The proof of the following lemma is trivial.

Lemma 1 *Let $\mathcal{A}$ and $\mathcal{B}$ be λ_1-structures. Then $\mathcal{A} \to \mathcal{B}$ if, and only if, $\dot{\mathcal{A}} \to \dot{\mathcal{B}}$.* □

Proposition 2 *The problem Hom-Alg*$_1$ *is in* $\mathbf{L}$.

Proof By Lemma 1, we can assume that we are given pairs of digraphs of unary functions as instances rather than pairs of unary functions.

Let $\dot{\mathcal{A}}$ be the digraph of some unary function $\mathcal{A}$. Then, in general, $\dot{\mathcal{A}}$ consists of a collection of connected components where each component is a directed cycle, which may have any length greater than 0 (and so may be a self-loop), some of whose vertices are roots of in-trees (our notion of connectivity in a digraph is with respect to the underlying undirected graph). These components can be visualized as in Fig. 1. We call these components *cycles with pendant in-trees*. We define the *length* of a cycle with pendant in-trees as the length of the directed cycle.

Figure 1. The components of the digraph of a unary function.

Let $(\dot{\mathcal{A}}, \dot{\mathcal{B}})$ be a pair of digraphs of unary functions where $|\dot{\mathcal{A}}| + |\dot{\mathcal{B}}|$ is n. Suppose that there is a homomorphism taking some connected component $\mathcal{A}_0$ of $\dot{\mathcal{A}}$ to a connected component $\mathcal{B}_0$ of $\dot{\mathcal{B}}$. If $\mathcal{A}_0$ is a cycle with pendant in-trees of length γ then $\mathcal{B}_0$ must be a cycle with pendant in-trees of length δ where δ divides γ. Furthermore, if $\mathcal{A}_0$ and $\mathcal{B}_0$ are cycles with pendant in-trees of lengths γ and δ, respectively, and δ divides γ then there is a homomorphism from $\mathcal{A}_0$ to $\mathcal{B}_0$. Hence, the following is a necessary and sufficient condition for a homomorphism from $\dot{\mathcal{A}}$ to $\dot{\mathcal{B}}$ to exist.

- For every cycle with pendant in-trees of length γ in $\dot{\mathcal{A}}$, there must exist a cycle with pendant in-trees of length δ in $\dot{\mathcal{B}}$ where δ divides γ.

This condition can easily be verified using $O(\log n)$ space. For example, we can ascertain whether a vertex u lies on the cycle of a cycle with pendant in-trees in $\dot{\mathcal{A}}$ by walking along the path emanating from u and stopping after n moves (when u doesn't lie on a cycle) or after we have returned to u (when u does lie on a cycle). By counting as we walk, we obtain the length of the cycle (if u lies on a cycle). We can then work through the vertices of $\dot{\mathcal{B}}$ checking to see whether they lie on the cycle of a cycle with pendant in-trees in $\dot{\mathcal{B}}$; and if a vertex does lie on the cycle of a cycle with pendant in-trees then we can check whether the length of this cycle divides c. Hence, the problem *Hom-Alg*$_1 \in \mathbf{L}$. □

Proposition 3 *The problem Hom-Alg*$_1$ *is* $\mathbf{L}$*-hard (via quantifier-free projections).*

Proof Let $\mathcal{A}$ be an instance of DTC_1 (and so is a σ_2-structure). Define the unary function $f_{\mathcal{A}}$ as follows. The domain of $f_{\mathcal{A}}$ is $|\mathcal{A}|^2 \times \{0,1\}$ and:

- if $c = d$ then:
  - $f_{\mathcal{A}}((u,w,b)) = (c,c,0)$, for all $(u,w,b) \in |\mathcal{A}|^2 \times \{0,1\}$;
- if $c \neq d$ then:
  - if $(u,w) \in E$ where $u \neq d$, $w \neq c$ and $u \neq w$ then $f_{\mathcal{A}}((u,u,0)) = (u,w,0)$ and $f_{\mathcal{A}}((u,w,0)) = (w,w,0)$
  - if $(u,u) \in E$ where $u \neq d$ then $f_{\mathcal{A}}((u,u,0)) = (u,u,1)$ and $f_{\mathcal{A}}((u,u,1)) = (u,u,0)$
  - $f_{\mathcal{A}}((d,d,0)) = (c,c,0)$
  - for any element $(u,w,b) \in |\mathcal{A}|^2 \times \{0,1\} \setminus \{(d,c,0)\}$ for which $f_{\mathcal{A}}((u,w,b))$ is still undefined, define $f_{\mathcal{A}}((u,w,b)) = (d,c,0)$, and define $f_{\mathcal{A}}((d,c,0)) = (d,c,1)$.

Essentially, apart from the trivial case where $c = d$, the digraph of $f_{\mathcal{A}}$ is obtained from the digraph whose edge relation is E as follows:

- take a copy of the digraph (with self-loops) whose edge relation is E, and replace any edge emanating from vertex d with the edge (d,c); and
- replace every edge (u,w), apart from the edge (d,c), by a pair of edges $(u,e_{u,w})$ and $(e_{u,w},w)$, where $e_{u,w}$ is a new vertex.

Other vertices are actually introduced in the formal constructive process (defined above), with two of these vertices being $(d,c,0)$ and $(d,c,1)$. The construction is completed by introducing edges from all vertices, apart from $(d,c,0)$, to $(d,c,0)$; and also an edge from $(d,c,0)$ to $(d,c,1)$. Now define the function $g_{\mathcal{A}}$ to have domain $\{0,1\}$ and to be such that $g_{\mathcal{A}}(0) = 1$ and $g_{\mathcal{A}}(1) = 0$. We claim that $\mathcal{A} \in \mathrm{DTC}_1$ if, and only if, $(f_{\mathcal{A}}, g_{\mathcal{A}}) \notin \textit{Hom-Alg}_1$.

The trivial case is straightforward (note that if the digraph of $f_{\mathcal{A}}$ has a self-loop then there is not a homomorphism from $f_{\mathcal{A}}$ to $g_{\mathcal{A}}$): so suppose henceforth that $c \neq d$. Suppose that there is a path in the digraph whose edge relation is E from vertex c to vertex d. Then in the digraph of $f_{\mathcal{A}}$, there is a odd length cycle with pendant in-trees of length greater than 1. Hence, there is no homomorphism from $f_{\mathcal{A}}$ to $g_{\mathcal{A}}$.

Suppose that there is not a path in the digraph whose edge relation is E from vertex c to vertex d. Then all components of the digraph of $f_{\mathcal{A}}$ are even length cycles with pendant in-trees. Hence, there is a homomorphism from $f_{\mathcal{A}}$ to $g_{\mathcal{A}}$.

The construction of the unary functions $f_{\mathcal{A}}$ and $g_{\mathcal{A}}$ from $\mathcal{A}$ can easily be described by quantifier-free projections (see, *e.g.*, [20] for concrete illustrations of logical formulae describing translations between problems) and so the result follows as DTC_1 is **L**-complete via quantifier-free projections. $\square$

The following is now immediate from Propositions 2 and 3.

Theorem 4 *The problem Hom-Alg*$_1$ *is* **L**-*complete* (*via quantifier-free projections*). $\square$

The problem *Hom-Alg*$_1$ is uniform in the sense that any unary function can appear as either the first or second component of an instance. We obtain non-uniform versions of *Hom-Alg*$_1$ by fixing the second component. The problem *Hom-Alg*$_1(\mathcal{T})$, for some λ_1-structure $\mathcal{T}$, is the problem $\mathrm{CSP}(\mathcal{T})$.

The following is immediate from the proofs of Propositions 2 and 3.

Theorem 5 *Let* $\mathcal{T}$ *be the* λ_1-*structure corresponding to the unary function* f *whose domain is* $\{0,1\}$ *and* $f(0) = 1$ *and* $f(1) = 0$. *The problem Hom-Alg*$_1(\mathcal{T})$ *is* **L**-*complete* (*via quantifier-free projections*). $\square$

Hence, not only is the uniform problem *Hom-Alg*$_1$ **L**-complete, there are also non-uniform problems *Hom-Alg*$_1(\mathcal{T})$ that are **L**-complete (moreover, even when $\mathcal{T}$ has only two elements).

Actually, we can say more about non-uniform problems of the form *Hom-Alg*$_1(\mathcal{T})$. Whilst the proof of Proposition 3 is such that the template has a digraph that is a cycle of length 2, we can actually replace this template with any λ_1-structure $\mathcal{T}$ so long as the digraph of $\mathcal{T}$ has a cycle of pendant in-trees of length at least 2 as follows. Suppose that $\dot{\mathcal{T}}$ has cycles of pendant in-trees of lengths $\delta_1, \delta_2, \ldots, \delta_k$, for some $k > 0$. Adopting the terminology of the proof of Proposition 3 and with reference to this proof, in our construction process when we replace an edge of the digraph of $f_{\mathcal{A}}$ with a path of 2 edges, instead we replace the edge with a path of $\delta_1\delta_2\ldots\delta_k$ edges. So, if there is a path in the digraph whose edge relation is E from vertex c to vertex d then the digraph of $f_{\mathcal{A}}$ has a cycle with pendant in-trees of length $\gamma.\delta_1\delta_2\ldots\delta_k+1$, for some $\gamma \geq 1$, and all other cycles with pendant in-trees (if there are any) have lengths divisible by $\delta_1\delta_2\ldots\delta_k$. Also, if there is no such path in the digraph whose edge relation is E then the digraph of $f_{\mathcal{A}}$ is such that every cycle with pendant in-trees has length divisible by $\delta_1\delta_2\ldots\delta_k$. Hence, we obtain the following corollary.

Corollary 6 *Let $\mathcal{T}$ be any λ_1-structure that is without a fixed point. The problem Hom-Alg$_1(\mathcal{T})$ is* **L***-complete (via quantifier-free projections).* □

Trivially, if the λ_1-structure $\mathcal{T}$ has a fixed point then *Hom-Alg*$_1(\mathcal{T})$ consists of every λ_1-structure and is identical to the problem *Hom-Alg*$_1(T_1)$, where f^{T_1} is the function whose domain has one element. Note that whereas the 'trivial' cases of *Hom-Alg*$_1(\mathcal{T})$ are identical to *Hom-Alg*$_1(T_1)$, so there is an analogous remark to be made about Hell and Nešetřil's dichotomy: the 'trivial' cases, here the cases where the problem is solvable in polynomial-time, are identical to the case where the template graph consists of a solitary edge.

4 Two unary functions

We now consider the class of non-uniform constraint satisfaction problems with template two unary functions. Let λ_2 be the signature consisting of the two unary function symbols f and g. We define the decision problem *Hom-Alg*$_2$ to be the problem $\mathrm{CSP}_{\mathbf{C}}$, where $\mathbf{C}$ is the class of pairs of λ_2-structures; and for any λ_2-structure $\mathcal{T}$, we define the problem *Hom-Alg*$_2(\mathcal{T})$ to be the problem $\mathrm{CSP}(\mathcal{T})$. Our main result in this section is that whether this class of problems has a dichotomy is at least as hard as whether the class of non-uniform constraint satisfaction problems with template an arbitrary relational structure has a dichotomy; where by 'dichotomy' we mean that every such problem is either in **P** or complete for **NP** via polynomial-time Turing-reductions (we actually prove the converse result later). We then go on to prove a restricted version of this result.

In future, for brevity, we shall refer to whether the class of non-uniform constraint satisfaction problems with template two unary functions has a dichotomy as the problem $\Delta(\lambda_2)$, and do likewise for other similar problems. For example, whether the class of non-uniform constraint satisfaction problems with template an arbitrary relational structure has a dichotomy is the problem $\Delta(\sigma : \sigma \text{ relational})$; and whether the class of non-uniform constraint satisfaction problems with template a digraph has a dichotomy is the problem $\Delta(\sigma_2)$. We denote the fact that the problem $\Delta(\lambda_2)$ is at least as hard as $\Delta(\sigma: \text{ relational})$, for example, by $\Delta(\sigma: \text{ relational}) \leq \Delta(\lambda_2)$.

Remark 7 From now on, whenever we talk about a dichotomy, we mostly mean whether every problem is either in **P** or is complete for **NP** via polynomial-time Turing-reductions. We reserve the notation '**NP**-complete' to mean complete for **NP** via polynomial-time (many-one) reductions. Most dichotomy results, notably those due to Schaefer and to Hell and Nešetřil mentioned in the Introduction, are such that every problem is in **P** or is **NP**-complete. If a problem is **NP**-complete then it is complete for **NP** via polynomial-time Turing-reductions but it is unknown whether the converse is true. Furthermore, all reductions in Section 5 of [9], where various dichotomy problems over relational signatures are considered, involve polynomial-time reductions; and so the resulting dichotomy problems therein are with respect to **P** and **NP**-completeness via polynomial-time reductions (as opposed to **P** and **NP**-completeness via polynomial-time Turing-reductions as it is for us in this paper).

We shall begin by detailing a transformation of a σ_2-structure $\mathcal{G}$, which we regard as a simple digraph with vertex set $V(\mathcal{G})$ and edge set $E(\mathcal{G})$ (and possibly with self-loops), into a λ_2-structure $\lambda_2(\mathcal{G})$. The λ_2-structure $\lambda_2(\mathcal{G})$ is defined as follows.

- The domain of $\lambda_2(\mathcal{G})$ consists of:

$$\{1,2,3,4\} \cup \{v_u : u \in |\mathcal{G}|\} \cup \{e_{u,w} : u, w \in |\mathcal{G}| \text{ are such that } E^{\mathcal{G}}(u,w) \text{ holds}\}.$$

- $f(1) = 2$, $f(2) = 1$, $f(3) = 2$, $f(4) = 1$, $g(1) = 4$, $g(2) = 3$, $g(3) = 4$ and $g(4) = 3$ in $\lambda_2(\mathcal{G})$.
- If $v_u \in |\lambda_2(\mathcal{G})|$ then $f(v_u) = 3$ and $g(v_u) = 1$ in $\lambda_2(\mathcal{G})$.
- If $e_{u,w} \in |\lambda_2(\mathcal{G})|$ then $f(e_{u,w}) = v_u$ and $g(e_{u,w}) = v_w$ in $\lambda_2(\mathcal{G})$.

The λ_2-structure $\lambda_2(\mathcal{G})$ can be visualized in Fig. 2.

Figure 2. The λ_2-structure $\lambda_2(\mathcal{G})$.

Consider an arbitrary λ_2-structure $\mathcal{F}$. We write the *digraph* (resp. *graph*) $\mathcal{F}$ to denote the digraph (resp. graph) with vertex set $|\mathcal{F}|$ where there is a directed (resp. undirected) edge (x, y) if, and only if, either $f(x) = y$ or $g(x) = y$ in $\mathcal{F}$ (unlike in the previous section, we no longer take the trouble to differentiate between a λ_2-structure and the digraph it describes). The *f-digraph* (resp. *f-graph*) of $\mathcal{F}$ is the subdigraph (resp. subgraph) of the digraph (resp. graph) $\mathcal{F}$ obtained by retaining only the edges resulting from the unary function f (the *g-digraph* and the *g-graph* of $\mathcal{F}$ are defined similarly).

The λ_2-structure $\mathcal{F}$ is *bipartite* if it is bipartite as a graph, and *connected* if it is connected as a graph. A *neighbour* of a vertex x in the digraph $\mathcal{F}$ is any vertex y for which there is an edge (x, y) or (y, x).

Lemma 8 *Let $\mathcal{G}$ and $\mathcal{H}$ be σ_2-structures. Then $\mathcal{G} \to \mathcal{H}$ if, and only if, $\lambda_2(\mathcal{G}) \to \lambda_2(\mathcal{H})$.*

Proof Suppose that $\psi : \mathcal{G} \to \mathcal{H}$. Define the map $\varphi : |\lambda_2(\mathcal{G})| \to |\lambda_2(\mathcal{H})|$ as follows:

- $\varphi(i) = i$, for any $i \in \{1, 2, 3, 4\}$;
- $\varphi(v_u) = v_{\psi(u)}$, for any $v_u \in |\lambda_2(\mathcal{G})|$; and
- $\varphi(e_{u,w}) = e_{\psi(u),\psi(w)}$, for any $e_{u,w} \in |\lambda_2(\mathcal{G})|$.

It is straightforward to check that the map φ is a homomorphism.

Conversely, suppose that $\varphi : \lambda_2(\mathcal{G}) \to \lambda_2(\mathcal{H})$ is a homomorphism. In both $\lambda_2(\mathcal{G})$ and $\lambda_2(\mathcal{H})$, the elements 1 and 2 are the only elements x for which $f(f(x)) = x$; so, in particular, $\varphi(1) \in \{1, 2\}$. If $v_u \in |\lambda_2(\mathcal{G})|$ then $g^{\lambda_2(\mathcal{G})}(v_u) = 1$; and so $\varphi(1) = \varphi(g^{\lambda_2(\mathcal{H})}(v_u)) = g^{\lambda_2(\mathcal{H})}(\varphi(v_u)) \neq 2$, because 2 is not the image of any element of $|\lambda_2(\mathcal{H})|$ under g. Hence, $\varphi(1) = 1$; and consequently $\varphi(2) = 2$, $\varphi(3) = 3$ and $\varphi(4) = 4$.

If $v_u \in |\lambda_2(\mathcal{G})|$ then $g^{\lambda_2(\mathcal{G})}(v_u) = 1$. Hence, $g^{\lambda_2(\mathcal{H})}(\varphi(v_u)) = 1$, with the result that $\varphi(v_u) = v_{u'}$, for some $u' \in |\mathcal{H}|$. If $e_{u,w} \in |\lambda_2(\mathcal{G})|$ then $f^{\lambda_2(\mathcal{G})}(e_{u,w}) = v_u$ and $g^{\lambda_2(\mathcal{G})}(e_{u,w}) = v_w$. Hence, $f^{\lambda_2(\mathcal{H})}(\varphi(e_{u,w})) = \varphi(v_u)$ and $g^{\lambda_2(\mathcal{H})}(\varphi(e_{u,w})) = \varphi(v_w)$, with the result that $\varphi(e_{u,w}) = e_{u',w'}$, for some $u', w' \in |\mathcal{H}|$.

Define the map $\psi : |\mathcal{G}| \to |\mathcal{H}|$ by $\psi(u) = u'$, where $\varphi(v_u) = v_{u'}$, for each $u \in |\mathcal{G}|$. Consider $u, w \in |\mathcal{G}|$ such that $E^{\mathcal{G}}(u, w)$ holds. By above, $\varphi(e_{u,w}) = e_{u',w'}$, for some $e_{u',w'} \in |\lambda_2(\mathcal{H})|$. By above, $f^{\lambda_2(\mathcal{H})}(\varphi(e_{u,w})) = \varphi(v_u) = v_{\psi(u)}$ and $g^{\lambda_2(\mathcal{H})}(\varphi(e_{u,w})) = \varphi(v_w) = v_{\psi(w)}$. Hence, $u' = \psi(u)$ and $w' = \psi(w)$, with $E^{\mathcal{H}}(\psi(u), \psi(w))$ holding. Consequently, ψ is a homomorphism. $\square$

Proposition 9 *Let $\mathcal{H}$ be a σ_2-structure. There is a polynomial-time Turing-reduction from Hom-Alg$_2$($\lambda_2(\mathcal{H})$) to CSP($\mathcal{H}$).*

Proof Let $\mathcal{F}$ be a λ_2-structure. Since $\lambda_2(\mathcal{H})$ is bipartite, there exists no homomorphism from $\mathcal{F}$ to $\lambda_2(\mathcal{H})$ unless $\mathcal{F}$ is also bipartite; so let us assume that $\mathcal{F}$ is bipartite (with some unique bipartition). We may also assume that $\mathcal{F}$ is connected, as otherwise we deal with each connected component in turn.

Colour the elements of $|\mathcal{F}|$ black or white according to the side of the bipartition they lie on. If there exists a homomorphism φ from $\mathcal{F}$ to $\lambda_2(\mathcal{H})$ then there are two possibilities: all the elements coloured black are mapped to 2, 4 or an element corresponding to a vertex of $\mathcal{H}$, and all the elements coloured white are mapped to 1, 3 or an element corresponding to an edge of $\mathcal{H}$; or all the elements coloured black are mapped to 1, 3 or an element corresponding to an edge of $\mathcal{H}$, and all the elements coloured white are mapped to 2, 4 or an element corresponding to a vertex of $\mathcal{H}$. We start with the first possibility (and subsequently handle the second possibility in exactly the same way).

Let $\mathcal{H}_1$ be the digraph, *i.e.*, σ_2-structure, having a single vertex ν and a single edge $\epsilon = (\nu, \nu)$. Since there is a homomorphism from $\mathcal{H}$ to $\mathcal{H}_1$, the proof of Lemma 8 yields that there is a homomorphism τ from $\lambda_2(\mathcal{H})$ to $\lambda_2(\mathcal{H}_1)$, given by:

- $\tau(i) = i$, for $i \in \{1, 2, 3, 4\}$;
- $\tau(v_u) = \nu$, for $u \in |\mathcal{H}|$; and
- $\tau(e_{u,w}) = \epsilon$, for $e_{u,w} \in \{e_{u,w} : u, w \in |\mathcal{H}| \text{ and } E^{\mathcal{H}}(u, w) \text{ holds}\}$.

Suppose that φ is a homomorphism from $\mathcal{F}$ to $\lambda_2(\mathcal{H})$. The homomorphism $\tau\varphi$ enables us to label every element of $\mathcal{F}$, in the natural way, with an element of $\{1, 2, 3, 4, \nu, \epsilon\}$.

Our strategy in developing our algorithm (to satisfy the statement of the theorem) is essentially to initially label every element x of $|\mathcal{F}|$ with a set of elements of $\{1, 2, 3, 4, \nu, \epsilon\}$ so that each element might constitute a legitimate value for $\tau\varphi(x)$, via any resulting homomorphism φ, and then to reason so that we can reduce these label sets and:

- either obtain an element with an empty label set, in which case a homomorphism φ cannot exist; or
- obtain a collection of digraphs whereby answers as to whether these digraphs are in $\mathrm{CSP}(\mathcal{H})$ will enable us to ascertain whether there exists a homomorphism φ.

Consider the f-digraph of $\mathcal{F}$. This is of the form of a disjoint union of directed cycles with pendant in-trees. For each element x in this f-digraph, define $\delta_f(x)$ to be the length of the longest walk whose end-point is the element x (note that we allow elements to be repeated on a walk). So, in particular, if x is an element in a directed cycle then $\delta_f(x) = \infty$. For each element x of our f-digraph, consider the possible label set $\Lambda(x)$ of x. There are a number of possibilities.

(1) If $\delta_f(x) > 1$ and x is coloured black then $\Lambda(x) \subseteq \{2\}$.

(2) If $\delta_f(x) = 1$ and x is coloured black then $\Lambda(x) \subseteq \{2, \nu\}$.

(3) If $\delta_f(x) = 0$ and x is coloured black then $\Lambda(x) \subseteq \{2, 4, \nu\}$.

(4) If $\delta_f(x) > 2$ and x is coloured white then $\Lambda(x) \subseteq \{1\}$.

(5) If $\delta_f(x) \in \{1, 2\}$ and x is coloured white then $\Lambda(x) \subseteq \{1, 3\}$.

(6) If $\delta_f(x) = 0$ and x is coloured white then $\Lambda(x) \subseteq \{1, 3, \epsilon\}$.

Suppose that we are in Case (3). There are two further cases: either there exists some y such that $g^{\mathcal{F}}(y) = x$ or there doesn't. In the first case, the value 2 cannot lie in $\Lambda(x)$; and in the second case, we may as well remove the value 4 from $\Lambda(x)$, as we could amend any corresponding homomorphism φ so that $\varphi(x) = 2$.

Suppose that we are in Case (6). There are two further cases: either there exists some y such that $g^{\mathcal{F}}(y) = x$ or there doesn't. In the first case, the value ϵ cannot lie in $\Lambda(x)$; and in the second case, we may as well remove the value 3 from $\Lambda(x)$, as we could amend any corresponding homomorphism φ so that $\varphi(x) = 1$.

So, the result is that every element x of $|\mathcal{F}|$ can be labelled with one of the following sets:

$$\{1\}, \{2\}, \{2, \nu\}, \{4, \nu\}, \{1, 3\}, \{1, \epsilon\}.$$

We can consider some connected g-digraph within $\mathcal{F}$ in exactly the same way. Defining δ_g with respect to this connected g-digraph rather than a connected f-digraph yields a number of possibilities.

(1) If $\delta_g(x) > 1$ and v is coloured black then $\Lambda(x) \subseteq \{4\}$.

(2) If $\delta_g(x) = 1$ and x is coloured black then $\Lambda(x) \subseteq \{4, \nu\}$.

(3) If $\delta_g(x) = 0$ and x is coloured black then $\Lambda(x) \subseteq \{2, 4, \nu\}$.

(4) If $\delta_g(x) > 2$ and x is coloured white then $\Lambda(x) \subseteq \{3\}$.

(5) If $\delta_g(x) \in \{1, 2\}$ and x is coloured white then $\Lambda(x) \subseteq \{1, 3\}$.

(6) If $\delta_g(x) = 0$ and x is coloured white then $\Lambda(x) \subseteq \{1, 3, \epsilon\}$.

Suppose that we are in Case (3). There are two further cases: either there exists some y such that $f^{\mathcal{F}}(y) = x$ or there doesn't. In the first case, the value 4 cannot lie in $\Lambda(x)$; and in the second case, we may as well remove the value 2 from $\Lambda(x)$, as we could amend any corresponding homomorphism φ so that $\varphi(x) = 4$.

Suppose that we are in Case (6). There are two further cases: either there exists some y such that $f^{\mathcal{F}}(y) = x$ or there doesn't. In the first case, the value ϵ cannot lie in $\Lambda(x)$; and in the second case, we may as well remove the value 3 from $\Lambda(x)$, as we could amend any corresponding homomorphism φ so that $\varphi(x) = 1$.

So, the result is that every element x of $|\mathcal{F}|$ can be labelled with one of the following sets:

$$\{3\}, \{4\}, \{2, \nu\}, \{4, \nu\}, \{1, 3\}, \{1, \epsilon\}.$$

Note that evaluating the two analyses above might result in an element labelled by the empty set. If this is so then there can exist no homomorphism φ (from $\mathcal{F}$ to $\lambda_2(\mathcal{H})$ where all the elements coloured black are mapped to 2, 4 or an element corresponding to a vertex of $\mathcal{H}$, and all the elements coloured white are mapped to 1, 3 or an element corresponding to a edge of $\mathcal{H}$). Hence, we may assume that every element of $\mathcal{F}$ is labelled with one of the following sets:

$$\{1\}, \{2\}, \{3\}, \{4\}, \{\nu\}, \{\epsilon\}, \{2, \nu\}, \{4, \nu\}, \{1, 3\}, \{1, \epsilon\}$$

(note that our evaluation of the two analyses might enable us to deduce that $\Lambda(x) = \{\nu\}$ or $\{\epsilon\}$).

Recall our λ_2-structure $\lambda_2(\mathcal{H}_1)$ (with domain $\{1, 2, 3, 4, \nu, \epsilon\}$). Abbreviate f and g in $\lambda_2(\mathcal{H}_1)$ by f_1 and g_1, respectively; and for any set $S \subseteq \{1, 2, 3, 4, \nu, \epsilon\}$, define $f_1(S) = \{f_1(i) : i \in S\}$ and $g_1(S) = \{g_1(i) : i \in S\}$.

Suppose that $x, y \in |\mathcal{F}|$ (where possibly $x = y$) and $f^{\mathcal{F}}(x) = y$. Then any label set of y must be contained in $f_1(\Lambda(x))$. Hence, we can replace $\Lambda(y)$ with $\Lambda(y) \cap f_1(\Lambda(x))$. The same can be said if $x, y \in |\mathcal{F}|$ and $g^{\mathcal{F}}(x) = y$. We now repeatedly perform the above modification until either we obtain stability or an element with an empty label set. If we obtain an element with an empty label set then there can be no homomorphism φ: otherwise, there is still a chance that such a homomorphism φ exists.

We can now delete any element x of $|\mathcal{F}|$ that is labelled with the set $\{1\}$ (resp. $\{2\}$, $\{3\}$, $\{4\}$) as if the homomorphism φ exists then: the value of

φ at x is necessarily 1 (resp. 2, 3, 4); and this will be consistent with any resulting value of φ at any neighbouring element in the digraph $\mathcal{F}$. Hence, let us assume: that every element of $|\mathcal{F}|$ is labelled with one of the sets $\{\nu\}$, $\{\epsilon\}$, $\{2,\nu\}$, $\{4,\nu\}$, $\{1,3\}$ and $\{1,\epsilon\}$; and that the resulting digraph $\mathcal{F}$ is connected (as otherwise we deal with the connected components one-by-one).

We shall now prove the following claim: if $\mathcal{F}$ is such that every element is labelled either $\{\nu\}$ or $\{\epsilon\}$ then there is a homomorphism from $\mathcal{F}$ to $\lambda_2(\mathcal{H})$ which takes every element of $|\mathcal{F}|$ to an appropriate (w.r.t. its label set) element of $|\lambda_2(\mathcal{H})|$ if, and only if, there is a homomorphism from the following digraph $\mathcal{H}'$ to $\mathcal{H}$:

- the vertex set of $\mathcal{H}'$ is $\{x : x \in |\mathcal{F}| \text{ is labelled } \{\nu\}\}$; and
- for every $x, y \in |\mathcal{H}'|$, there is an edge (x,y) if, and only if, there is an element $z \in |\mathcal{F}|$ labelled $\{\epsilon\}$ such that $f^{\mathcal{F}}(z) = x$ and $g^{\mathcal{F}}(z) = y$.

Suppose that $\psi : \mathcal{F} \to \lambda_2(\mathcal{H})$ (where the homomorphism respects the label sets). Let $x \in |\mathcal{H}'|$. So, $x \in |\mathcal{F}|$ and is labelled $\{\nu\}$. Hence, $\psi(x) = v_u$, for some $u \in |\mathcal{H}|$. Define $\varphi : |\mathcal{H}'| \to |\mathcal{H}|$ by $\varphi(x) = u$.

Suppose that (x,y) is an edge of the digraph $\mathcal{H}'$. Hence, $\psi(x) = v_u$ and $\psi(y) = v_w$, for some $u, w \in |\mathcal{H}|$. By definition, there is $z \in |\mathcal{F}|$ labelled $\{\epsilon\}$ such that $f^{\mathcal{F}}(z) = x$ and $g^{\mathcal{F}}(z) = y$. So, $f^{\lambda_2(\mathcal{H})}(\psi(z)) = \psi(x) = v_u$ and $g^{\lambda_2(\mathcal{H})}(\psi(z)) = \psi(y) = v_w$; and by the definition of $\lambda_2(\mathcal{H})$, (u,w) is an edge of the digraph $\mathcal{H}$. Hence, $\varphi : \mathcal{H}' \to \mathcal{H}$.

Conversely, suppose that $\varphi : \mathcal{H}' \to \mathcal{H}$. Let $x \in |\mathcal{F}|$ be labelled $\{\nu\}$; so, $x \in |\mathcal{H}'|$. Define $\psi(x) = v_{\varphi(x)} \in |\lambda_2(\mathcal{H})|$. Let $z \in |\mathcal{F}|$ be labelled $\{\epsilon\}$; so, $f^{\mathcal{F}}(z) = x$, for some x with x labelled $\{\nu\}$, and $g^{\mathcal{F}}(z) = y$, for some y with y labelled $\{\nu\}$. Define $\psi(z) = e_{\varphi(x),\varphi(y)} \in |\lambda_2(\mathcal{H})|$.

Suppose that $f^{\mathcal{F}}(z) = x$; and so z is labelled $\{\epsilon\}$ and x is labelled $\{\nu\}$. We have that $f^{\lambda_2(\mathcal{H})}(\psi(z)) = f^{\lambda_2(\mathcal{H})}(e_{\varphi(x),\varphi(y)})$, for some $y \in |\mathcal{F}|$; and so $f^{\lambda_2(\mathcal{H})}(\psi(z)) = v_{\varphi(x)} = \psi(x)$. A similar argument for the case when $g^{\mathcal{F}}(z) = y$ yields that $\psi : \mathcal{F} \to \lambda_2(\mathcal{H})$, and so we have proved our claim.

Let us return to the situation prior to the statement of our claim. Suppose that $x \in |\mathcal{F}|$ is labelled with either $\{\nu\}$ or $\{\epsilon\}$. As $\mathcal{F}$ is connected, this effectively fixes every label set to consist of at most one element, and we can remove elements from $|\mathcal{F}|$ labelled $\{1\}$, $\{2\}$, $\{3\}$ or $\{4\}$ (as above). The above claim yields the result (note that we did not require $\mathcal{F}$ to be connected to prove our claim).

Alternatively, suppose that there does not exist an element $x \in |\mathcal{F}|$ labelled with either $\{\nu\}$ or $\{\epsilon\}$. Choose some element $x \in |\mathcal{F}|$ and remove one element from $\Lambda(x)$. This effectively forces every other label set to have size at most one, and we can remove elements from $|\mathcal{F}|$ labelled $\{1\}$, $\{2\}$, $\{3\}$ or $\{4\}$ (as above). Removing the alternative element from $\Lambda(x)$ results in a similar scenario. Regardless, the above claim yields the result. □

Hence, we obtain the main theorem of this section.

Theorem 10 *For any σ_2-structure $\mathcal{H}$, the problem CSP$(\mathcal{H})$ is polynomial-time Turing-equivalent to the problem Hom-Alg$_2(\lambda_2(\mathcal{H}))$. Thus, any non-uniform constraint satisfaction problem where the template is a relational structure is polynomial-time Turing-equivalent to a non-uniform constraint satisfaction problem where the template is a λ_2-structure of the form $\lambda_2(\mathcal{H})$, for some σ_2-structure $\mathcal{H}$.*

Proof The first part of the theorem follows immediately from Lemma 8 and Proposition 9. The second part follows from Theorem 10 of [9]. □

Essentially, Theorem 10 tells us that $\Delta(\sigma:\text{ relational}) \leq \Delta(\lambda_2)$ and that $\Delta(\sigma_2) \leq \Delta(\lambda_2)$ ([9] tells us that $\Delta(\sigma:\text{ relational}) \leq \Delta(\sigma_2)$ and $\Delta(\sigma_2) \leq \Delta(\sigma:\text{ relational})$, *i.e.*, that $\Delta(\sigma_2) \equiv \Delta(\sigma:\text{ relational})$).

We can obtain a yet more restricted version of Theorem 10, in terms of the structure $\lambda_2(\mathcal{H})$.

A *matching* is a unary function h such that $h(x) \neq x$ and $h(h(x)) = x$, for all x. Let $\mathcal{F}$ be a λ_2-structure. We shall define a λ_2-structure $\mu(\mathcal{F})$ that is bipartite and where the function g in $\mu(\mathcal{F})$ is a matching. The domain of $\mu(\mathcal{F})$ consists of two disjoint copies of $|\mathcal{F}|$; so, for every element $x \in |\mathcal{F}|$, there are two elements $x_1, x_2 \in |\mu(\mathcal{F})|$ (the former we call an index-1 element, the latter an index-2 element). We define $g(x_1) = x_2$ and $g(x_2) = x_1$ in $\mu(\mathcal{F})$. Also, if $f(x) = y$ in $\mathcal{F}$ then we define $f(x_1) = y_2$ in $\mu(\mathcal{F})$; and if $g(x) = y$ in $\mathcal{F}$ then we define $f(x_2) = y_1$ in $\mu(\mathcal{F})$.

Suppose that we construct the structure $\mu(\mathcal{F})$ from the structure $\mathcal{F}$. We can undo the construction of $\mu(\mathcal{F})$ to get not only $\mathcal{F}$ but also the λ_2-structure $\mathcal{F}_s$, by swapping the indices 1 and 2 on the corresponding pairs of elements of $\mu(\mathcal{F})$. So, for example, if $f(x_1) = y_2$ in $\mu(\mathcal{F})$ then $g(x) = y$ in $\mathcal{F}_s$; that is, $\mathcal{F}_s$ is $\mathcal{F}$ with f and g swapped.

Lemma 11 *For any λ_2-structures $\mathcal{F}$ and $\mathcal{F}'$:*

- *if $\mathcal{F} \to \mathcal{F}'$ then $\mu(\mathcal{F}) \to \mu(\mathcal{F}')$; and*
- *if $\mathcal{F}$ is connected and $\mu(\mathcal{F}) \to \mu(\mathcal{F}')$ then either $\mathcal{F} \to \mathcal{F}'$ or $\mathcal{F}_s \to \mathcal{F}'$.*

Proof Suppose that $\varphi : \mathcal{F} \to \mathcal{F}'$. Define $\psi : |\mu(\mathcal{F})| \to |\mu(\mathcal{F}')|$ by $\psi(x_1) = \varphi(x)_1$ and $\psi(x_2) = \varphi(x)_2$. Let $x_1, y_2 \in |\mu(\mathcal{F})|$ be such that $f^{\mu(\mathcal{F})}(x_1) = y_2$; and so $f^{\mathcal{F}}(x) = y$. Thus, $f^{\mathcal{F}'}(\phi(x)) = \phi(y)$, and so $f^{\mu(\mathcal{F}')}(\phi(x)_1) = \phi(y)_2$, with $f^{\mu(\mathcal{F}')}(\psi(x_1)) = \psi(y_2)$. The other case for $f^{\mu(\mathcal{F})}$ is handled similarly, and the fact that ψ preserves $g^{\mu(\mathcal{F})}$ is trivial. Hence, $\psi : \mu(\mathcal{F}) \to \mu(\mathcal{F}')$.

Suppose that $\mathcal{F}$ is connected and that $\psi : \mu(\mathcal{F}) \to \mu(\mathcal{F}')$. Define $\varphi : |\mathcal{F}| \to |\mathcal{F}'|$ by $\varphi(x) = y'$ if $\{\psi(x_1), \psi(x_2)\} = \{y'_1, y'_2\}$ (note that φ is well-defined). Suppose that $f^{\mathcal{F}}(x) = y$; and so $f^{\mu(\mathcal{F})}(x_1) = y_2$. Hence, $f^{\mu(\mathcal{F}')}(\psi(x_1)) = \psi(y_2)$. There are two possibilities: every index-1 element of $\mu(\mathcal{F})$ is mapped under ψ to an index-1 element of $\mu(\mathcal{F}')$; or every index-1 element of $\mu(\mathcal{F})$ is mapped under ψ to an index-2 element of $\mu(\mathcal{F}')$ (this is because $\mathcal{F}$ is connected). Suppose that $\{\psi(x_1), \psi(x_2)\} = \{x'_1, x'_2\}$ and $\{\psi(y_1), \psi(y_2)\} = \{y'_1, y'_2\}$. Thus, either $f^{\mathcal{F}'}(x') = y'$ or $g^{\mathcal{F}'}(x') = y'$. Moreover, whether $f^{\mathcal{F}'}(x') = y'$ or $g^{\mathcal{F}'}(x') = y'$ is independent of x and y. Similarly, if $g^{\mathcal{F}}(x) = y$ then either $g^{\mathcal{F}'}(x') = y'$ or $f^{\mathcal{F}'}(x') = y'$. Consequently, either $\mathcal{F} \to \mathcal{F}'$ or $\mathcal{F}_s \to \mathcal{F}'$. $\square$

Corollary 12 *For any σ_2-structure $\mathcal{H}$, the problem CSP($\mathcal{H}$) is polynomial-time Turing-equivalent to the problem Hom-Alg$_2$($\mu(\lambda_2(\mathcal{H}))$). Thus, any non-uniform constraint satisfaction problem where the template is a relational structure is polynomial-time Turing-equivalent to a non-uniform constraint satisfaction problem where the template is a λ_2-structure of the form $\mu(\lambda_2(\mathcal{H}))$, for some σ_2-structure $\mathcal{H}$.*

Proof Let $\mathcal{F}$ be a λ_2-structure. We shall transform $\mathcal{F}$ so that $\mathcal{F}$ is of the form $\mu(\mathcal{F}')$ and where $\mathcal{F} \to \mu(\lambda_2(\mathcal{H}))$ if, and only if, either $\mathcal{F}' \to \lambda_2(\mathcal{H})$ or $\mathcal{F}'_s \to \lambda_2(\mathcal{H})$.

We may assume that $\mathcal{F}$ is connected (otherwise we work with each connected component in turn) and bipartite (otherwise we define $\mathcal{F}'$ to be any λ_2-structure not in *Hom-Alg*$_2$($\mu(\lambda_2(\mathcal{H}))$)). If there is a homomorphism $\mathcal{F} \to \mu(\lambda_2(\mathcal{H}))$ then, for any element $x \in |\mathcal{F}|$, the elements x and $g^{\mathcal{F}}(g^{\mathcal{F}}(x))$ must map to the same element of $|\mu(\lambda_2(\mathcal{H}))|$. Thus, we can identify the elements x and $g^{\mathcal{F}}(g^{\mathcal{F}}(x))$ in $\mathcal{F}$ as a single element x. Of course, this may result in $f^{\mathcal{F}}$ or $g^{\mathcal{F}}$ being multiple-valued at x, and if this is the case then we identify the two values of $f^{\mathcal{F}}$ or the two values of $g^{\mathcal{F}}$. Note that there is a homomorphism from the amended $\mathcal{F}$ to $\mu(\lambda_2(\mathcal{H}))$ if, and only if, there is a homomorphism from the original $\mathcal{F}$ to $\mu(\lambda_2(\mathcal{H}))$ (a homomorphism from a multiple-valued $\mathcal{F}$ must respect all functional identities). We iterate this process until we can go no further. Eventually, we obtain that $\mathcal{F}$ is connected and bipartite, and that $g^{\mathcal{F}}$ is a matching.

We may pair up elements x_1 and x_2 for which $g^{\mathcal{F}}(x_1) = x_2$ and $g^{\mathcal{F}}(x_2) = x_1$ so that all elements indexed by 1 (resp. 2) are on the same side of the bipartition of $\mathcal{F}$. Hence, $\mathcal{F} = \mu(\mathcal{F}')$ where: $\mathcal{F}'$ has a single element x for each pair (x_1, x_2); $f^{\mathcal{F}'}(x) = y$ if $f^{\mathcal{F}}(x_1) = y_2$; and $g^{\mathcal{F}'}(x) = y$ if $g^{\mathcal{F}}(x_2) = y_1$. By Lemma 11, $\mathcal{F} \to \mu(\lambda_2(\mathcal{H}))$ if, and only if, either $\mathcal{F}' \to \lambda_2(\mathcal{H})$ or $\mathcal{F}'_s \to \lambda_2(\mathcal{H})$. So, by Proposition 9, there is a polynomial-time Turing-reduction from *Hom-Alg*$_2(\mu(\lambda_2(\mathcal{H})))$ to CSP$(\mathcal{H})$.

Conversely suppose that $\mathcal{G}$ is a σ_2-structure. Let $\mathcal{G}^*$ be the disjoint union of $\mathcal{G}$ and $\mathcal{H}$. We shall show that $\mathcal{G} \to \mathcal{H}$ if, and only if, $\mu(\lambda_2(\mathcal{G}^*)) \to \mu(\lambda_2(\mathcal{H}^*))$.

Suppose that $\mathcal{G} \to \mathcal{H}$. Then $\mathcal{G}^* \to \mathcal{H}$ and $\lambda_2(\mathcal{G}^*) \to \lambda_2(\mathcal{H})$ by Lemma 8. So, $\mu(\lambda_2(\mathcal{G}^*)) \to \mu(\lambda_2(\mathcal{H}))$ by Lemma 11.

Conversely, suppose that $\mu(\lambda_2(\mathcal{G}^*)) \to \mu(\lambda_2(\mathcal{H}))$. As $\lambda_2(\mathcal{G}^*)$ is connected, we can apply Lemma 11 to obtain that either $\lambda_2(\mathcal{G}^*) \to \lambda_2(\mathcal{H})$ or $\lambda_2(\mathcal{G}^*)_s \to \lambda_2(\mathcal{H})$. In the former case, we get that $\mathcal{G}^* \to \mathcal{H}$ by Lemma 8. In the latter case, we get that $\lambda_2(\mathcal{G}^*) \to \lambda_2(\mathcal{H})_s$. However, if we denote by $\mathcal{H}_s$ the digraph $\mathcal{H}$ where all edges are reversed then it is straightforward to see that $\lambda_2(\mathcal{H})_s = \lambda_2(\mathcal{H}_s)$; and so by Lemma 8, $\mathcal{G}^* \to \mathcal{H}_s$ and thus $\mathcal{H} \to \mathcal{H}_s$. If $\psi : \mathcal{H} \to \mathcal{H}_s$ then the same map ψ is a homomorphism from $\mathcal{H}_s$ to $\mathcal{H}$. Thus, composing homomorphisms gives us that $\mathcal{G}^* \to \mathcal{H}$. Hence, there is a polynomial-time reduction from CSP$(\mathcal{H})$ to *Hom-Alg*$_2(\mu(\lambda_2(\mathcal{H})))$, and the result follows. $\square$

Essentially, Corollary 12 tells us that $\Delta(\sigma$: relational$) \leq \Delta(\lambda_2$: bipartite, g a matching). Thus, if we wish to establish dichotomy results for non-uniform constraint satisfaction problems with template a λ_2-structure, but without resolving $\Delta(\sigma : \sigma$ relational), then we need to restrict the templates beyond being bipartite with one of the unary functions a matching. We shall consider restricted structures later in this paper.

5 Constants, unary functions and relations

In this section we prove what amounts to the converse of Theorem 10. We then go on to examine the complexity of problems of the form *Hom-Alg*$_2(\mathcal{T})$ for certain elementary λ_2-structures $\mathcal{T}$.

Theorem 13 *Let $\mathcal{B}$ be some σ-structure where σ consists of constant symbols, unary function symbols and relation symbols. There exists a relational signature σ' and a σ'-structure $\mathcal{B}'$ such that CSP$(\mathcal{B})$ and CSP$(\mathcal{B}')$ are polynomial-time equivalent.*

Proof Let $\sigma = \langle c_1, \dots, c_k, f_1, \dots, f_l, R_1, \dots, R_m \rangle$, where each c_i (resp. f_i, R_i) is a constant (resp. unary function, relation) symbol; and let $\sigma' = \langle C_1, \dots, C_k, F_1, \dots, F_l, R_1, \dots, R_m \rangle$, where each C_i is a unary relation symbol and each F_i is a binary relation symbol. Define the σ'-structure $\mathcal{B}'$ as follows:

- $|\mathcal{B}'| = |\mathcal{B}|$;
- for each $i \in \{1, 2, \dots, k\}$, $C_i^{\mathcal{B}'} = \{c_i^{\mathcal{B}}\}$;
- for each $i \in \{1, 2, \dots, l\}$ and for any $u, w \in |\mathcal{B}'|$, $F_i^{\mathcal{B}'}(u, w)$ holds if, and only if, $f_i^{\mathcal{B}}(u) = w$; and
- for each $i \in \{1, 2, \dots, m\}$, $R_i^{\mathcal{B}'} = R_i^{\mathcal{B}}$.

Suppose that $\mathcal{A}'$ is a σ'-structure. Suppose further than $C_i^{\mathcal{A}'} = \emptyset$, for some $i \in \{1, 2, \dots, k\}$: so $I = \{i : C_i^{\mathcal{A}'} = \emptyset \text{ where } 1 \leq i \leq k\}$ is non-empty. Add a new element to the domain of $\mathcal{A}'$ and for every $i \in I$, set the relation C_i of the amended structure to consist of the set containing only this new element (keep all the other relations as they are). Clearly, the amended structure $\mathcal{A}'$ is in CSP($\mathcal{B}'$) if, and only if, the original structure $\mathcal{A}'$ is in CSP($\mathcal{B}'$). Hence, we assume that every relation $C_i^{\mathcal{A}'}$, for $i \in \{1, 2, \dots, k\}$, is non-empty.

We shall now amend $\mathcal{A}'$ so that in the amended structure, $\mathcal{A}''$, for every $i \in \{1, 2, \dots, l\}$ and for every $u \in |\mathcal{A}''|$, there is at least one $w \in |\mathcal{A}''|$ such that $F_i^{\mathcal{A}''}(u, w)$ holds. We begin by defining the σ'-structure $\mathcal{B}'_p$ as follows:

- $|\mathcal{B}'_p| = |\mathcal{B}'|^p$, where $|\mathcal{B}'|$ contains p elements;
- for every $i \in \{1, 2, \dots, k\}$, $C_i^{\mathcal{B}'_p} = \emptyset$;
- for every $i \in \{1, 2, \dots, l\}$ and for any $\bar{u}, \bar{v} \in |\mathcal{B}'_p|$, $F_i^{\mathcal{B}'_p}(\bar{u}, \bar{v})$ holds if, and only if, $v_j = f^{\mathcal{B}}(u_j)$, for every $j \in \{1, 2, \dots, p\}$; and
- for every $i \in \{1, 2, \dots, m\}$, $R_i^{\mathcal{B}'_p} = \emptyset$.

Let $\bar{v} \in |\mathcal{B}'_p|$ be such that all components of $\bar{v}$ are distinct, and let $w \in |\mathcal{B}'|$. Then w appears as some coordinate in $\bar{v}$: w.l.o.g. assume that w appears as the first coordinate. Define the map $\varphi_{\bar{v}}^{w} : |\mathcal{B}'_p| \to |\mathcal{B}'|$ as:

$$\varphi_{\bar{v}}^{w}(\bar{u}) = u_1.$$

It is straightforward to verify that $\varphi_{\bar{v}}^{w}$ is a homomorphism from $\mathcal{B}'_p$ to $\mathcal{B}'$. Observe that there exists a homomorphism from $\mathcal{B}'_p$ to $\mathcal{B}'$ which maps *any* chosen element of $|\mathcal{B}'_p|$ whose coordinates are all distinct to *any* chosen element of $|\mathcal{B}'|$.

We construct the σ'-structure $\mathcal{A}''$ as follows. First, fix some element $\bar{u} \in |\mathcal{B}'_p|$ for which all the components are distinct. Define:

- $|\mathcal{A}''| = |\mathcal{A}'| \times |\mathcal{B}'_p|$;
- for each $i \in \{1, 2, \ldots, k\}$, $C_i^{\mathcal{A}''}((w, \bar{v}))$ holds if, and only if, $\bar{v} = \bar{u}$ and $C_i^{\mathcal{A}'}(w)$ holds;
- for each $i \in \{1, 2, \ldots, l\}$ and for any $(w_1, \bar{v}^1), (w_2, \bar{v}^2) \in |\mathcal{B}''|$, $F_i^{\mathcal{A}''}((w_1, \bar{v}^1), (w_2, \bar{v}^2))$ holds if, and only if, either ($\bar{v}^1 = \bar{v}^2 = \bar{u}$ and $F_i^{\mathcal{A}'}(w_1, w_2)$ holds) or ($w_1 = w_2$ and $F_i^{\mathcal{B}'_p}(\bar{v}^1, \bar{v}^2)$ holds); and
- for every $i \in \{1, 2, \ldots, m\}$ and for any $(w_1, \bar{v}^1), (w_2, \bar{v}^2), \ldots, (w_a, \bar{v}^a) \in |\mathcal{A}''|$, $R_i^{\mathcal{A}''}((w_1, \bar{v}^1), (w_2, \bar{v}^2), \ldots, (w_a, \bar{v}^a))$ holds if, and only if, $\bar{v}^1 = \bar{v}^2 = \ldots = \bar{v}^a = \bar{u}$ and $R_i^{\mathcal{A}'}(w_1, w_2, \ldots, w_a)$ holds (where a is the arity of R_i).

The structure $\mathcal{A}''$ can be depicted in Fig. 3, where the relations C_i and R_i are omitted and an arrow labelled i from one vertex to another denotes that the associated pair of vertices is in the binary relation F_i. Note that the substructure of $\mathcal{A}''$ induced by the set of elements $|\mathcal{A}'| \times \{\bar{u}\}$ is isomorphic to $\mathcal{A}'$.

Figure 3. The σ'-structure $\mathcal{A}''$.

Suppose that there is a homomorphism from $\mathcal{A}'$ to $\mathcal{B}'$. Then this homomorphism can be extended to a homomorphism of $\mathcal{A}''$ to $\mathcal{B}'$ by utilizing the observation at the end of the paragraph in which the construction of $\mathcal{B}'_p$ is detailed. Conversely, if there is a homomorphism from $\mathcal{A}''$ to $\mathcal{B}'$ then the restriction of this map to the set of elements $|\mathcal{A}'| \times \{\bar{u}\}$ yields a homomorphism from $\mathcal{A}'$ to $\mathcal{B}'$. So, $\mathcal{A}' \in \mathrm{CSP}(\mathcal{B}')$ if, and only if, $\mathcal{A}'' \in \mathrm{CSP}(\mathcal{B}')$.

Suppose that $w_1, w_2 \in |\mathcal{A}''|$ are such that $w_1 \neq w_2$ and both $C_i^{\mathcal{A}''}(w_1)$ and $C_i^{\mathcal{A}''}(w_2)$ hold, for some $i \in \{1, 2, \ldots, k\}$. The structure obtained from $\mathcal{A}''$ by identifying w_1 and w_2 is such that there is a homomorphism to $\mathcal{B}'$ if, and only if, there is homomorphism from $\mathcal{A}''$ to $\mathcal{B}'$ (as $\mathcal{B}'$ has been derived from $\mathcal{B}$ by 'converting' the constants and functions into relations mirroring these constants and functions). Suppose that $w_1, w_2, w_3 \in |\mathcal{A}''|$ are such that $w_2 \neq w_3$ and both $F_i^{\mathcal{A}''}(w_1, w_2)$ and $F_i^{\mathcal{A}''}(w_1, w_3)$ hold, for some $i \in \{1, 2, \ldots, l\}$. The structure obtained from $\mathcal{A}''$ by identifying w_2 and w_3 is such that there is a homomorphism to $\mathcal{B}'$ if, and only if, there is homomorphism from $\mathcal{A}''$

to $\mathcal{B}'$. By repeatedly identifying vertices (if necessary) as above, we may assume that $\mathcal{A}''$ is such that: for every $i \in \{1, 2, \ldots, k\}$, there is exactly one element in $C_i^{\mathcal{A}''}$; and for every $i \in \{1, 2, \ldots, l\}$ and for every $w_1 \in |\mathcal{A}''|$, there is exactly one element $w_2 \in |\mathcal{A}''|$ such that $F_i^{\mathcal{A}''}(w_1, w_2)$ holds. Finally, let $\mathcal{A}$ be the σ-structure obtained from $\mathcal{A}''$ in the natural way. Clearly, $\mathcal{A}'' \in \mathrm{CSP}(\mathcal{B}')$ if, and only if, $\mathcal{A} \in \mathrm{CSP}(\mathcal{B})$. Moreover, $\mathcal{A}$ can be constructed from $\mathcal{A}'$ in time polynomial in the size of $\mathcal{A}'$.

Conversely, let $\mathcal{A}$ be a σ-structure. Let $\mathcal{A}'$ be the σ'-structure obtained from $\mathcal{A}$ in the natural way. Clearly, $\mathcal{A} \in \mathrm{CSP}(\mathcal{B})$ if, and only if, $\mathcal{A}' \in \mathrm{CSP}(\mathcal{B}')$; and $\mathcal{A}'$ can be constructed from $\mathcal{A}$ in time polynomial in the size of $\mathcal{A}$. The result follows. $\square$

A corollary of Theorems 10 and 13 is that $\Delta(\sigma : \text{ relational}) \equiv \Delta(\lambda_2)$.

Now let us turn to non-uniform constraint satisfaction problems $\mathrm{CSP}(\mathcal{T})$ where the λ_2-structure $\mathcal{T}$ is bipartite and where the g-digraph is a single directed cycle. For any such λ_2-structure $\mathcal{T}$, let us define a λ_3-structure $\tilde{\mathcal{T}}$, where λ_3 consists of the unary function symbols f, g_1 and g_2, with the property that both g_1 and g_2 are matchings. Let T_1 and T_2 form the bipartition of $|\mathcal{T}|$. The domain of $\tilde{\mathcal{T}}$ is $|\mathcal{T}|$; if $x_1 \in T_1$ then $g_1^{\tilde{\mathcal{T}}}(x_1) = g^{\mathcal{T}}(x_1)$, with $g_1^{\tilde{\mathcal{T}}}(g^{\mathcal{T}}(x_1)) = x_1$; if $x_2 \in T_2$ then $g_2^{\tilde{\mathcal{T}}}(x_2) = g^{\mathcal{T}}(x_2)$, with $g_2^{\tilde{\mathcal{T}}}(g^{\mathcal{T}}(x_2)) = x_2$; and $f^{\tilde{\mathcal{T}}}$ is defined identically to $f^{\mathcal{T}}$. The problem *Hom-Alg*$_3(\tilde{\mathcal{T}})$ is $\mathrm{CSP}(\tilde{\mathcal{T}})$.

Theorem 14 *Let $\mathcal{T}$ be a λ_2-structure that is bipartite and whose g-digraph is a directed cycle. The problem Hom-Alg$_2(\mathcal{T})$ is polynomial-time Turing-equivalent to the problem Hom-Alg$_3(\tilde{\mathcal{T}})$. In consequence, there exists such a λ_2-structure $\mathcal{T}$ with the property that Hom-Alg$_2(\mathcal{T})$ is complete for* **NP** *via polynomial-time Turing-reductions.*

Proof Let $\mathcal{S}$ be a λ_3-structure. We may clearly assume that $\mathcal{S}$ is bipartite (as otherwise $\mathcal{S} \notin$ *Hom-Alg*$_3(\tilde{\mathcal{T}})$) and connected (as otherwise we take each connected component in turn). Furthermore, just as we did in the proof of Corollary 12, we may assume that $g_1^{\mathcal{S}}$ and $g_2^{\mathcal{S}}$ are matchings. Fix some element $u \in |\mathcal{S}|$. First, we shall ascertain (using an oracle for *Hom-Alg*$_2(\mathcal{T})$) whether there is a homomorphism from $\mathcal{S}$ to $\tilde{\mathcal{T}}$ where u is mapped to an element of T_1 (as in the construction of $\tilde{\mathcal{T}}$ above: we proceed similarly when looking for a homomorphism mapping u to an element of T_2).

We begin by constructing a λ_2-structure $\mathcal{F}_u$. Begin with the domain $|\mathcal{S}|$ (we shall extend it presently). Assume that a homomorphism of $\mathcal{S}$ to $\tilde{\mathcal{T}}$ (if it exists) takes u to an element of T_1. As $\mathcal{S}$ is connected and bipartite, we can now determine which side of the bipartition every element of $|\mathcal{S}|$ is mapped to. If $w \in |\mathcal{S}|$ maps to T_1 then we define $g^{\mathcal{F}_u}(w) = g_1^{\mathcal{S}}(w)$; otherwise we define $g^{\mathcal{F}_u}(w) = g_2^{\mathcal{S}}(w)$. The function $f^{\mathcal{F}_u}$ is identical to $f^{\mathcal{S}}$. Trivially, φ is a homomorphism from $\mathcal{S}$ to $\tilde{\mathcal{T}}$ mapping u to an element of T_1 if, and only if, φ is a homomorphism from $\mathcal{F}_u$ to $\mathcal{T}$ mapping u to an element of T_1 (recall that $g_1^{\mathcal{S}}$ and $g_2^{\mathcal{S}}$ are matchings.

For every $v \in T_1$, form the structure $\mathcal{F}_u + \mathcal{T}_v$ by taking disjoint copies of $\mathcal{F}_u$ and $\mathcal{T}$ and identifying the elements $u \in |\mathcal{F}_u|$ and $v \in |\mathcal{T}|$. Note that this yields a 'structure' where the functions f and g have multiple values; and any homomorphism from $\mathcal{F}_u + \mathcal{T}_v$ to $\mathcal{T}$, say, must respect all functional identities (as in the proof of Corollary 12).

Suppose that $\mathcal{F}_u \to \mathcal{T}$ where u is mapped to an element $v \in T_1$. Then $\mathcal{F}_u + \mathcal{T}_v \to \mathcal{T}$. Conversely, suppose that $\psi : \mathcal{F}_u + \mathcal{T}_v \to \mathcal{T}$, for some $v \in T_1$. Suppose further that the element u of $\mathcal{F}_u + \mathcal{T}_v$ (that is, the element resulting from the identification) is mapped under ψ to an element of T_2. Thus, ψ restricted to the copy of $\mathcal{T}$ in $\mathcal{F}_u + \mathcal{T}_v$ is a homomorphism of $\mathcal{T}$ to $\mathcal{T}$ taking every element of T_1 to T_2 and *vice versa*. Hence, by composing maps we get that there is a homomorphism of $\mathcal{F}_u + \mathcal{T}_v$ to $\mathcal{T}$ taking u to an element of T_1; and so $\mathcal{F}_u \to \mathcal{T}$ where u is mapped to an element of T_1. Hence, $\mathcal{F}_u \to \mathcal{T}$ where u is mapped to an element of T_1 if, and only if, $\mathcal{F}_u + \mathcal{T}_v \to \mathcal{T}$, for some $v \in T_1$.

Just as we did in the proof of Corollary 12, we resolve the multiple values of any $\mathcal{F}_u + \mathcal{T}_v$ by repeatedly identifying any multiple values until the functions f and g become single-valued. We denote the resulting λ_2-structure by $\mathcal{F}_u(\mathcal{T}_v)$. Just as in Corollary 12, there is a homomorphism from $\mathcal{F}_u + \mathcal{T}_v$ to $\mathcal{T}$ if, and only if, there is a homomorphism from $\mathcal{F}_u(\mathcal{T}_v)$ to $\mathcal{T}$. Hence, there is a polynomial-time Turing-reduction from $\textit{Hom-Alg}_3(\tilde{\mathcal{T}})$ to $\textit{Hom-Alg}_2(\mathcal{T})$.

Conversely, let $\mathcal{F}$ be a λ_2-structure. Again, we can assume that $\mathcal{F}$ is connected and bipartite. We proceed as we did above and fix some element $u \in |\mathcal{F}|$. We shall ascertain whether there is a homomorphism from $\mathcal{F}$ to $\mathcal{T}$ where u is mapped to an element of T_1. This time, we construct a λ_3-structure $\mathcal{S}_u$. If $w \in |\mathcal{F}|$ maps to an element of T_1 in $\mathcal{T}$ then we define $g_1^{\mathcal{S}_u}(w) = g^{\mathcal{F}}(w)$; otherwise we define that $g_2^{\mathcal{S}_u}(w) = g^{\mathcal{F}}(w)$. The function $f^{\mathcal{S}_u}$ is identical to $f^{\mathcal{F}}$. This only partially defines the functions g_1 and g_2 of $\mathcal{S}_u$. Nevertheless, it is straightforward to see that there is a homomorphism $\mathcal{F}$ to $\mathcal{T}$ where u is mapped to an element of T_1 if, and only if, there is a homomorphism $\mathcal{S}_u$ to $\tilde{\mathcal{T}}$ where u is mapped to an element of T_1. However, by proceeding as we did in the proof of Theorem 13 (where we built the structure $\mathcal{A}''$ from $\mathcal{A}'$ and $\mathcal{B}'$), from $\mathcal{S}_u$ and $\tilde{\mathcal{T}}$ we can build a λ_3-structure, call it $\mathcal{S}_u$ also, where the functions g_1 and g_2 are fully defined and where:

there is a homomorphism from the partially defined $\mathcal{S}_u$ to $\tilde{\mathcal{T}}$ where u is mapped to an element of T_1 if, and only if, there is a homomorphism from the fully defined $\mathcal{S}_u$ to $\tilde{\mathcal{T}}$ where u is mapped to an element of T_1.

Let $v \in T_1$. Just as we did above, we can form the structure $\mathcal{S}_u + \tilde{\mathcal{T}}_v$ by identifying $u \in |\mathcal{S}|$ and $v \in |\tilde{\mathcal{T}}|$. Again, the functions g_1 and g_2 of $\mathcal{S}_u + \tilde{\mathcal{T}}_v$ may have multiple values and again we use our process of identification to yield the structure $\mathcal{S}_u(\tilde{\mathcal{T}}_v)$. Reasoning as above yields that there is a polynomial-time Turing-reduction from *Hom-Alg*$_2(\mathcal{T})$ to *Hom-Alg*$_3(\tilde{\mathcal{T}})$.

By Corollary 12, there exists a bipartite λ_2-structure $\mathcal{F}'$ for which the function g is a matching and for which *Hom-Alg*$_2(\mathcal{F}')$ is complete for **NP** via polynomial-time Turing reductions. Define the bipartite λ_3-structure $\mathcal{S}$ from $\mathcal{F}'$ by setting: $|\mathcal{S}| = |\mathcal{F}'|$; $f^{\mathcal{S}} = f^{\mathcal{F}'}$; $g_1^{\mathcal{S}} = g^{\mathcal{F}'}$; and $g_2^{\mathcal{S}}$ is a matching such that the $g_1^{\mathcal{S}}$-graph and the $g_2^{\mathcal{S}}$-graph form a single cycle. By above, there exists a λ_2-structure $\mathcal{F}$ such that: the $g^{\mathcal{F}}$-digraph is a directed cycle; $\mathcal{S} = \tilde{\mathcal{F}}$; and *Hom-Alg*$_2(\mathcal{F})$ and *Hom-Alg*$_3(\mathcal{S})$ are polynomial-time Turing-equivalent.

An instance of the problem *Hom-Alg*$_2(\mathcal{F}')$ is an instance of the problem *Hom-Alg*$_3(\mathcal{S})$ where the function g_2 is not defined. As above, we can modify this instance so that g_2 is defined and there is a homomorphism from this amended instance to $\mathcal{S}$ if, and only if, there is a homomorphism from the original instance to $\mathcal{S}$. Hence, *Hom-Alg*$_3(\mathcal{S})$ is complete for **NP** via polynomial-time Turing-reductions; and so *Hom-Alg*$_2(\mathcal{F})$ is complete for **NP** via polynomial-time Turing-reductions. □

Theorem 14 is a general result that relates non-uniform constraint satisfaction problems with templates λ_2-structures in which one of the unary functions is a cycle with non-uniform constraint satisfaction problems with templates λ_3-structures in which two of the unary functions are matchings. It also tells us that even when we restrict our template $\mathcal{T}$ in the problem *Hom-Alg*$_2(\mathcal{T})$ so that it is bipartite and its g-digraph is a directed cycle, the resulting problem can still be complete for **NP** via polynomial-time Turing-reductions (note that we make no statement as regards the degree of difficulty of the problem $\Delta(\lambda_2:$ bipartite, g-digraph a directed cycle$)$). In fact, we can restrict the size of $\mathcal{T}$ to be as small as possible and, further, obtain a complete problem for **NP** via polynomial-time reductions.

Corollary 15 *There exist λ_2-structures $\mathcal{T}$ and $\hat{\mathcal{T}}$ such that:*

- $|\mathcal{T}| = 3$, *the g-digraph is a directed cycle, the f-digraph is connected and Hom-Alg$_2(\mathcal{T})$ is* **NP***-complete*; *and*

- $|\hat{\mathcal{T}}| = 6$, *the* g*-digraph is a directed cycle, the* f*-digraph is connected,* $\hat{\mathcal{T}}$ *is bipartite and Hom-Alg*$_2(\hat{\mathcal{T}})$ *is* **NP***-complete.*

Proof Define the λ_2-structure $\mathcal{T}$ to have: domain $\{0,1,2\}$; function $f^{\mathcal{T}}$ defined as $f^{\mathcal{T}}(0) = 1$, $f^{\mathcal{T}}(1) = 0$ and $f^{\mathcal{T}}(2) = 0$; and function $g^{\mathcal{T}}$ defined as $g^{\mathcal{T}}(0) = 1$, $g^{\mathcal{T}}(1) = 2$ and $g^{\mathcal{T}}(2) = 0$. Note that for any element $x \in |\mathcal{T}|$, the possible values for $(f^{\mathcal{T}}(x), f^{\mathcal{T}}(g^{\mathcal{T}}(x)), f^{\mathcal{T}}(g^{\mathcal{T}}(g^{\mathcal{T}}(x))))$ are $(1,0,0)$, $(0,1,0)$ and $(0,0,1)$, when x is 0, 2 and 1, respectively.

Consider the problem 1-in-3-SAT, defined over the signature $\sigma_3 = \langle R \rangle$, where R is a relation symbol of arity 3: an instance is a σ_3-structure $\mathcal{A}$, which is viewed as a collection of clauses of 3 distinct boolean variables via $\{X_i, X_j, X_k\}$ is a clause if, and only if, $R^{\mathcal{A}}(i,j,k)$ holds, where $i \neq j \neq k \neq i$; and a yes-instance is an instance for which there is a truth assignment making exactly one of the 3 boolean variables in each clause *true*. This problem is **NP**-complete [19].

Suppose we are given an instance $\mathcal{A}$ of 1-in-3-SAT where the clauses are $C_1, C_2, \ldots, C_m$ over the boolean variables $X_1, X_2, \ldots, X_n$. Define functions f and g over

$$\{1,2,\ldots,m\} \cup \{\dot{1}, \dot{2}, \ldots, \dot{n}\} \cup S,$$

where S is as yet undefined, as follows. For each: $i \in \{1,2,\ldots,m\}$:

if clause $C_i = \{X_{a(i)}, X_{b(i)}, X_{c(i)}\}$ then define f and g such that

$$f(i) = \dot{a(i)}; f(g(i)) = \dot{b(i)}; \text{ and } f(g(g(i))) = \dot{c(i)}.$$

The set S is chosen so as to facilitate these conditions. In more detail, $S = \{r_i, s_i : i = 1,2,\ldots,m\}$ and for each $i = 1,2,\ldots,m$:

- $f(i) = \dot{a(i)}$;
- $g(i) = r_i$; $f(r_i) = \dot{b(i)}$;
- $g(r_i) = s_i$; $f(s_i) = \dot{c(i)}$;

and $f(u) = u$ and $g(u) = u$, for any as yet undefined u. Denote by $\mathcal{F}$ the λ_2-structure whose functions f and g are as defined above.

Suppose that $\varphi : \{X_1, X_2, \ldots, X_n\} \to \{0,1\}$ is a satisfying truth assignment for the instance $\mathcal{A}$ of 1-in-3-SAT. Define the function ψ from $\{1,2,\ldots,m\} \cup \{\dot{1}, \dot{2}, \ldots, \dot{n}\} \cup S$ to $\{0,1,2\}$ as follows.

- For each $i \in \{1,2,\ldots,m\}$:
  - if $(\varphi(X_{a(i)}), \varphi(X_{b(i)}), \varphi(X_{c(i)})) = (1,0,0)$ then $\psi(i) = 0$;

- if $(\varphi(X_{a(i)}), \varphi(X_{b(i)}), \varphi(X_{c(i)})) = (0,1,0)$ then $\psi(i) = 2$;
- if $(\varphi(X_{a(i)}), \varphi(X_{b(i)}), \varphi(X_{c(i)})) = (0,0,1)$ then $\psi(i) = 1$.

- For each $i \in \{1, 2, \ldots, m\}$:
  - $\psi(r_i) = g^{\mathcal{T}}(\psi(i))$;
  - $\psi(s_i) = g^{\mathcal{T}}(g^{\mathcal{T}}(\psi(i)))$.

- For each $j \in \{1, 2, \ldots, n\}$, $\psi(\dot{j}) = \varphi(X_j)$.

We claim that ψ is a homomorphism from $\mathcal{F}$ to $\mathcal{T}$.

Suppose that the truth assignment φ is such that

$$(\varphi(X_{a(i)}), \varphi(X_{b(i)}), \varphi(X_{c(i)})) = (1,0,0).$$

Consider the identity $f^{\mathcal{F}}(i) = \dot{a(i)}$. Then $f^{\mathcal{T}}(\psi(i)) = f^{\mathcal{T}}(0) = 1$; and $\psi(\dot{a(i)}) = \varphi(X_{a(i)}) = 1$. Similarly, all other functional identities can be so verified. A similar analysis follows should the truth assignment φ be such that $(\varphi(X_{a(i)}), \varphi(X_{b(i)}), \varphi(X_{c(i)}))$ is $(0,1,0)$ or $(0,0,1)$. Thus, $\psi : \mathcal{F} \to \mathcal{T}$.

Conversely, suppose that $\psi : \mathcal{F} \to \mathcal{T}$. Define the truth assignment $\varphi : \{X_1, X_2, \ldots, X_n\} \to \{0,1\}$ by:

- $\varphi(X_{a(i)}) = 1$ if $\psi(i) = 0$, and 0 otherwise;
- $\varphi(X_{b(i)}) = 1$ if $\psi(i) = 2$, and 0 otherwise;
- $\varphi(X_{c(i)}) = 1$ if $\psi(i) = 1$, and 0 otherwise.

The map φ is clearly well-defined and a satisfying truth assignment for the instance $\mathcal{A}$. Hence, *Hom-Alg*$_2(\mathcal{T})$ is **NP**-complete.

Let $\mathcal{A}$ be an arbitrary λ_2-structure. Define the λ_2-structure $\hat{\mathcal{A}}$ as follows:

- the domain $|\hat{\mathcal{A}}| = \{x : x \in |\mathcal{A}|\} \cup \{\hat{x} : x \in |\mathcal{A}|\}$;
- if $f^{\mathcal{A}}(x) = y$ then $f^{\hat{\mathcal{A}}}(x) = \hat{x}$ and $f^{\hat{\mathcal{A}}}(\hat{x}) = y$, for $x \in |\mathcal{A}|$; and
- if $g^{\mathcal{A}}(x) = y$ then $g^{\hat{\mathcal{A}}}(x) = \hat{x}$ and $g^{\hat{\mathcal{A}}}(\hat{x}) = y$, for $x \in |\mathcal{A}|$.

For any $x \in |\mathcal{A}|$, $\hat{\hat{x}}$ denotes the element x.

Suppose that $\psi : \mathcal{A} \to \mathcal{T}$. Define the map $\hat{\psi} : |\hat{\mathcal{A}}| \to |\hat{\mathcal{T}}|$ by $\hat{\psi}(x) = \psi(x)$ and $\hat{\psi}(\hat{x}) = \widehat{\psi(x)}$, for $x \in |\mathcal{A}|$. Let $x, y \in |\hat{\mathcal{A}}|$ be such that $f^{\hat{\mathcal{A}}}(x) = y$. Then either $x \in |\mathcal{A}|$ and $y = \hat{x}$; or $x = \hat{u}$, for some $u \in |\mathcal{A}|$, and $f^{\mathcal{A}}(u) = y \in |\mathcal{A}|$. In the former case, as $\hat{\psi}(x) = \psi(x) \in |\mathcal{T}|$, we have that $f^{\hat{\mathcal{T}}}(\hat{\psi}(x)) =$

$\widehat{\hat{\psi}(x)} = \widehat{\psi(x)} = \hat{\psi}(\hat{x}) = \hat{\psi}(y)$. In the latter case, $f^{\mathcal{A}}(u) = y \in |\mathcal{A}|$; and so $f^{\mathcal{T}}(\psi(u)) = \psi(y)$. Thus, $f^{\hat{\mathcal{T}}}(\widehat{\psi(u)}) = \psi(y)$ and $f^{\hat{\mathcal{T}}}(\hat{\psi}(\hat{u})) = \hat{\psi}(y)$; that is, $f^{\hat{\mathcal{T}}}(\hat{\psi}(x)) = \hat{\psi}(y)$. The case for when $g^{\hat{\mathcal{A}}}(x) = y$ is identical; and so we have $\hat{\psi} : \hat{\mathcal{A}} \to \hat{\mathcal{T}}$.

Conversely, suppose that $\hat{\psi} : \hat{\mathcal{A}} \to \hat{\mathcal{T}}$. There are two cases: when $\hat{\psi}$ maps elements of $|\mathcal{A}|$ to elements of $|\mathcal{T}|$; and when $\hat{\psi}$ maps elements of $\{\hat{x} : x \in |\mathcal{A}|\}$ to elements of $|\mathcal{T}|$. Let us consider the first case. Define $\psi : |\mathcal{A}| \to |\mathcal{T}|$ by $\psi(x) = \hat{\psi}(x)$. Let $x, y \in |\mathcal{A}|$ be such that $f^{\mathcal{A}}(x) = y$. So, $f^{\hat{\mathcal{A}}}(\hat{x}) = y$ and $f^{\hat{\mathcal{T}}}(\hat{\psi}(\hat{x})) = \hat{\psi}(y)$. But $f^{\hat{\mathcal{A}}}(x) = \hat{x}$ and so $f^{\hat{\mathcal{T}}}(\hat{\psi}(x)) = \hat{\psi}(\hat{x})$. Thus, $f^{\mathcal{T}}(\psi(x)) = \psi(y)$. The situation when $g^{\mathcal{A}}(x) = y$ is similar.

Alternatively, when $\hat{\psi}$ maps elements of $\{\hat{x} : x \in |\mathcal{A}|\}$ to elements of $|\mathcal{T}|$, we define $\psi(x) = \widehat{\hat{\psi}(x)}$. Let $x, y \in |\mathcal{A}|$ be such that $f^{\mathcal{A}}(x) = y$. So, $f^{\hat{\mathcal{A}}}(x) = \hat{x}$ and $f^{\hat{\mathcal{A}}}(\hat{x}) = y$, with $f^{\hat{\mathcal{T}}}(\hat{\psi}(x)) = \hat{\psi}(\hat{x})$ and $f^{\hat{\mathcal{T}}}(\hat{\psi}(\hat{x})) = \hat{\psi}(y)$. But as $\hat{\psi}(\hat{x}) \in |\mathcal{T}|$, $f^{\hat{\mathcal{T}}}(\hat{\psi}(\hat{x})) = \widehat{\hat{\psi}(\hat{x})}$; and so $\widehat{\hat{\psi}(\hat{x})} = \hat{\psi}(y)$ with $\psi(y) = \widehat{\hat{\psi}(y)} = \hat{\psi}(\hat{x})$. Also, as $\hat{\psi}(x) \notin |\mathcal{T}|$, $f^{\hat{\mathcal{T}}}(\widehat{\hat{\psi}(x)}) = \hat{\psi}(x)$; so $f^{\mathcal{T}}(\widehat{\hat{\psi}(x)}) = \hat{\psi}(\hat{x})$ with $f^{\mathcal{T}}(\psi(x)) = \psi(y)$. The situation when $g^{\mathcal{A}}(x) = y$ is similar. Hence, $\psi : \mathcal{A} \to \mathcal{T}$ and consequently, *Hom-Alg*$_2(\hat{\mathcal{T}})$ is **NP**-complete and the result follows. $\square$

We stated prior to Corollary 15 that this corollary was optimal in terms of the size of the structures $\mathcal{T}$ and $\hat{\mathcal{T}}$. That this is true for $\mathcal{T}$ is obvious; and we now show that this is true for $\hat{\mathcal{T}}$ also.

Lemma 16 *Let $\hat{\mathcal{T}}$ be a bipartite λ_2-structure of size 4 whose g-digraph is a directed cycle. The problem Hom-Alg$_2(\hat{\mathcal{T}})$ is solvable in polynomial-time.*

Proof Let $\hat{\mathcal{T}}$ be a bipartite λ_2-structure of size 4, with domain $\{0, 1, 2, 3\}$, and whose g-digraph is the directed cycle $(0, 1, 2, 3)$. Let $\mathcal{A}$ be an arbitrary λ_2-structure. Without loss of generality, we may assume that $\mathcal{A}$ is connected (otherwise we handle each connected component separately). Furthermore, suppose that φ is a homomorphism from $\mathcal{A}$ to $\hat{\mathcal{T}}$ taking the element $u \in |\mathcal{A}|$ to the element $0 \in |\hat{\mathcal{T}}|$: this effectively determines which side of the bipartition of $\mathcal{T}$ every element of $|\mathcal{A}|$ must map to under φ (if we can determine, in polynomial-time, whether such a homomorphism exists then this will clearly suffice to yield the result).

We shall build a set of clauses, C, corresponding to $\mathcal{A}$, each clause consisting of 2 Boolean literals. The underlying set of Boolean variables is $\{X_v : v \in |\mathcal{A}|\}$. For every instantiation $g^{\mathcal{A}}(v, w)$, where $v, w \in |\mathcal{A}|$, include clauses in C as follows.

(*a*) If $\varphi(v) \in \{0,2\}$ then include the clauses $\{\neg X_v, X_w\}$ and $\{X_v, \neg X_w\}$ in C.

(*b*) If $\varphi(v) \in \{1,3\}$ then include the clauses $\{\neg X_v, \neg X_w\}$ and $\{X_v, X_w\}$ in C.

Suppose that ψ is a satisfying truth assignment for the resulting set of clauses C. Suppose further that v is such that $\varphi(v) \in \{0,2\}$ (resp. $\varphi(v) \in \{1,3\}$): if $\psi(X_v)$ is $true$ then we interpret this as dictating that $\varphi(v) = 0$ (resp. $\varphi(v) = 1$); and if $\psi(X_v)$ is $false$ then we interpret this as dictating that $\varphi(v) = 2$ (resp. $\varphi(v) = 3$). If $g^{\mathcal{A}}(v,w)$ holds then the fact that ψ is a satisfying truth assignment ensures that $g^{\hat{\mathcal{T}}}(\varphi(v), \varphi(w))$ holds. Extending C by similarly including clauses for every instantiation of the form $f^{\mathcal{A}}(v,w)$ yields that if the resulting set of clauses C is satisfiable then there is homomorphism from $\mathcal{A}$ to $\hat{\mathcal{T}}$. Conversely, if there is a homomorphism from $\mathcal{A}$ to $\hat{\mathcal{T}}$ then this homomorphism can be translated, using our interpretation above, into a satisfying truth assignment for C. Thus, the problem *Hom-Alg*$_2(\hat{\mathcal{T}})$ can be reduced, in polynomial-time, to the problem of whether a collection of clauses of size 2 is satisfiable, which is well-known to be solvable in polynomial-time (see, for example, [16]). The result follows. □

6 Reversing problems

We have seen so far that even restrictions of the problem $\Delta(\lambda_2)$ can be as hard to resolve as the (very general) problem $\Delta(\sigma : \sigma$ relational$)$. In this section, we prove a dichotomy result for the class of non-uniform constraint satisfaction problems with template a pair of unary functions one of which is the reverse of the other.

Recall that the digraph of a unary function, f, say, consists of a directed cycle with pendant in-trees. Suppose that the unary function g is such that its digraph is obtained from that of f by reversing the directions of the edges of the directed cycle and leaving all other edges as they are. Then we say that g is the *reverse* of f. If $\mathcal{T}$ is a λ_2-structure for which $f^{\mathcal{T}}$ is the reverse of $g^{\mathcal{T}}$ then the problem *Hom-Alg*$_2(\mathcal{T})$ is called a *reversing problem*.

When the unary functions f and g are the reverse of one another, for any element u on the directed cycle in the digraph of f (or g), define the *height* of u to be the length of the longest path ending at u but so that no element on this path apart from u lies on the directed cycle (there are similarities between the height of u and the values $\delta_f(u)$ and $\delta_g(u)$ from the proof of Proposition 9).

Theorem 17 *The problem CSP($\mathcal{T}$) can be solved in polynomial-time when $\mathcal{T}$ is a structure over a signature containing only constant symbols and unary function symbols, so long as:*

(*a*) *every unary function of $\mathcal{T}$ has a digraph that is a disjoint union of directed cycles;*

(*b*) *there is one unary function symbol in the signature; or*

(*c*) *there are two unary function symbols, f and g, in the signature, and the functions $f^{\mathcal{T}}$ and $g^{\mathcal{T}}$ are the reverse of one another such that the corresponding directed cycles, (t_0, t_1, t_2, t_3) and (t_3, t_2, t_1, t_0), have length 4 and either:*

– *the height of each t_i is at most 1; or*

– *t_0 and t_2 have height at most 2 and t_1 and t_3 have height 0.*

Proof (*a*) Let $\mathcal{T}$ be over the signature σ consisting of a number of unary function symbols and constant symbols, where every unary function has a digraph that is a disjoint union of directed cycles. Let $\mathcal{A}$ be a σ-structure and fix some element $u \in |\mathcal{A}|$. Suppose that there is a homomorphism $\varphi : \mathcal{A} \to \mathcal{T}$ taking u to t, for some $t \in |\mathcal{T}|$. This fixes the value (under φ) of each element in the connected component of the digraph $\mathcal{A}$ containing u; and it is straightforward to check whether the resulting map is indeed a homomorphism (of this connected component of the digraph $\mathcal{A}$: how we consider $\mathcal{A}$ as a digraph, and the notion of connected component, should be clear given the definitions prior to Lemma 8). We do this for each element $t \in |\mathcal{T}|$ and for each connected component of $\mathcal{A}$ to obtain our polynomial-time algorithm.

(*b*) Let $\mathcal{T}$ and $\mathcal{A}$ be over the signature σ consisting of one function symbol and a number of constant symbols. The σ-structures $\mathcal{A}$ and $\mathcal{T}$ are both disjoint unions of directed cycles with pendant in-trees (with possibly additional constants). We take each directed cycle with pendant in-trees in $\mathcal{A}$ in turn. Fix some element $u \in |\mathcal{A}|$ which lies on the directed cycle of some directed cycle with pendant in-trees $\mathcal{B}$. Suppose that there is a homomorphism $\varphi : \mathcal{B} \to \mathcal{T}$ taking u to t, for some $t \in |\mathcal{T}|$ (and so t must lie on the cycle of some directed cycle with pendant in-trees in $\mathcal{T}$). This fixes the value of each element in the directed cycle of $\mathcal{B}$. Let us assume that such a partial homomorphism exists. If there are no constants lying on the pendant in-trees in $\mathcal{B}$ then it is trivial to determine whether φ can extended to a homomorphism of $\mathcal{B}$ to $\mathcal{T}$ (see the proof of Proposition 2).

Suppose that $c^{\mathcal{B}}$ is a constant lying on some pendant in-tree in $\mathcal{B}$ such that in the path from the constant to the root of the in-tree (that is, the element on both the in-tree and the directed cycle) there are no other constants. The map φ must take $c^{\mathcal{B}}$ to $c^{\mathcal{T}}$ and whether this is possible can easily be verified. Moreover, if this is possible then the image of every element of $|\mathcal{B}|$ on the path in $\mathcal{B}$ from $c^{\mathcal{B}}$ to the root of the pendant in-tree is fixed. If there are no other constants lying on some pendant in-tree in $\mathcal{B}$ then the map φ can trivially be extended to a homomorphism of $\mathcal{B}$ to $\mathcal{T}$.

Let us suppose that $d^{\mathcal{B}}$ is some other constant lying on some pendant in-tree in $\mathcal{B}$ such that in the path p from $d^{\mathcal{B}}$ to an element of $|\mathcal{B}|$ for which the value under φ has already been set, there are no other constants. The map φ must take $d^{\mathcal{B}}$ to $d^{\mathcal{T}}$ and whether this is possible can easily be verified. Moreover, if this is possible then the image of every element of $|\mathcal{B}|$ on the path p in $\mathcal{B}$ is fixed. We continue similarly with other constants in $\mathcal{B}$ until either we obtain a homomorphism of $\mathcal{B}$ to $\mathcal{T}$ or we show that one cannot exist. We repeat the above for each element $t \in |\mathcal{T}|$ and for each directed cycle with pendant in-trees of $\mathcal{A}$ to obtain our polynomial-time algorithm.

(c) Let $\mathcal{T}$ be over the signature σ consisting of the unary function symbols f and g and a number of constant symbols. Suppose further that $f^{\mathcal{T}}$ and $g^{\mathcal{T}}$ are the reverse of one another where the corresponding directed cycles are (t_0, t_1, t_2, t_3) and (t_3, t_2, t_1, t_0), respectively, and the height of each t_i is at most 1. Let $\mathcal{A}$ be some σ-structure and let $\mathcal{B}$ be the connected component of the digraph $\mathcal{A}$ containing some fixed element $u \in |\mathcal{A}|$ which lies on a cycle in the digraph $\mathcal{A}$ ($\mathcal{B}$ may also have some constants). Suppose that there is a homomorphism $\varphi : \mathcal{A} \to \mathcal{T}$ taking u to t_i, for some $i \in \{0, 1, 2, 3\}$. Either such a φ can easily be verified not to exist or the values under φ of the elements of $|\mathcal{B}|$ with non-zero in-degree are fixed. Let us assume that such a partial homomorphism exists.

Take elements $v, w \in |\mathcal{B}|$ of in-degree 0. As to whether the partial homomorphism φ can be extended to $|\mathcal{B}| \cup \{v\}$ is independent from whether the partial homomorphism φ can be extended to $|\mathcal{B}| \cup \{w\}$, since the values $\varphi(f^{\mathcal{B}}(v))$, $\varphi(g^{\mathcal{B}}(v))$, $\varphi(f^{\mathcal{B}}(w))$ and $\varphi(g^{\mathcal{B}}(w))$ are fixed. Hence, we can efficiently verify whether the homomorphism φ exists. We repeat the above for each element t_i, for $i = 0, 1, 2, 3$ (with u fixed), and for each connected component of $\mathcal{A}$ (with suitably chosen fixed element) to obtain our polynomial-time algorithm.

Now suppose that instead of every t_i having height at most 1 in $\mathcal{T}$, for $i = 0, 1, 2, 3$, t_0 and t_2 have height at most 2 and t_1 and t_3 have height 0. Let $\mathcal{A}$ be some σ-structure, and let $\mathcal{B}$ be the connected component of the digraph $\mathcal{A}$ containing some fixed element $u \in |\mathcal{A}|$ which lies on a cycle in the digraph $\mathcal{A}$ ($\mathcal{B}$ may also have some constants). Suppose that there is a homomorphism $\varphi : \mathcal{A} \to \mathcal{T}$ taking u to t_i, for some $i \in \{0, 1, 2, 3\}$. Either such a φ can easily be verified not to exist or the values under φ of many of the elements of $|\mathcal{B}|$ are effectively fixed (using simple reasoning, like if $\varphi(v) = t_1$ and $f^{\mathcal{B}}(w) = v$ then $\varphi(w) = t_0$). Let us assume that such a partial homomorphism exists.

Let $F \subseteq |\mathcal{B}|$ be the set of elements whose values under φ are fixed. Note that given $v, v', v'' \in |\mathcal{B}| \setminus F$, it cannot be the case that v, v' and v'' are all distinct and (v, v') and (v', v'') are edges in the digraph $\mathcal{B}$ (as this would mean that the value $\varphi(v'')$ is effectively fixed). So, partition $|\mathcal{B}| \setminus F$ into the disjoint sets: L_1, consisting of those elements v for which either $f^{\mathcal{B}}(v)$ or $g^{\mathcal{B}}(v)$ is in F; and L_2, consisting of those elements v for which neither $f^{\mathcal{B}}(v)$ nor $g^{\mathcal{B}}(v)$ is in F.

Suppose that $v \in L_1$. Note that either $f^{\mathcal{B}}(v) \notin F$ or $g^{\mathcal{B}}(v) \notin F$ (as otherwise the value of $\varphi(v)$ would effectively be fixed). Suppose, for example, that $f^{\mathcal{B}}(v) \in F$. Then it must be the case that $\varphi(f^{\mathcal{B}}(v)) \in \{t_0, t_2\}$. If $\varphi(f^{\mathcal{B}}(v)) = t_0$ then $\varphi(v)$ must either be t_3 or in the pendant in-tree with root t_0 (where there is an edge $(\varphi(v), t_0)$ in the digraph $\mathcal{T}$); and if $\varphi(f^{\mathcal{B}}(v)) = t_2$ then $\varphi(v)$ must either be t_1 or in the pendant in-tree with root t_2 (where there is an edge $(\varphi(v), t_2)$ in the digraph $\mathcal{T}$). In the former case, label v with the set $\{t_3, p_0\}$; and in the latter case, with the set $\{t_1, p_2\}$. Do likewise for every element of L_1: the label sets are drawn from $\{t_3, p_0\}$, $\{t_1, p_2\}$, $\{t_3, p_2\}$ and $\{t_1, p_0\}$. Note that this process might even enable us to deduce that the partial homomorphism φ cannot be extended to a full homomorphism (assume in the following that this is not the case).

As $\mathcal{B}$ must necessarily be bipartite if there is a homomorphism to $\mathcal{T}$, we may assume that there are no functional relationships between the elements of L_1. Hence, we may assume that in the digraph $\mathcal{B}$, there is an edge from every element of L_1 to an element of L_2. In the same vein, we may assume that there are no functional relationships between elements of L_2; and so both $f^{\mathcal{B}}(v)$ and $g^{\mathcal{B}}(v)$ are in L_1, for every element $v \in L_2$.

We shall now construct a conjunction π of clauses of at most 2 literals (over a set of boolean variables) in such a way that π is satisfiable if, and only if, the partial homomorphism φ can be extended to a full homomorphism from $\mathcal{B}$ to $\mathcal{T}$. Suppose that the element $v \in L_1$ is labelled with the set $\{t_3, p_0\}$, for example. Then the clauses

$$\{t_3^v, p_0^v\} \text{ and } \{\neg t_3^v, \neg p_0^v\}$$

are clauses in the conjunction; and we add clauses likewise for every labelling of every element of L_1. Note that any truth assignment satisfying both

$\{t_3^v, p_0^v\}$ and $\{\neg t_3^v, \neg p_0^v\}$ makes exactly one of t_3^v and p_0^v *true*; which we will interpret as meaning that element v is mapped to t_3 or an element in the pendant in-tree with root t_0 (where there is an edge $(\varphi(v), t_0)$).

Let $v \in L_1$ be such that, for example, v is labelled with the set $\{t_3, p_0\}$ and $g^{\mathcal{B}}(v) = w$, where $w \in L_2$. Include in π the clauses

$$\{\neg t_3^v, t_2^w\}, \{\neg p_0^v, t_0^w\}, \{t_2^w, t_0^w\} \text{ and } \{\neg t_2^w, \neg t_0^w\}$$

to reflect the facts that: if v maps to t_3 then w must map to t_2; if v maps to an element in the pendant in-tree with root t_0 (where there is an edge $(\varphi(v), t_0)$) then w must map to t_0; and v must map to either t_0 or t_2. Do likewise for every edge in the digraph $\mathcal{B}$ from an element of L_1 to an element of L_2.

Let $v \in L_2$ be such that, for example, $g^{\mathcal{B}}(v) = w$, where $w \in L_1$ and w is labelled with the set $\{t_3, p_0\}$. Include in π the clauses

$$\{\neg t_3^w, t_0^v\}, \{\neg p_0^w, q_0^v\}, \{t_0^v, q_0^v\} \text{ and } \{\neg t_0^v, \neg q_0^v\}$$

to reflect the facts that: if w maps to t_3 then v must map to t_0; if w maps to an element in the pendant in-tree with root t_0 (where there is an edge $(\varphi(v), t_0)$) then v must map to an element in the pendant in-tree with root t_0 (where there is a path of length 2 to the root t_0); and v must map to either t_0 or an element in the pendant in-tree with root t_0 (where there is a path of length 2 to the root t_0). Do likewise for every edge in the digraph $\mathcal{B}$ from an element of L_2 to an element of L_1.

Finally, if the pendant in-tree in $\mathcal{T}$ with root t_0 (resp. t_2) has height 0 then assign every boolean variable of the form p_0^v and q_0^v (resp. p_2^v and q_2^v) the value $false$; and if the pendant in-tree in $\mathcal{T}$ with root t_0 (resp. t_2) has height 1 then assign every boolean variable of the form q_0^v (resp. q_2^v) the value $false$ (this restricts the maps allowed). Denote the resulting conjunction by π.

Let s_0 be an element in $\mathcal{T}$ in the pendant in-tree with root t_0 for which (s_0, t_0) is an edge in the digraph $\mathcal{T}$ (if such an element exists); and let r_0 be an element in $\mathcal{T}$ in the pendant in-tree with root t_0 for which (r_0, s_0) is an edge in the digraph $\mathcal{T}$ (if such an element exists). Nodes s_2 and r_2 are defined likewise.

Suppose that there is a satisfying truth assignment of π. Interpreting this truth assignment as described above, except that if p_0^v (p_2^v, q_0^v, q_2^v) is set at *true* under this assignment then map the element v to the element s_0 (resp. s_2, r_0, r_2), yields an extension of our partial homomorphism φ which obeys all functional dependencies in $\mathcal{B}$, *i.e.*, a homomorphism from $\mathcal{B}$ to $\mathcal{T}$.

Conversely, suppose that there is an extension of our partial homomorphism φ to a homomorphism from $\mathcal{B}$ to $\mathcal{T}$. Then there exists an extension of our partial homomorphism φ to a homomorphism φ' from $\mathcal{B}$ to $\mathcal{T}$ so that all elements of $|\mathcal{B}| \setminus F$ that are mapped to an element of $|\mathcal{T}| \setminus \{t_0, t_1, t_2, t_3\}$ are mapped to one of the elements of $\{s_0, s_2, r_0, r_2\}$ (this is because the constants in $\mathcal{B}$ and $\mathcal{T}$ have to real role to play as to whether an extension of φ exists or not). Consequently, such a homomorphism φ' immediately translates into a satisfying truth assignment of π. As the satisfiability problem for clauses of size at most 2 is solvable in polynomial-time (see, for example, [16]), as to whether there is a homomorphism from $\mathcal{B}$ to $\mathcal{T}$, taking u to t_i, can be decided in polynomial-time. We repeat the above for each element t_i, for $i = 0, 1, 2, 3$ (with u fixed), and for each connected component of $\mathcal{A}$ (with suitably chosen fixed element) to obtain our polynomial-time algorithm. $\square$

It turns out that the reversing problems not covered by Theorem 17 are **NP**-complete (via polynomial-time reductions).

Theorem 18 *The problem Hom-Alg*$_2(\mathcal{T})$*, where* $f^{\mathcal{T}}$ *and* $g^{\mathcal{T}}$ *are the reverse of one another, is* **NP***-complete if one of the following conditions holds*:

- (a) *the directed cycles* $(t_0, t_1, \ldots, t_{l-1})$ *and* $(t_{l-1}, t_{l-2}, \ldots, t_0)$ *corresponding to* $f^{\mathcal{T}}$ *and* $g^{\mathcal{T}}$*, respectively, are such that* t_0 *has height* 1 *and either* $l = 3$ *or* $l \geq 5$;
- (b) *the directed cycles* (t_0, t_1, t_2, t_3) *and* (t_3, t_2, t_1, t_0) *corresponding to* $f^{\mathcal{T}}$ *and* $g^{\mathcal{T}}$*, respectively, are such that* t_0 *has height* 3; *or*
- (c) *the directed cycles* (t_0, t_1, t_2, t_3) *and* (t_3, t_2, t_1, t_0) *corresponding to* $f^{\mathcal{T}}$ *and* $g^{\mathcal{T}}$*, respectively, are such that* t_0 *has height* 2 *and* t_1 *has height at least* 1.

Proof (a) Consider the retract problem for an undirected cycle of length l, RETCYC(l) for short, where: an instance is an undirected graph $\mathcal{G}$ and a subgraph $\mathcal{H}$, a cycle of length l; and a yes-instance is an instance for which there is a homomorphism φ from $\mathcal{G}$ to $\mathcal{H}$ such that $\varphi(u) = u$, for every vertex of the cycle C. This problem is known to be **NP**-complete when $l = 3$ or when $l \geq 5$ (the case for cycles of odd length is a special case of the proof of the main result in [11]; and the case for cycles of even length at least 6 can be found in [5]). We shall describe a polynomial-time reduction from RETCYC(l) to *Hom-Alg*$_2(\mathcal{T})$.

Let $(\mathcal{G}, \mathcal{H})$ be an instance of RETCYC(l). We shall construct a λ_2-structure $\mathcal{F}$ corresponding to $(\mathcal{G}, \mathcal{H})$. Actually, our structure $\mathcal{F}$ will be only partially defined; but by proceeding as we did in the proof of Theorem 13, we can obtain a fully defined λ_2-structure for which there is a homomorphism to $\mathcal{T}$ if, and only if, there is a homomorphism from the partially defined structure to $\mathcal{T}$. Take a set S of elements isomorphic to the vertex set of $\mathcal{G}$. For every element of $s \in S$, introduce l additional elements $u_0^s, u_1^s, \ldots, u_{l-1}^s$ and define $f(u_i^s) = u_{(i+1)\text{mod } l}^s$, for each $l \in \{0, 1, \ldots, l-1\}$. Furthermore, define $f(s) = u_0^s$ (all elements so introduced are distinct). Note that this has the effect that in any homomorphism from the eventual structure $\mathcal{F}$ to $\mathcal{T}$, every element so far introduced must map to t_i or to some element t for which $f^{\mathcal{T}}(t) = g^{\mathcal{T}}(t) = t_i$, for some $i \in \{0, 1, \ldots, l-1\}$.

Let u and v be distinct elements of S. We now detail three constructions to be subsequently used in order to obtain our eventual structure $\mathcal{F}$. Note that in each of these constructions, we introduce additional elements into the domain of $\mathcal{F}$. Using the above construction, we may clearly assume that in any homomorphism from the eventual structure $\mathcal{F}$ to $\mathcal{T}$, every such element introduced must map to t_i or to some element t for which $f^{\mathcal{T}}(t) = g^{\mathcal{T}}(t) = t_i$, for some $i \in \{0, 1, \ldots, l-1\}$.

<u>Construction 1</u> By *add the α-edge* (u, v) we mean introduce additional elements u' and v' to $|\mathcal{F}|$ and define that $g^{\mathcal{F}}(u') = u$, $g^{\mathcal{F}}(v') = v$ and $f^{\mathcal{F}}(u') = v'$. This ensures that in any homomorphism from the eventual structure $\mathcal{F}$ to $\mathcal{T}$, u and v must map to some t_i and t_j, respectively, where: if t_i has height 0 then $j = (i+1) \bmod l$ is the only possibility; and if t_i has non-zero height then both $j = (i-1) \bmod l$ and $j = (i+1) \bmod l$ are possibilities. Furthermore, any i and j are viable (subject to the above constraints).

<u>Construction 2</u> Consider the directed cycle $(t_0, t_1, \ldots, t_{l-1})$ in $\mathcal{T}$ (corresponding to $f^{\mathcal{T}}$). Suppose that there is a path of length r in this cycle from t_i such that some element t_j on this path has height 1. If this is the case then we say that t_i can *r-see* an element of height 1. Let r be the smallest non-negative integer with the property that every t_i can r-see an element of height 1 but, if $r \neq 0$, there is some element t_j which cannot $(r-1)$-see an element of height 1. If $r > 0$ then fix j_0 as some index for which t_{j_0} cannot $(r-1)$-see an element of height 1.

By *add the β-edge* (u, v) we mean introduce additional elements so that there is a path of $(r+1)$ α-edges from u to some (new) element w, followed by a path of r g-edges from w to v (a g-edge is just a pair of elements u' and v' for which $g^{\mathcal{F}}(u') = v'$, and similarly for an f-edge: if $r = 0$ then $u = w$). This ensures that in any homomorphism from the eventual structure $\mathcal{F}$ to $\mathcal{T}$, it is possible for u to map to any t_i and for v to map to $t_{(i-1)\text{mod } l}$ and

$t_{(i+1) \bmod l}$: furthermore, these are the only two possibilities if $i = j_0$. Note that any i and j are viable (subject to the above constraints).

Construction 3 By *add the γ-edge* (u, v) we mean introduce additional sets of elements Γ_p so that for every $p \in \{0, 1, \ldots, l-1\}$: there is an f-edge from a new element u_0^p to u; there is a path of p f-edges from u to some (new) element w^p; followed by a β-edge (w^p, x^p), where x^p is a new element; followed by a path of p g-edges from x^p to v; and there is an f-edge from a new element v_0^p to v (the different sets Γ_p of new elements introduced are disjoint). It is certainly the case that there is a homomorphism from any γ-edge to $\mathcal{T}$; and our construction ensures that in any homomorphism from the eventual structure $\mathcal{F}$ to $\mathcal{T}$, u and v must map to some elements t_i and t_j. However, we can say more.

Suppose that u and v map to elements t_i and t_j, respectively, in some eventual homomorphism φ from the eventual structure $\mathcal{F}$ to $\mathcal{T}$. Consider the distinguished element t_{j_0}. There exists a value p so that the element $w^p \in \Gamma_p$ must map to t_{j_0} via φ; and consequently, from above, x^p must map to either $t_{(j_0-1) \bmod l}$ or $t_{(j_0+1) \bmod l}$, with both these values possible. Thus, v must map to either $t_{(i-1) \bmod l}$ or $t_{(i+1) \bmod l}$, with both these values possible. That is, either $j = (i-1) \bmod l$ or $j = (i+1) \bmod l$; and any i and j are viable (subject to the above constraints).

We now extend $\mathcal{F}$ using these constructions (recall that $\mathcal{F}$'s current state is as it was prior to the descriptions of our constructions). For every pair of elements $u, v \in S$ for which (u, v) is an undirected edge of $\mathcal{G} \setminus \mathcal{H}$, add the γ-edge (u, v) to $\mathcal{F}$ (to be more specific, add either the γ-edge (u, v) or the γ-edge (v, u)). The fact that there is a homomorphism of $\mathcal{G}$ to $\mathcal{H}$ fixing each vertex of $\mathcal{H}$ if, and only if, there is a homomorphism from $\mathcal{F}$ to $\mathcal{T}$, is immediate from our constructions; and this, allied with the fact that $\mathcal{F}$ can be constructed from $(\mathcal{G}, \mathcal{H})$ in polynomial-time, yields the result.

(b) Consider some instance $\mathcal{A}$ of the problem 1-in-3-SAT; that is, a collection $C_1, C_2, \ldots, C_m$ of subsets of $\{X_1, X_2, \ldots, X_n\}$ of size 3. We shall describe an instance $\mathcal{F}$ of *Hom-Alg*$_2(\mathcal{T})$, where $\mathcal{T}$ is as in the statement of the theorem.

There are elements c_0, c_1, c_2 and c_3 in $|\mathcal{F}|$ for which:

$$f^{\mathcal{F}}(c_0) = c_1; f^{\mathcal{F}}(c_1) = c_2; f^{\mathcal{F}}(c_2) = c_3; \text{ and } f^{\mathcal{F}}(c_3) = c_0.$$

There are elements d_0, d_1 and d_2 for which:

$$f^{\mathcal{F}}(d_0) = d_1; f^{\mathcal{F}}(d_1) = d_2; \text{ and } f^{\mathcal{F}}(d_2) = g^{\mathcal{F}}(d_2) = c_0.$$

In particular, we may assume that any eventual homomorphism from $\mathcal{F}$ to $\mathcal{T}$ must map c_0 to t_0.

For every $i \in \{1, 2, \ldots, m\}$, there is an element C_i in $|\mathcal{F}|$ as well as other elements so that $f^{\mathcal{F}} g^{\mathcal{F}} f^{\mathcal{F}}(C_i) = c_0$ (all nodes so introduced are distinct). This has the effect that in any eventual homomorphism from $\mathcal{F}$ to $\mathcal{T}$, the element C_i must be mapped to:

(1) an element not in the directed cycle in $\mathcal{T}$ but for which in the digraph $\mathcal{T}$ there is a path of length 3 to t_0 where no element on this path is in the directed cycle apart from t_0;

(2) an element not in the directed cycle in $\mathcal{T}$ but for which in the digraph $\mathcal{T}$ there is an edge to t_0; or

(3) the element t_3.

For every $j \in \{1, 2, \ldots, n\}$, introduce an element X_j to $|\mathcal{F}|$. Fix an ordering on the Boolean variables in each clause in $\mathcal{A}$. If clause $C_i = \{X_i^1, X_i^2, X_i^3\}$ then introduce additional elements so that:

$$f^{\mathcal{F}} f^{\mathcal{F}} g^{\mathcal{F}}(C_i) = X_i^1; f^{\mathcal{F}} f^{\mathcal{F}} f^{\mathcal{F}} f^{\mathcal{F}} f^{\mathcal{F}}(C_i) = X_i^2; \text{ and } g^{\mathcal{F}} f^{\mathcal{F}} g^{\mathcal{F}} f^{\mathcal{F}} g^{\mathcal{F}}(C_i) = X_i^3$$

(again, all nodes so introduced are distinct). Note that the resulting λ_2-structure $\mathcal{F}$ is partially defined. This does not create any difficulties as, again, by proceeding as we did in the proof of Theorem 13, we can obtain a fully defined λ_2-structure for which there is a homomorphism to $\mathcal{T}$ if, and only if, there is a homomorphism from the partially defined structure to $\mathcal{T}$.

Let φ be a homomorphism from $\mathcal{F}$ to $\mathcal{T}$; and let $C_i = \{X_i^1, X_i^2, X_i^3\}$ be a clause of $\mathcal{A}$, for some $i \in \{1, 2, \ldots, m\}$. From above, there are three different possibilities for $\varphi(C_i)$ (here, C_i refers to the element C_i of $|\mathcal{F}|$ and not the clause C_i of $\mathcal{A}$). In case (1), $(\varphi(X_i^1), \varphi(X_i^2), \varphi(X_i^3)) = (t_0, t_2, t_0)$; in case (2), $(\varphi(X_i^1), \varphi(X_i^2), \varphi(X_i^3)) = (t_2, t_0, t_0)$; and in case (3), $(\varphi(X_i^1), \varphi(X_i^2), \varphi(X_i^3)) = (t_0, t_0, t_2)$. Interpreting t_0 as *true* and t_2 as *false* yields a satisfying truth assignment for $\mathcal{A}$.

Conversely, if there is a satisfying truth assignment for $\mathcal{A}$ then applying reasoning similar to that above yields a homomorphism from $\mathcal{F}$ to $\mathcal{T}$. As $\mathcal{F}$ can be constructed in polynomial-time, the result follows.

(c) Consider some instance $\mathcal{A}$ of the problem 1-in-3-SAT; that is, a collection $C_1, C_2, \ldots, C_m$ of subsets of $\{X_1, X_2, \ldots, X_n\}$ of size 3. We shall describe an instance $\mathcal{F}$ of *Hom-Alg*$_2(\mathcal{T})$, where $\mathcal{T}$ is as in the statement of the theorem.

There are elements c_0, c_1, c_2 and c_3 for which

$$f^{\mathcal{F}}(c_0) = c_1; f^{\mathcal{F}}(c_1) = c_2; f^{\mathcal{F}}(c_2) = c_3; \text{ and } f^{\mathcal{F}}(c_3) = c_0.$$

There are elements d_0, d_1 and d_2 for which

$$f^{\mathcal{F}}(d_0) = d_1; f^{\mathcal{F}}(d_1) = g^{\mathcal{F}}(d_1) = c_0; \text{ and } f^{\mathcal{F}}(d_2) = g^{\mathcal{F}}(d_2) = c_1.$$

In particular, we may assume that any eventual homomorphism from $\mathcal{F}$ to $\mathcal{T}$ must map c_0 to t_0.

For every $i \in \{1, 2, \ldots, m\}$, there is an element C_i in $|\mathcal{F}|$ as well as other elements so that $g^{\mathcal{F}} f^{\mathcal{F}}(C_i) = c_0$ (all nodes so introduced are distinct). This has the effect that in any eventual homomorphism from $\mathcal{F}$ to $\mathcal{T}$, the element C_i must be mapped to:

(1) an element not in the directed cycle in $\mathcal{T}$ but for which in the digraph $\mathcal{T}$ there is a path of length 2 to t_0 where no element on this path is in the directed cycle apart from t_0;

(2) an element not in the directed cycle in $\mathcal{T}$ but for which in the digraph $\mathcal{T}$ there is an edge to t_1; or

(3) the element t_0.

For every $j \in \{1, 2, \ldots, n\}$, introduce an element X_j into $|\mathcal{F}|$. Fix an ordering on the Boolean variables in each clause in $\mathcal{A}$. If clause $C_i = \{X_i^1, X_i^2, X_i^3\}$ then introduce additional elements so that:

$$g^{\mathcal{F}} g^{\mathcal{F}}(C_i) = X_i^1; f^{\mathcal{F}} f^{\mathcal{F}} f^{\mathcal{F}} f^{\mathcal{F}}(C_i) = X_i^2; \text{ and } f^{\mathcal{F}} g^{\mathcal{F}}(C_i) = X_i^3$$

(again, all nodes so introduced are distinct). Note that the resulting λ_2-structure $\mathcal{F}$ is partially defined. This does not create any difficulties as, again, by proceeding as we did in the proof of Theorem 13, we can obtain a fully defined λ_2-structure for which there is a homomorphism to $\mathcal{T}$ if, and only if, there is a homomorphism from the partially defined structure to $\mathcal{T}$.

Let φ be a homomorphism from $\mathcal{F}$ to $\mathcal{T}$; and let $C_i = \{X_i^1, X_i^2, X_i^3\}$ be a clause of $\mathcal{A}$, for some $i \in \{1, 2, \ldots, m\}$. From above, there are three different possibilities for $\varphi(C_i)$. In case (1), $(\varphi(X_i^1), \varphi(X_i^2), \varphi(X_i^3)) = (t_0, t_2, t_0)$; in case (2), $(\varphi(X_i^1), \varphi(X_i^2), \varphi(X_i^3)) = (t_0, t_0, t_2)$; and in case (3), $(\varphi(X_i^1), \varphi(X_i^2), \varphi(X_i^3)) = (t_2, t_0, t_0)$. Interpreting t_0 as *true* and t_2 as $false$ yields a satisfying truth assignment for $\mathcal{A}$.

Conversely, if there is a satisfying truth assignment for $\mathcal{A}$ then applying reasoning similar to that above yields a homomorphism from $\mathcal{F}$ to $\mathcal{T}$. As $\mathcal{F}$ can be constructed in polynomial-time, the result follows. □

An immediate corollary from Theorems 17 and 18 is the following dichotomy result.

Corollary 19 *Every reversing problem is either solvable in polynomial-time or* **NP***-complete.*

Proof Consider some reversing problem $Hom\text{-}Alg_2(\mathcal{T})$. Let l denote the length of the directed cycles corresponding to the unary functions $f^{\mathcal{T}}$ and $g^{\mathcal{T}}$.

If $l = 1$ or $l = 2$ then $f^{\mathcal{T}}$ and $g^{\mathcal{T}}$ are identical; and part (b) of Theorem 17 yields that our reversing problem is solvable in polynomial-time.

If $l = 3$ or $l \geq 5$ and, further, every element on the directed cycles corresponding to $f^{\mathcal{T}}$ and $g^{\mathcal{T}}$ has height 0 then part (a) of Theorem 17 yields that our reversing problem is solvable in polynomial-time.

If $l = 3$ or $l \geq 5$ and, further, some element on the directed cycles corresponding to $f^{\mathcal{T}}$ and $g^{\mathcal{T}}$ has height greater than 0 then part (a) of Theorem 18 yields that our reversing problem is **NP**-complete.

If $l = 4$ then part (c) of Theorem 17 and parts (b) and (c) of Theorem 18 cover all other cases. $\square$

7 Adding lists

Consider some structure $\mathcal{T}$. A *list* in $\mathcal{T}$ is a unary relation over $|\mathcal{T}|$ (describing a set of elements of $|\mathcal{T}|$). In this section we consider the addition of a list or lists to templates. We shall show that some of our results can be extended when the templates are augmented with lists.

We begin with some definitions. Let $\mathcal{H}$ be some digraph. A *choice majority function* χ is a ternary function on $|\mathcal{H}|$ with the property that if $\chi(x, y, z) = w$ and $\chi(x', y', z') = w'$ then:

(i) if (x, x'), (y, y') and (z, z') are edges of $\mathcal{H}$ then (w, w') is too;

(ii) $w \in \{x, y, z\}$; and

(iii) if two out of x, y and z are equal to the same value v then $w = v$.

We can define a choice majority function for structures over arbitrary relational signatures in a similar way (we only remark that every one of the structures relations must satisfy condition (i) of the definition). A crucial result in what follows will be Theorem 25 of [9] where it is proven that if $\mathcal{T}$ is a relational structure which admits a choice majority function then the problem CSP$(\mathcal{T})$ is uniformly polynomial-time solvable.

Let $\mathcal{H}$ be a digraph. We say that $\mathcal{H}$ is *levelled* if we can assign a non-negative integer $lev(v)$ to each vertex v so that for every edge (u, v) of $\mathcal{H}$,

$lev(u) = lev(v) + 1$. We call $lev(v)$ the *level* of v. We say that a levelled digraph is *orderly levelled* if every vertex v can be assigned an integer $ord(v)$, the *order*, so that: if u and v are distinct vertices then $ord(u) \neq ord(v)$; and if (u, v) and (u', v') are edges of $\mathcal{H}$ with $lev(u) = lev(u')$ and $ord(u) < ord(u')$ then $ord(v) \leq ord(v')$. We say that a digraph $\mathcal{H}$ is *almost orderly levelled* if there exists an orderly levelled digraph $\mathcal{H}'$ such that $\mathcal{H}$ is obtained from $\mathcal{H}'$ by identifying two vertices u and v in $\mathcal{H}'$, where: u is the unique vertex of $\mathcal{H}'$ with $lev(u) = 0$; and if $lev(v) = m$, with $m > 0$, then v has the smallest order from all vertices of level m.

We begin by considering the problem CSP($\mathcal{T}$), where $\mathcal{T}$ is over the signature consisting of one unary function symbol and an arbitrary number of unary relation symbols; that is, an extension of the scenario studied in Section 3 (recall that in the absence of any unary relation symbols, CSP($\mathcal{T}$) is in $\mathbf{L}$). Actually, the consideration of such a problem CSP($\mathcal{T}$) is really an extension of the scenario in part (b) of Theorem 17, as a constant can be modelled as a unary relation containing exactly one element.

Referring to the proof of Theorem 17, one can easily verify that the proofs of parts (a) and (c) of that theorem hold when the signature contains, in addition, an arbitrary number of unary relation symbols (in brief, the proofs of (a) and (c) are based around a systematic construction of a homomorphism, where the homomorphism is incrementally constructed according to 'deterministic' reasoning, and the presence of lists just rules out some potential homomorphisms). However, the proof of part (b) of Theorem 17 does not hold when there are unary relation symbols present. Nevertheless, it turns out that there is an alternative proof of the analogue to part (b) of Theorem 17 when there are unary relation symbols present.

Theorem 20 *Let $\mathcal{H}$ be a digraph.*

- (a) *If $\mathcal{H}$ is almost orderly levelled then H admits a choice majority function.*
- (b) *If $\mathcal{H}$ is the digraph of a unary function f then each connected component of $\mathcal{H}$ is almost orderly levelled.*

Consequently, if $\mathcal{T}$ is a σ-structure where σ consists of a unary function and an arbitrary number of unary relations then the problem CSP($\mathcal{T}$) can be solved uniformly in polynomial-time.

Proof (a) Let $\mathcal{H}'$ be an orderly levelled digraph such that $\mathcal{H}$ is obtained from $\mathcal{H}'$ by identifying the vertices u and v of $\mathcal{H}'$, where u is the unique vertex of level 0, and where $lev(v) = m > 0$ and v has the smallest order from all vertices of level m. Let us order the set of pairs $\{(lev(v), ord(v)) : v \text{ is a vertex of } \mathcal{H}'\}$ lexicographically. Furthermore, we refer to the sets of vertices whose levels differ by a multiple of m as *layers*.

Define the function $\chi(x, y, z)$ as follows.

(1) If x, y and z are in three different layers then $\chi(x, y, z) = x$.

(2) If exactly two of x, y and z are in different layers then $\chi(x, y, z)$ is the value from this pair of values occurring first in the list x, y, z.

(3) If x, y and z are in the same layer then $\chi(x, y, z)$ is the value w out of $\{x, y, z\}$ for which $(lev(w), ord(w))$ is neither the minimum nor the maximum out of $\{(lev(x), ord(x)), (lev(y), ord(y)), (lev(z), ord(z))\}$.

For any triple of vertices (x, y, z) of $\mathcal{H}'$, define the *type* of (x, y, z) as (1), (2) or (3), according to the definition of χ.

We shall now verify that χ is a choice majority function (clearly, it suffices to verify condition (i) of the definition). Let (x, x'), (y, y') and (z, z') be edges of $\mathcal{H}'$. First, note that the type of (x, y, z) is identical to the type of (x', y', z'); and that if (x, y, z) and (x', y', z') have type (1) or (2) then trivially $(\varphi(x, y, z), \varphi(x', y', z'))$ is an edge of $\mathcal{H}'$. Suppose that both (x, y, z) and (x', y', z') have type (3). Consider, for example, (x, x') and (y, y'). Suppose that $(lev(x), ord(x)) \leq (lev(y), ord(y))$. If $lev(x) < lev(y)$ then trivially $lev(x') < lev(y')$, with $(lev(x'), ord(x')) < (lev(y'), ord(y'))$. If $lev(x) = lev(y)$ and $ord(x) < ord(y)$ then, by definition, $ord(x') \leq ord(y')$, with $(lev(x'), ord(x')) \leq (lev(y'), ord(y'))$. Consequently, $(lev(x), ord(x)) \leq (lev(y), ord(y))$ implies that $(lev(x'), ord(x')) \leq (lev(y'), ord(y'))$, and so $(\chi(x, y, z), \chi(x', y', z'))$ is an edge of $\mathcal{H}'$. Thus, χ is a choice majority function for $\mathcal{H}'$.

Finally, u and v are the two smallest elements of the layer to which they belong, with respect to the lexicographic level/order ordering. Hence, when u and v are identified (to obtain $\mathcal{H}$ from $\mathcal{H}'$) the above argument applies; and we obtain a choice majority function for $\mathcal{H}$.

(b) Assume that $\mathcal{H}$ is connected. Let x be any vertex of $\mathcal{H}$ belonging to the unique cycle in the digraph $\mathcal{H}$, which has length $m \geq 1$, say, and let $y = f(x)$ (note that y may be identical to x). Define the digraph $\mathcal{H}'$ by including a new vertex x', removing the edge (x, y) and including the edge (x', y). So, x' has in-degree 0 and x has out-degree 0 in $\mathcal{H}'$; and $\mathcal{H}'$ is an in-tree with root x. Define $lev(z)$ to be the depth of any vertex z in the in-tree $\mathcal{H}'$ (rooted at x); and define $ord(z)$ to be the order z is visited in a breadth-first search of $\mathcal{H}'$ starting at the root x, taking care to ensure that x' has the least order of all vertices whose level is m, where m is the length of the unique cycle. Thus, $\mathcal{H}'$ is orderly levelled and $\mathcal{H}$ is almost orderly levelled.

Now, consider some instance $\mathcal{F}$ of the problem CSP($\mathcal{T}$), where $\mathcal{T}$ is a σ-structure and σ consists of a unary function f and an arbitrary number of unary relations. Suppose, for the moment, that the digraphs $\mathcal{F}$ and $\mathcal{T}$ are both connected (that is, the digraphs described by $f^{\mathcal{F}}$ and $f^{\mathcal{T}}$ are both connected). Let $\mathcal{F}_1$ and $\mathcal{T}_1$ be the restrictions of $\mathcal{F}$ and $\mathcal{T}$ to the signature $\lambda_1 = \langle f \rangle$, respectively. By (a) and (b), $\mathcal{T}_1$ is almost orderly levelled and $\mathcal{T}_1$ admits a choice majority function. By Theorem 25 of [9], if any digraph $\mathcal{H}$ admits a choice majority function then the problem CSP($\mathcal{H}'$), where $\mathcal{H}'$ is $\mathcal{H}$ augmented with an arbitrary number of lists, can be solved in time polynomial in the size of an instance of the problem and also the size of $\mathcal{H}$. Hence, as to whether there is a homomorphism from $\mathcal{F}_1$ to $\mathcal{T}_1$ can be determined in time polynomial in $|\mathcal{F}_1|$ and $|\mathcal{T}_1|$.

There is a homomorphism from $\mathcal{F}$ to $\mathcal{T}$ if, and only if, there is a homomorphism from every connected component of $\mathcal{F}$ to some connected component of $\mathcal{T}$. Thus, we obtain that the problem CSP($\mathcal{T}$) can be solved in time polynomial in both the size of the instance and the size of $\mathcal{T}$. $\square$

Note that even though we have shown, in Section 3, that $Hom\text{-}Alg_1(\mathcal{T})$ is in **L**, we do not know whether the analogous problem where lists are present is in **L**: all we have been able to show is that it is in **P**.

The fact that the proof of part (c) of Theorem 17 holds in the presence of lists, and the trivial observation that the proof of Theorem 18 holds in the presence of lists, yield the following corollary.

Corollary 21 *Let $\mathcal{T}$ be a structure over the signature consisting of two unary function symbols and an arbitrary number of unary relation symbols, where the two functions of $\mathcal{T}$ are the reverse of one another. The problem CSP($\mathcal{T}$) is either solvable in polynomial-time or* **NP***-complete. That is, every reversing problem with lists is either solvable in polynomial-time or* **NP***-complete.* $\square$

Note that Theorem 20 can be applied to other classes of digraphs. An *unbalanced oriented cycle* is a cycle with each edge oriented in some direction but which is not a levelled digraph. It was shown in [5] that unbalanced oriented cycles are almost orderly levelled digraphs. Thus, Theorem 20 and the result from [9] cited in the proof of Theorem 20 yield the following corollary.

Corollary 22 *Let $\mathcal{T}$ be a structure over the signature consisting of a binary relation symbol and an arbitrary number of unary relation symbols, where the digraph described by the binary relation is an unbalanced oriented cycle. The problem CSP$(\mathcal{T})$ can be solved uniformly in polynomial-time.* $\square$

We end this section by considering constraint satisfaction problems with a template $\mathcal{T}$, over some signature σ, but where any instance of the problem is a σ-structure $\mathcal{A}$ augmented with a list U_u of elements of $|\mathcal{T}|$, for every element of $u \in |\mathcal{A}|$. Such an instance is a yes-instance if there exists a homomorphism $\varphi : \mathcal{A} \to \mathcal{T}$ for which $\varphi(u) \in U_u$, for every $u \in |\mathcal{A}|$. Essentially, the lists attached to $\mathcal{A}$ restrict the set of potential homomorphisms. We denote the above problem by CSP$(\mathcal{T} + \text{all lists})$.

A digraph $\mathcal{H}$ is *N-free* if it does not have vertices x, x', y and y', with $x \neq x'$ and $y \neq y'$, such that (x, y), (x, y') and (x', y') are edges of $\mathcal{H}$.

Theorem 23 *Let $\mathcal{T}$ be over a structure σ which consists of the binary relation symbols $H_1, H_2, \ldots, H_r$, where $r \geq 1$.*

- *(a) If $\mathcal{T}$ admits a choice majority function then the problem CSP$(\mathcal{T} +$ all lists$)$ can be solved in polynomial-time.*
- *(b) If the digraph described by $H_i^{\mathcal{T}}$ is N-free, for all $i \in \{1, 2, \ldots, r\}$, and $\mathcal{T}$ does not admit a choice majority function then the problem CSP$(\mathcal{A} +$ all lists$)$ is* **NP***-complete.*

Furthermore, every digraph associated with a unary function is N-free. Consequently, if $\mathcal{T}'$ is a structure over the signature consisting of the unary function symbols $f_1, f_2, \ldots, f_r$ then the problem CSP$(\mathcal{T}' +$ all lists$)$ is either solvable in polynomial-time or **NP***-complete.*

Proof (a) This follows immediately from Theorem 25 of [9].

(b) Suppose that for each $i \in \{1, 2, \ldots, r\}$, the digraph described by $H_i^{\mathcal{T}}$ is N-free. We shall now define a particular σ-structure $\mathcal{F}$ with attached lists (henceforth, when we refer to $\mathcal{F}$ we shall mean $\mathcal{F}$ with these attached lists). The domain of $\mathcal{F}$ consists of $\{(u, v, w) : u, v, w \in |\mathcal{T}|\}$.

- If $u, v, w \in |\mathcal{T}|$ are all distinct then the list associated with the element (u, v, w) of $|\mathcal{F}|$, that is, $U_{(u,v,w)}$, is defined as $\{u, v, w\}$.
- If $u, v, w \in |\mathcal{T}|$ are such that at least two of $\{u, v, w\}$ are identical, with this value being x, then the list associated with the element (u, v, w) of $|\mathcal{F}|$, that is, $U_{(u,v,w)}$, is defined as $\{x\}$.

- For each $i \in \{1, 2, \ldots, r\}$, if $u, v, w, u', v', w' \in |\mathcal{T}|$ then $((u, v, w), (u', v', w')) \in H_i^{\mathcal{F}}$ if, and only if, $(u, u'), (v, v'), (w, w') \in H_i^{\mathcal{T}}$ (note that if $((u, v, w), (u', v', w')) \in H_i^{\mathcal{F}}$ then, since $H_i^{\mathcal{T}}$ is N-free, the only pairs in $H_i^{\mathcal{T}}$ where the first component is in $\{u, v, w\}$ and the second component is in $\{u', v', w'\}$ are (u, u'), (v, v') and (w, w')).

The construction of $\mathcal{F}$ is such that there is a correspondence between the set of homomorphisms from $\mathcal{F}$ to $\mathcal{T}$ and the set of choice majority functions admitted by $\mathcal{T}$. To see this, note that the list associated with any element $(u, v, w) \in |\mathcal{F}|$ contains the potential images of the element (u, v, w) under some choice majority function (see conditions (ii) and (iii) of the definition of a choice majority function); and the relations $H_1^{\mathcal{F}}, H_2^{\mathcal{F}}, \ldots, H_r^{\mathcal{F}}$ ensure that condition (i) of the definition of a choice majority function is adhered to.

There is a straightforward algorithm to determine whether there is a homomorphism from $\mathcal{F}$ to $\mathcal{T}$; and consequently whether $\mathcal{T}$ admits a choice majority function. We think of $H_1^{\mathcal{F}}, H_2^{\mathcal{F}}, \ldots, H_r^{\mathcal{F}}$ as describing digraphs super-imposed on the same set of vertices: so, let us denote the resulting underlying undirected graph by $\mathcal{G}$. Our algorithm can be described as a repeated application of the procedure below (until it can no longer be applied).

> Take some $(u, v, w) \in |\mathcal{F}|$, where u, v and w are all distinct. If there is a path of (undirected) edges in $\mathcal{G}$ from (u, v, w) to some vertex (v', v', u') (resp. (v', u', v'), (u', v', v')), where $u' \neq v'$ and where each vertex on the path apart from the last is such that all components are distinct, then remove the element w (resp. v, u) from $U_{(u,v,w)}$.

The algorithm is such that if some list $U_{(u,v,w)}$ becomes empty, through repeated applications of the above procedure, then the algorithm outputs that no homomorphism exists; otherwise the output is that a homomorphism does exist.

If our algorithm responds that no homomorphism exists then we must verify whether this is truly the case. There is clearly a sub-tree of $\mathcal{G}$ (not necessarily induced) consisting of three paths from a vertex (u, v, w), where u, v and w are all distinct, to vertices (v_0, v_0, u_0), (v_1, u_1, v_1) and (u_2, v_2, v_2), where $u_i \neq v_i$, for $i = 0, 1, 2$, and where these paths may have vertices in common. Any homomorphism φ from $\mathcal{F}$ to $\mathcal{T}$ must be such that $\varphi((v_0, v_0, u_0)) = v_0$; and so if (u', v', w') is the penultimate vertex on the path from (u, v, w) to (v_0, v_0, u_0) (hence, u', v' and w' are all distinct) then it cannot be the case that $\varphi((u', v', w')) = w'$, as this would imply that either

(v', v_0), (w', u_0) and (w', v_0) are in $H_i^{\mathcal{F}}$, for some $i \in \{1, 2, \ldots, r\}$, or (v_0, v'), (u_0, w') and (v_0, w') are in $H_i^{\mathcal{F}}$, for some $i \in \{1, 2, \ldots, r\}$ (and so contradict the fact that the digraph described by each $H_i^{\mathcal{F}}$ is N-free). Reasoning similarly back up the path to vertex (u, v, w) yields that it cannot be the case that $\varphi((u, v, w)) = w$. Analogous arguments for the paths from (u, v, w) to (v_1, u_1, v_1) and (u_2, v_2, v_2) yield that the algorithm's response is correct.

Conversely, suppose that our algorithm responds that a homomorphism exists. Then, on termination of our algorithm, the list associated with every vertex is non-empty. Let $(u, v, w), (u', v', w') \in |\mathcal{F}|$ be such that u, v and w are all distinct and u', v' and w' are all distinct. Clearly: $u \notin U_{(u,v,w)}$ if, and only if, $u' \notin U_{(u',v',w')}$; $v \notin U_{(u,v,w)}$ if, and only if, $v' \notin U_{(u',v',w')}$; and $w \notin U_{(u,v,w)}$ if, and only if, $w' \notin U_{(u',v',w')}$. Thus, define the function φ from $|\mathcal{F}|$ to $\mathcal{T}$ as follows. If $(u, v, w) \in |\mathcal{F}|$ then define $\varphi((u, v, w))$ as the first element (with respect to the order u, v, w) from $U_{(u,v,w)}$ on termination. This function φ is well-defined and, by construction of $\mathcal{F}$, a homomorphism from $\mathcal{F}$ to $\mathcal{T}$. Hence, our algorithm correctly determines whether there exists a homomorphism from $\mathcal{F}$ to $\mathcal{T}$.

Now suppose that $\mathcal{T}$ does not admit a choice majority function. From above, there is a σ-structure $\mathcal{H}$ corresponding to the undirected tree in $\mathcal{G}$, alluded to above, which consists of three paths from a vertex (w_0, w_1, w_2), where w_0, w_1 and w_2 are all distinct (and were formerly called u, v and w), to vertices $x_0 = (v_0, v_0, u_0)$, $x_1 = (v_1, u_1, v_1)$ and $x_2 = (u_2, v_2, v_2)$, where $u_i \neq v_i$, for $i = 0, 1, 2$. Amend $\mathcal{H}$ so that the list associated with element x_0 (resp. x_1, x_2) is $\{u_0, v_0\}$ (resp. $\{u_1, v_1\}$, $\{u_2, v_2\}$).

Lemma 24 (*a*) *If* $\varphi : \mathcal{H} \to \mathcal{T}$ *then*:

- $\varphi(x_0) = u_0$, $\varphi(x_1) = v_1$ *and* $\varphi(x_2) = v_2$;
- $\varphi(x_0) = v_0$, $\varphi(x_1) = u_1$ *and* $\varphi(x_2) = v_2$; *or*
- $\varphi(x_0) = v_0$, $\varphi(x_1) = v_1$ *or* $\varphi(x_2) = u_2$.

(*b*) *There exist homomorphisms* $\varphi : \mathcal{H} \to \mathcal{T}$ *extending the partial maps*:

- $\varphi(x_0) = u_0$, $\varphi(x_1) = v_1$ *and* $\varphi(x_2) = v_2$;
- $\varphi(x_0) = v_0$, $\varphi(x_1) = u_1$ *and* $\varphi(x_2) = v_2$; *and*
- $\varphi(x_0) = v_0$, $\varphi(x_1) = v_1$ *or* $\varphi(x_2) = u_2$.

Proof (*a*) Suppose that $\varphi : \mathcal{H} \to \mathcal{T}$. If $\varphi(x_i) = v_i$, for some $i \in \{0, 1, 2\}$, then by arguing as we did above when we proved our algorithm correct, $\varphi((w_0, w_1, w_2)) \neq w_i$. Without loss of generality, suppose that $\varphi(x_0) = u_0$.

Let (u', v', w') be the penultimate vertex on the path in $\mathcal{G}$ from (w_0, w_1, w_2) to (v_0, v_0, u_0). As φ is a homomorphism and the digraph described by each $H_i^{\mathcal{F}}$ is N-free (see also the parenthetical remark made in the definition of $\mathcal{F}$, above), we must have that $\varphi((u', v', w')) = w'$. Reasoning similarly back up the path to (w_0, w_1, w_2) yields that $\varphi((w_0, w_1, w_2)) = w_2$. Hence, again by reasoning as we did when we proved our algorithm correct, we must have that $\varphi(x_1) = v_1$ and $\varphi(x_2) = v_2$.

(b) Without loss of generality, suppose that φ is a partial map from $|\mathcal{H}|$ to $|\mathcal{T}|$ such that $\varphi(x_0) = u_0$, $\varphi(x_1) = v_1$ and $\varphi(x_2) = v_2$. For any element (u', v', w') of $|\mathcal{H}|$, define $\varphi((u', v', w')) = w'$. By arguing as we have done previously, this yields a homomorphism $\varphi : \mathcal{H} \to \mathcal{T}$. □

Now, consider the following σ-structure $\mathcal{K}_0^1$. Take the distinct set of elements $\{p_i, q_i, z_i : i = 0, 1, 2\}$ and disjoint copies of $\mathcal{H}$ as follows.

- For each $i \in \{0, 1, 2\}$, include a copy of $\mathcal{H}$ by identifying the elements: p_i and x_0; q_i and x_1; and z_i and x_2.
- For each $i \in \{0, 1, 2\}$, include a copy of $\mathcal{H}$ by identifying the elements: p_i and x_0; q_{i-1} and x_1; and z_{i-2} and x_2, where subtraction is modulo 3.
- The list associated with element z_0 is $\{u_2\}$, with all other lists being as in $\mathcal{H}$ (and so, for example: the list associated with p_0, p_1 and p_2 is $\{u_0, v_0\}$; the list associated with q_0, q_1 and q_2 is $\{u_1, v_1\}$; and the list associated with z_1 and z_2 is $\{u_2, v_2\}$).

Rename the element p_1 with the name y_0 and the element q_2 with the name y_1. The construction of $\mathcal{K}_0^1$ can be visualized as in Fig. 4.

Lemma 25 (a) *If* $\varphi : \mathcal{K}_0^1 \to \mathcal{T}$ *then*:

- $\varphi(y_0) = u_0$ *and* $\varphi(y_1) = u_1$; *or*
- $\varphi(y_0) = v_0$ *and* $\varphi(y_1) = v_1$.

(b) *There exist homomorphisms* $\varphi : \mathcal{K}_0^1 \to \mathcal{T}$ *extending the partial maps*:

- $\varphi(y_0) = u_0$ *and* $\varphi(y_1) = u_1$; *and*
- $\varphi(y_0) = v_0$ *and* $\varphi(y_1) = v_1$.

Proof Throughout we reason as we did earlier and in Lemma 24 but with respect to the different copies of $\mathcal{H}$, which are composed to form $\mathcal{K}_0^1$, and using Lemma 24.

Figure 4. The σ-structure $\mathcal{K}_0^1$.

(a) Suppose that $\varphi : \mathcal{K}_0^1 \to \mathcal{T}$. As $\varphi(z_0) = u_2$, we must have that $\varphi(q_1) = v_1$ and $\varphi(p_2) = v_0$ (by considering the copy of $\mathcal{H}$ corresponding to p_2, q_1 and z_0). Also, $\varphi(q_0) = v_1$ and $\varphi(p_0) = v_0$ (by considering the copy of $\mathcal{H}$ corresponding to p_0, q_0 and z_0).

Suppose that $\varphi(y_0) = u_0$. Thus, $\varphi(z_1) = v_2$ and $\varphi(z_2) = v_2$. Hence, $\varphi(y_1) = u_1$.

Suppose that $\varphi(y_0) = v_0$. Thus, $\varphi(z_1) = u_2$ and $\varphi(y_1) = v_1$ (and $\varphi(z_2) = u_2$).

(b) Follows by extending the partial maps of part (a) to homomorphisms of $\mathcal{K}_0^1$ to $\mathcal{T}$ using Lemma 24. $\square$

There is an analogous result to Lemma 25 except where: the σ-structure is $\mathcal{K}_1^2$; the elements of $\mathcal{T}$ are $\{u_1, v_1, u_2, v_2\}$; and the distinguished elements are y_1 and y_2, respectively.

Let $\mathcal{A}$ be an instance of 1-in-3-SAT; that is, a collection of clauses $C_1, C_2, \ldots, C_m$ of 3 Boolean variables over the set $\{X_1, X_2, \ldots, X_n\}$. Construct the σ-structure $\mathcal{B}$ (where every element has an associated list) as follows.

- For every $i \in \{1, 2, \ldots, n\}$, there are elements X_i^0, X_i^1 and X_i^2.
- For every $i \in \{1, 2, \ldots, n\}$, introduce a copy of $\mathcal{K}_0^1$ (resp. $\mathcal{K}_1^2$) by identifying X_i^0 with y_0 and X_i^1 with y_1 (resp. X_i^1 with y_1 and X_i^2 with y_2).
- For every $j \in \{1, 2, \ldots, m\}$, if $C_j = \{X_{j_0}, X_{j_1}, X_{j_2}\}$ then introduce a copy of $\mathcal{H}$ by identifying $X_{j_0}^0$ with x_0, $X_{j_1}^1$ with x_1 and $X_{j_2}^2$ with x_2.

All elements and copies so introduced are distinct (apart from where they are identified) and the lists at each element are inherited from the particular copies of $\mathcal{H}$, $\mathcal{K}_0^1$ and $\mathcal{K}_1^2$. The σ-structure $\mathcal{B}$ can be visualized as in Fig. 5 where we show the portion of $\mathcal{B}$ corresponding to a clause $\{X_{j_0}, X_{j_1}, X_{j_2}\}$.

Figure 5. The σ-structure $\mathcal{B}$.

Suppose that $\mathcal{A}$ is a yes-instance of 1-in-3-SAT; and so there is a satisfying truth assignment ψ, making exactly one variable in every clause *true*. If, for some $i \in \{1, 2, \ldots, n\}$: $\psi(X_i) = true$ then set $\varphi(X_i^0) = u_0$, $\varphi(X_i^1) = u_1$ and $\varphi(X_i^2) = u_2$; otherwise set $\varphi(X_i^0) = v_0$, $\varphi(X_i^1) = v_1$ and $\varphi(X_i^2) = v_2$.

Lemmas 24 and 25 yield that the partial map φ can be extended to a homomorphism from $\mathcal{B}$ to $\mathcal{T}$. Conversely, suppose that there is a homomorphism φ from $\mathcal{B}$ to $\mathcal{T}$. If, for some $i \in \{1, 2, \ldots, n\}$: $\varphi(X_i^0) = u_0$, $\varphi(X_i^1) = u_1$ and $\varphi(X_i^2) = u_2$ then define $\psi(X_i) = true$; and if $\varphi(X_i^0) = v_0$, $\varphi(X_i^1) = v_1$ and $\varphi(X_i^2) = v_2$ then define $\psi(X_i) = false$ (by Lemma 25, this truth assignment ψ is well defined). By Lemma 24, ψ is a satisfying truth assignment. Consequently part (b) of the theorem follows.

Every vertex in the digraph of a unary function has out-degree 1; and so such a digraph is N-free. The result follows. $\square$

Note that, from Theorem 23, whether the problem CSP($\mathcal{T}'$ + all lists) is solvable in polynomial-time or **NP**-complete is solely determined by whether $\mathcal{T}$ has a choice majority function or not.

Acknowledgement

The second two authors would like to thank Eric Allender for a useful discussion and guidance.

References

[1] A.A. Bulatov, A.A. Krokhin and P. Jeavons, Constraint satisfaction problems and finite algebras, *Proceedings of 27th International Colloquium on Automata, Languages and Programming, ICALP'00* (ed. U. Montanari, J.D.P. Rolim and E. Welzl), Lecture Notes in Computer Science Vol. 1853, Springer-Verlag, Berlin (2000) 272–282.

[2] A.K. Chandra and P.M. Merlin, Optimal implementation of conjunctive queries in relational databases, *Proceedings of 9th ACM Symposium on Theory of Computing* (1977) 77–90.

[3] J. Diaz, M. Serna and D.M. Thilikos, The complexity of restrictive H-coloring, *Discrete Applied Mathematics*, to appear.

[4] M. Dyer and C. Greenhill, The complexity of counting graph homomorphisms, *Random Structures and Algorithms* **17** (2000) 260–289.

[5] T. Feder, Classification of homomorphisms to oriented cycles and of k-partite satisfiability, *SIAM Journal of Discrete Mathematics* **14** (2001) 471–480.

[6] T. Feder and P. Hell, List homomorphisms to reflexive graphs, *Journal of Combinatorial Theory* (*Series B*) **72** (1998) 236–250.

[7] T. Feder, P. Hell and J. Huang, List homomorphisms and circular-arc graphs, *Combinatorica* **19** (1999) 487–505.

[8] T. Feder, P. Hell and J. Huang, Bi-arc graphs and the complexity of list homomorphisms, *Journal of Graph Theory*, to appear.

[9] T. Feder and M. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory, *SIAM Journal of Computing* **28** (1998) 57–104.

[10] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman (1979).

[11] P. Hell and J. Nešetřil, On the complexity of H-colouring, *Journal of Combinatorial Theory* (*Series B*) **48** (1990) 92–110.

[12] N. Immerman and S. Landau, The complexity of iterated multiplication, *Information and Computation* **116** (1995) 103–116.

[13] N. Immerman, *Descriptive Complexity*, Springer-Verlag (1998).

[14] P.G. Jeavons, D.A. Cohen and J.K. Pearson, Constraints and universal algebra, *Annals of Mathematics and Artificial Intelligence* **24** (1998) 51–67.

[15] P.G. Kolaitis and M.Y. Vardi, Conjunctive-query containment and constraint satisfaction, *Proceedings of* 17*th ACM Symposium on Principles of Database Systems* (1998) 205–213.

[16] C.H. Papadimitriou, *Computational Complexity*, Addison Wesley, 1994.

[17] A. Puricella and I.A. Stewart, A generic greedy algorithm, partially-ordered graphs and NP-completeness, *Proceedings of* 27*th International Workshop on Graph-Theoretic Concepts in Computer Science, WG*'01 (ed. A. Brandstaedt and V.B. Le), Lecture Notes in Computer Science Vol. 2204, Springer-Verlag, Berlin (2001) 306-316.

[18] A. Puricella and I.A. Stewart, Greedy algorithms, H-colourings and a complexity-theoretic dichotomy, *Theoretical Computer Science*, to appear.

[19] T.J. Schaefer, The complexity of satisfiability problems, *Proceedings of 10th ACM Symposium on Theory of Computing* (1978) 216–226.

[20] I.A. Stewart, Completeness of path-problems via logical reductions, *Information and Computation* **121** (1995) 123-134.

[21] E.P.K. Tsang, *Foundations of Constraint Satisfaction*, Academic Press (1993).