GRAPHS ADMITTING k-NU OPERATIONS. PART 2: THE IRREFLEXIVE CASE

TOMÁS FEDER, PAVOL HELL, BENOÎT LAROSE, MARK SIGGERS, AND CLAUDE TARDIF

Abstract. We describe a generating set for the variety of simple graphs that admit a k-ary near-unanimity polymorphism. The result follows from an analysis of NU polymorphisms of *strongly bipartite* digraphs, i.e. whose vertices are either a source or a sink but not both. We show that the retraction problem for a strongly bipartite digraph $\mathbb{H}$ has finite duality if and only if $\mathbb{H}$ admits a near-unanimity polymorphism. This result allows the use of tree duals to generate the variety of digraphs admitting a k-NU polymorphism.

Contents

Date: March 12, 2013.

2000 *Mathematics Subject Classification.* Primary 05C75; Secondary 08B05.

Key words and phrases. Near-unanimity polymorphism, graphs, strongly bipartite digraphs, finite duality.

The second author's research is partially supported by a Discovery Grant from NSERC, and by the grant ERCCZ LL 1201. The third author's research is partially supported by FQRNT and NSERC. The fourth author was supported by the Kyungpook National University Research Fund. The fifth author's research is partially supported by NSERC and ARP. The second and third authors wish to thank the fourth author and Kyungpook National University for their hospitality. The fourth author wishes to thank the third author and CRM for their hospitality.

1. Introduction

The present paper is a companion to [5] where the class of reflexive graphs admitting a compatible k-NU operation is described by a simple set of generators via products and retracts; we refer to that paper for a thorough discussion of the motivations behind the study of such operations on graphs. In the present paper, we describe a nice generating set for the class of simple graphs admitting compatible k-NU operations. This result is obtained as a corollary of an analogous result (Theorem 4.2) for digraphs we call *strongly bipartite*, i.e. digraphs for which every vertex is a source or a sink (Definition 2.2). For this class of digraphs, we prove that the existence of a compatible NU operation is equivalent to the corresponding constraint satisfaction problem having finite duality (Theorem 3.1). Although the main idea underlying our characterization of k-NU graphs is essentially the same as in the case of reflexive graphs, namely by using duals of coloured trees, the present paper differs from its companion in at least two important aspects: firstly, in the reflexive case, the equivalence of finite duality and existence of an NU polymorphism was already known [11] whereas in the case of strongly bipartite digraphs this had to be proved (Theorem 3.1); secondly, in [5] we introduced the notion of *reflexive duals*, a simple analogue of tree dual tailor-made for the study of reflexive graphs; in the present paper we create our building blocks directly from digraph duals but the details are a bit more involved. The results for undirected graphs follow from the strongly bipartite case without much difficulty.

As a consequence of our work, it is shown in [9] that graphs admitting an NU polymorphism give rise to constraint satisfaction problems solvable in Logspace via the language Symmetric Datalog. The bipartite analog of our results in [5] on reflexive absolute retracts can be found in [12] where a more thorough study of k-absolute retracts for graphs and digraphs is developed.

We now outline the contents of the paper. In the next section we describe the preliminary notions we'll require. In Section 3 we prove that a connected, strongly bipartite digraph admits an NU polymorphism if and only if its retraction problem has finite duality (Theorem 3.1). In Section 4 we prove the main theorem for strongly bipartite digraphs, describing the building blocks that generate the class of k-NU digraphs (Theorem 4.2). In Section 5 we apply the above results to simple (irreflexive) graphs (Theorem 5.2). We conclude in Section 6 with a short discussion of the question of recognizing k-NU graphs.

2. Preliminaries

2.1. **Structures and homomorphisms.** We refer the reader to [4] for basic notation and terminology. In the present paper we'll use blackboard fonts such as $\mathbb{G}$, $\mathbb{H}$, etc. to denote relational structures and their latin equivalent G, H, etc. to denote their respective universes. A signature τ is a (finite) set of relation symbols with associated arities. We say that $\mathbb{H} = \langle H; R(\mathbb{H})(R \in \tau)\rangle$ is a *relational structure of signature* τ if $R(\mathbb{H})$ is a relation on H of the corresponding arity, for each relation symbol $R \in \tau$.

Let $\mathbb{G}$ and $\mathbb{H}$ be structures of signature τ. A *homomorphism* from $\mathbb{G}$ to $\mathbb{H}$ is a map f from G to H such that $f(R(\mathbb{G})) \subseteq R(\mathbb{H})$ for each $R \in \tau$. We write $\mathbb{G} \to \mathbb{H}$ to indicate there exists a homomorphism from $\mathbb{G}$ to $\mathbb{H}$. A homomorphism $r : \mathbb{H} \to \mathbb{R}$ is a *retraction* if there exists a homomorphism (called a *coretraction*) $e : \mathbb{R} \to \mathbb{H}$ such that $r \circ e$ is the identity on $\mathbb{R}$ and we say $\mathbb{R}$ is a *retract* of $\mathbb{H}$ and write $\mathbb{R} \trianglelefteq \mathbb{H}$.

Notice that the relation $\trianglelefteq$ is transitive. A structure $\mathbb{H}$ is called a *core* if every homomorphism from $\mathbb{H}$ to itself is a permutation of H; note that a retract of $\mathbb{H}$ of minimum cardinality is a core, and is unique up to isomorphism, and hence we may speak of *the core* of the structure $\mathbb{H}$.

Let $\mathbb{H}$ be a τ-structure. We denote by $\mathrm{CSP}(\mathbb{H})$ the class of all τ-structures $\mathbb{G}$ that admit a homomorphism to $\mathbb{H}$, and by $\neg\,\mathrm{CSP}(\mathbb{H})$ the complement of this class. We say that $\mathrm{CSP}(\mathbb{H})$ has *finite duality* if there exist finitely many τ-structures $\mathbb{T}_1, \ldots, \mathbb{T}_s$ such that the following holds: for every τ-structure $\mathbb{G}$, there is *no* homomorphism from $\mathbb{G}$ to $\mathbb{H}$ precisely if there is some $\mathbb{T}_i$ that admits a homomorphism to $\mathbb{G}$. The set $\{\mathbb{T}_1, \ldots, \mathbb{T}_s\}$ is called a *duality* for $\mathbb{H}$. From a complexity-theoretic point of view, CSPs with finite duality are the simplest of all constraint satisfaction problems.

Throughout this paper we consider the usual product of τ-structures, namely if $\mathbb{G}$ and $\mathbb{H}$ are τ-structures then their product is the τ-structure $\mathbb{G} \times \mathbb{H}$ with universe $G \times H$ and where $R(\mathbb{G} \times \mathbb{H}) = R(\mathbb{G}) \times R(\mathbb{H})$ for all $R \in \tau$. We shall consider the notations $\prod_{i=1}^{n} \mathbb{G}_i$ and such to be self-evident.

Let $\mathbb{H}$ be a τ-structure. The *retraction problem for* $\mathbb{H}$ is the following: given a structure $\mathbb{G}$ containing a copy of $\mathbb{H}$, decide if $\mathbb{G}$ retracts to $\mathbb{H}$. It is in fact equivalent under positive first-order reductions to the *one-or-all list-homomorphism* problem for $\mathbb{H}$: an input consists of a τ-structure $\mathbb{G}$ with certain vertices coloured by a pre-assigned value from H, and the problem is to determine if there exists a homomorphism from $\mathbb{G}$ to $\mathbb{H}$ that extends these values. For brevity's sake we shall still refer to the latter as the retraction problem. Formally, we "add constants" to structures, i.e. we add, as basic unary relations to a given structure, each of its one-element sets:

Definition 2.1. *Let $\mathbb{H}$ be a τ-structure. For each $h \in H$, let S_h be a unary relation symbol. Let $\tau_{\mathbb{H}} = \tau \cup \{S_h : h \in H\}$, and let $\mathbb{H}^c$ denote the $\tau_{\mathbb{H}}$-structure obtained from $\mathbb{H}$ by adding all relations $S_h(\mathbb{H}^c) = \{h\}$. The problem $CSP(\mathbb{H}^c)$ is called the* retraction problem for $\mathbb{H}$. *Let $\mathbb{G}$ be a $\tau_{\mathbb{H}}$-structure. We say that a vertex $x \in G$ is* coloured *if it belongs to some unary relation $S_h(\mathbb{G})$ and refer to h as its colour (a vertex may have several colours). Let $\mathbb{G}^\tau$ denote the (reduct) τ-structure obtained from $\mathbb{G}$ by simply removing the relations indexed by the S_h.*

Let H be a non-empty set, let θ be an m-ary relation on H and let $f : H^k \to H$ be a k-ary operation on H. We say that f *preserves* θ if the following holds: if a $k \times m$ matrix M has each column in θ, then applying f to the rows of M yields a tuple of θ. If $\mathbb{H}$ is a τ-structure, and f preserves each of its basic relations (equivalently, if f is a homomorphism from $\mathbb{H}^k$ to $\mathbb{H}$), we say that f is a *polymorphism* of $\mathbb{H}$, or that $\mathbb{H}$ *admits* f; one also says that f is *compatible* with $\mathbb{H}$, see [4] and [17] for instance. Recall that f is a *k-ary near-unanimity (k-NU)* operation if it satisfies, for every $1 \leq i \leq k$ the identity

$$f(x, \ldots, x, \underbrace{y}_{i}, x, \ldots, x) = x.$$

A structure is said to be *k-NU* if it admits a k-NU polymorphism.

2.2. **Graphs and digraphs.** A *digraph* is a relational structure equipped with a single, binary relation. A *graph* $\mathbb{H}$ is a relational structure $\mathbb{H} = \langle H, E(\mathbb{H})\rangle$ where $E(\mathbb{H})$ is a binary relation which is *symmetric*, i.e. $(x, y) \in E(\mathbb{H})$ if and only if $(y, x) \in E(\mathbb{H})$. The graph $\mathbb{H}$ is *reflexive (irreflexive)* if $(x, x) \in E(\mathbb{H})$ ($(x, x) \notin E(\mathbb{H})$) for all $x \in H$.

In the present paper, it will be convenient for us to consider both undirected and directed bipartite graphs; for this purpose we'll require a bit of notation and terminology.

Definition 2.2. *A digraph* $\mathbb{H}$ *is* strongly bipartite *if its vertex set can be partitioned into two sets* D *and* U *such that every arc* (x, y) *of* $\mathbb{H}$ *satisfies* $x \in D$ *and* $y \in U$.

Notice that $\mathbb{H}$ is strongly bipartite precisely if it admits a homomorphism to a single arc; or equivalently, if each vertex is a source or a sink but not both.

Let $\mathbb{G}$ be any digraph; then we let $\mathbb{G}^u$ denote the underlying undirected graph, i.e. $E(\mathbb{G}^u)$ consists of all pairs (x, y) such that (x, y) or (y, x) is in $E(\mathbb{G})$. Let $\mathbb{H}$ be a digraph. We say $\mathbb{H}$ is *connected* if the graph $\mathbb{H}^u$ is connected; if $\mathbb{H}$ is not connected, a set G of vertices of $\mathbb{H}$ is a *connected component* of $\mathbb{H}$ if it is a connected component of $\mathbb{H}^u$.

The next two results will allow us to consider only connected digraphs in what follows.

Lemma 2.3. *Let* $\mathbb{H}$ *be a bipartite graph or strongly bipartite digraph with colour classes* D *and* U*, and assume* $\mathbb{H}$ *is connected. Let* $k \geq 3$*, let* $(x_1, \ldots, x_k) \in H^k$ *and let* $\Delta \subseteq \{1, \ldots, k\}$ *denote the set of indices* i *such that* $x_i \in D$.

(1) *If* $(x_1, \ldots, x_k)$ *and* $(y_1, \ldots, y_k)$ *are in the same connected component of* $\mathbb{H}^k$ *then* $\Delta = \{i : y_i \in D\}$ *or* $\Delta = \{i : y_i \in U\}$;
(2) *if* $\mathbb{H}$ *is a graph or if* $|\Delta| \in \{0, k\}$ *then the above condition is also sufficient;*
(3) *if* $\mathbb{H}$ *is a digraph and* $0 < |\Delta| < k$ *then* $(x_1, \ldots, x_k)$ *is an isolated vertex of* $\mathbb{H}^k$.

Proof. The first and third statements are straightforward. For the second, let $\bar{x} = (x_1, \ldots, x_k)$ and $\bar{y} = (y_1, \ldots, y_k)$ be tuples whose entries in D are precisely those with a coordinate in Δ. Since $\mathbb{H}$ is connected we may find for each i a path from x_i to y_i, and by repeating vertices if necessary we may choose these paths to be of the same length; it follows that we have a path in $\mathbb{H}^k$ from $\bar{x}$ to $\bar{y}$. The argument if $y_i \in D$ precisely for coordinates $i \notin \Delta$ is similar. $\square$

Proposition 2.4. *Let* $\mathbb{H}$ *be an irreflexive graph or a strongly bipartite digraph. Then* $\mathbb{H}$ *admits a* k*-NU polymorphism if and only if each of its subgraphs induced by a connected component does.*

Proof. Let f be a k-NU polymorphism of $\mathbb{H}$. We modify f so that it preserves every connected component H_1 of $\mathbb{H}$, i.e. we redefine f on H_1^k, if necessary, so that $f(H_1^k) \subseteq H_1$. Let K be a connected component of H_1^k. If K contains a tuple of the form $(x, \ldots, x, y, x, \ldots, x)$, then clearly $f(K) \subseteq H_1$ so there is nothing to do. Otherwise, redefine f on K as the projection on the first coordinate. The resulting map is clearly an NU polymorphism that preserves H_1. This can be done independently for each component of $\mathbb{H}$.

For the converse, suppose that $\mathbb{H}$ has connected components $H_1, \ldots, H_n$ and that f_i is a k-NU polymorphism of the induced subgraph $\mathbb{H}_i$ for each $1 \leq i \leq n$. Define a k-ary operation f on $\mathbb{H}$ by defining it on each block $\mathbb{H}_{i_1} \times \cdots \times \mathbb{H}_{i_k}$ as follows: if all the indices are equal to i let $f = f_i$; if exactly $k - 1$ of the indices are identical, let f be the projection on the leftmost repeated index; otherwise, let f be the first projection. It is easy to see that this is a homomorphism and that it is NU. $\square$

For the remainder of the paper, unless otherwise stated, the signature τ is that of digraphs, i.e. it consists of a single, binary relation symbol.

3. NU Polymorphisms of Strongly Bipartite Digraphs and Finite Duality

In this section we link the existence of NU polymorphisms of strongly bipartite digraphs to finite duality. It is known that if $\mathbb{H}$ is a core and $CSP(\mathbb{H})$ has finite duality, then $\mathbb{H}$ admits an NU polymorphism [10]. Although the converse does not hold in general (not even for reflexive digraphs), it is known to hold for reflexive graphs [11]; the main result of this section states that this is also true for strongly bipartite digraphs:

Theorem 3.1. *Let $\mathbb{H}$ be a connected, strongly bipartite digraph. Then* $\mathrm{CSP}(\mathbb{H}^c)$ *has finite duality if and only if $\mathbb{H}$ admits an NU polymorphism.*

To prove this result we'll appeal to an analog for posets proved in [13]. Recall that a digraph is a *poset* if it is reflexive, antisymmetric and transitive. In [13], it is proved that a finite connected poset admits an NU polymorphism if and only if it has finitely many poset "zigzags", i.e. critical poset obstructions. We restate this result for our needs. Let $\mathbb{P}$ be a τ-structure which is a connected poset. We'll say that a $\tau_{\mathbb{P}}$-structure $\mathbb{Q}$ is a $\mathbb{P}$-*coloured poset* if $\mathbb{Q}^\tau$ is a poset. Let $CSP_p(\mathbb{P}^c)$ denote the class of $\mathbb{P}$-coloured posets in $CSP(\mathbb{P}^c)$. We'll say that $CSP_p(\mathbb{P}^c)$ *has finite duality* if there exist finitely many $\mathbb{P}$-coloured posets $\mathbb{Z}_1, \dots, \mathbb{Z}_s$ such that, for any $\mathbb{P}$-coloured poset $\mathbb{Q}$, $\mathbb{Q} \not\to \mathbb{P}^c$ if and only if there exists i such that $\mathbb{Z}_i \to \mathbb{Q}$. In other words, $CSP_p(\mathbb{P}^c)$ has finite duality if the decision problem has finite duality when considering only instances that are $\mathbb{P}$-coloured posets.

Theorem 3.2 ([13])**.** *Let $\mathbb{P}$ be a connected poset. Then $CSP_p(\mathbb{P}^c)$ has finite duality if and only if $\mathbb{P}$ admits an NU polymorphism.*

Let $\mathbb{H}$ be a connected strongly bipartite digraph. Let D and U denote the colour classes of $\mathbb{H}$ such that $E(\mathbb{H}) \subseteq D \times U$. If $X \subseteq D$ let $N(X)$ denote the set of $y \in U$ such that $(x, y) \in E(\mathbb{H})$ for all $x \in X$; similarly if $X \subseteq U$ let $N(X)$ denote the set of $y \in D$ such that $(y, x) \in E(\mathbb{H})$ for all $x \in X$. If $X = \{x\}$ we write $N(x)$ instead of $N(\{x\})$. If $X \subseteq D$ or $X \subseteq U$ let $N^2(X) = N(N(X))$. It is immediate that if $X \subseteq Y$ then $N(X) \supseteq N(Y)$ and $N^2(X) \subseteq N^2(Y)$. Notice finally that $A \times B \subseteq D \times U$ is a maximal bipartite clique of $\mathbb{H}^u$ precisely when $N(A) = B$ and $A = N(B)$. Define a poset $\mathbb{P} = \mathbb{P}_{\mathbb{H}}$ as follows: its vertices are the sets $A \times B \subseteq D \times U$ such that either

(1) $A \times B = \{d\} \times N(d)$ for some $d \in D$, or
(2) $A \times B = N(u) \times \{u\}$ for some $u \in U$, or
(3) $|A|, |B| \geq 2$ and $A \times B$ is a maximal bipartite clique of $\mathbb{H}^u$.

The ordering of $\mathbb{P}$ is defined by $A \times B \leq C \times D$ if $A \subseteq C$ and $B \supseteq D$. In particular, vertices of type (1) are the minimal elements of $\mathbb{P}$ and vertices of type (2) are the maximal elements of $\mathbb{P}$. Notice that the substructure of $\mathbb{P}$ induced by its extremal elements is isomorphic to the digraph $\mathbb{H}$. Through this isomorphism, any $\tau_{\mathbb{H}}$-structure can be viewed as a $\tau_{\mathbb{P}}$-structure by renaming the colours in the obvious way.

Definition 3.3. *Let $\mathbb{H}$ be a strongly bipartite digraph, and let $\mathbb{G}$ be a $\tau_{\mathbb{H}}$-structure. We say that $\mathbb{G}$ is an* $\mathbb{H}$-coloured, strongly bipartite digraph *if $\mathbb{G}^\tau$ is a strongly* *bipartite digraph; furthermore we'll say it is* consistently coloured *if its vertex set can be partitioned into sets A and B such that*

(1) $E(\mathbb{G}) \subseteq A \times B$, *and*
(2) if a vertex of A is coloured then it has a unique colour and it belongs to D, and
(3) if a vertex of B is coloured then it has a unique colour and it belongs to U.

Lemma 3.4. *Let $\mathbb{H}$ be a connected strongly bipartite digraph, and let $\mathbb{P} = \mathbb{P}_{\mathbb{H}}$. If $CSP_p(\mathbb{P}^c)$ has finite duality, then $CSP(\mathbb{H}^c)$ has finite duality.*

Proof. Let $\mathbb{G}$ be a consistently $\mathbb{H}$-coloured strongly bipartite digraph. As we remarked after the definition of $\mathbb{P}$ earlier, we can interpret such a $\mathbb{G}$ as an input to $CSP(\mathbb{P}^c)$: we claim that in fact $\mathbb{G} \to \mathbb{H}^c$ if and only if $\mathbb{G} \to \mathbb{P}^c$. Indeed, if $f : \mathbb{G} \to \mathbb{H}^c$, define $F : \mathbb{G} \to \mathbb{P}^c$ by

$$F(x) = \begin{cases} \{f(x)\} \times N(\{f(x)\}), & \text{if } x \in A, \\ N(\{f(x)\}) \times \{f(x)\}, & \text{if } x \in B. \end{cases}$$

Clearly F is both edge- and colour-preserving (it is just a reinterpretation of f within $\mathbb{P}^c$). Conversely, let $f : \mathbb{G} \to \mathbb{P}^c$. For every $x \in A$ let $f'(x)$ be any minimal element of $\mathbb{P}^c$ below $f(x)$ and for $x \in B$ let $f'(x)$ be any maximal element of $\mathbb{P}^c$ above $f(x)$. It is clear that f' is a homomorphism from $\mathbb{G}$ to $\mathbb{P}^c$ whose image is contained in the set of extremal elements of $\mathbb{P}^c$; it is then immediate that f' can be interpreted as a homomorphism from $\mathbb{G}$ to $\mathbb{H}^c$.

For any $\tau_{\mathbb{P}}$-structure $\mathbb{G}$, let $\mathbb{G}^l$ denote the $\tau_{\mathbb{P}}$-structure obtained from $\mathbb{G}$ by adding all loops to the underlying digraph $\mathbb{G}^\tau$. Let $\mathbb{Z}_1, \ldots, \mathbb{Z}_s$ be $\mathbb{P}$-coloured posets that witness the fact that $CSP_p(\mathbb{P}^c)$ has finite duality. Consider the class $\mathcal{F}$ that consists of the following $\tau_{\mathbb{P}}$-structures (see Figure 2):

(1) for all distinct $a, b \in H$, a single vertex coloured with a and b;
(2) for every $u \in U$, a single directed edge (x, y), with x coloured by u;
(3) for every $d \in D$, a single directed edge (x, y), with y coloured by d;
(4) the directed path of length 2 (with no colours);
(5) all structures $\mathbb{Q}$ such that $\mathbb{Q}^l$ is a homomorphic image of some $\mathbb{Z}_i$.

We claim that $\mathcal{F}$ is a finite duality for $\mathbb{H}^c$. Indeed, it is is clear that $\mathcal{F}$ is finite, and that no member of it admits a homomorphism to $\mathbb{H}^c$. Now suppose that a structure $\mathbb{G}$ admits no homomorphism to $\mathbb{H}^c$. If one of the structures in (1), (2), (3) or (4) maps to $\mathbb{G}$ we are done. Otherwise, $\mathbb{G}^\tau$ is a consistently coloured strongly bipartite digraph. By the above argument, $\mathbb{G} \not\to \mathbb{P}^c$, and in fact $\mathbb{G}^l \not\to \mathbb{P}^c$. Clearly $\mathbb{G}^l$ is a $\mathbb{P}$-coloured poset, thus there exists some $\mathbb{Z}_i \to \mathbb{G}^l$. Let $\mathbb{Q}$ denote the $\tau_{\mathbb{P}}$-structure obtained from the homomorphic image of $\mathbb{Z}_i$ in $\mathbb{G}$ by removing all its loops. Then $\mathbb{Q}^l$ is a homomorphic image of $\mathbb{Z}_i$ and thus $\mathbb{Q} \in \mathcal{F}$, and $\mathbb{Q} \to \mathbb{G}$. $\square$

Lemma 3.5. *Let $\mathbb{H}$ be a connected strongly bipartite digraph. If $\mathbb{H}$ admits a k-ary NU polymorphism then so does $\mathbb{P}_{\mathbb{H}}$.*

Proof. Let D and U denote the colour classes of $\mathbb{H}$ such that $E(\mathbb{H}) \subseteq D \times U$. Let $f : \mathbb{H}^k \to \mathbb{H}$ be a k-ary NU polymorphism. Notice that since f is idempotent and $\mathbb{H}$ is connected we have $f(D^k) \subseteq D$ and $f(U^k) \subseteq U$.

For every $d \in D$ define I_d to be the following set of vertices of $\mathbb{P}_{\mathbb{H}}^k$

$I_d = \{(A_1 \times B_1, \ldots, A_k \times B_k) : A_i \times B_i = \{d\} \times N(d)$ for at least $k-1$ values of $i\}$, and for every $u \in U$ define

$$J_u = \{(A_1 \times B_1, \ldots, A_k \times B_k) : A_i \times B_i = N(u) \times \{u\} \text{ for at least } k-1 \text{ values of } i\}.$$

Since $k \geq 3$ it is clear that all these subsets of $\mathbb{P}^k$ are disjoint. Notice also that for any d, I_d is an ideal of $\mathbb{P}^k$, i.e. if $x \leq y \in I_d$ then $x \in I_d$; dually, for every u the set J_u is a filter of $\mathbb{P}^k$ i.e. if $x \geq y \in J_u$ then $x \in J_u$.

If $A_1, \ldots, A_n \subseteq D$, let $f(A_1, \ldots, A_k)$ denote the set

$$\{f(a_1, \ldots, a_k) : a_i \in A_i, i = 1, \ldots, k\}.$$

Define $F : \mathbb{P}^k \to \mathbb{P}$ as follows:

$$F(A_1 \times B_1, \ldots, A_k \times B_k) = \begin{cases} \{d\} \times N(d), & \text{if } (A_1 \times B_1, \ldots, A_k \times B_k) \in I_d, \\ N(u) \times \{u\}, & \text{if } (A_1 \times B_1, \ldots, A_k \times B_k) \in J_u, \\ N^2(f(A_1, \ldots, A_k)) \times N(f(A_1, \ldots, A_k)) & \text{otherwise.} \end{cases}$$

We first prove that F is well-defined: we must verify that in the case the input of F does not lie in some I_d nor in a J_u the value is indeed in $\mathbb{P}$. Let $\alpha = f(A_1, \ldots, A_k)$. Suppose first that $|N^2(\alpha)| = 1$. Then $N^2(\alpha) \supseteq \alpha$ implies $N^2(\alpha) = \alpha$ so $N^2(\alpha) \times N(\alpha) = \{d\} \times N(d)$ for some $d \in D$. If on the other hand $|N(\alpha)| = 1$, then clearly $N^2(\alpha) \times N(\alpha) = N(u) \times \{u\}$ for some $u \in U$. So we may now suppose that $|N^2(\alpha)|, |N(\alpha)| \geq 2$, so we must prove that $N^2(\alpha) \times N(\alpha)$ is a maximal bipartite clique of $\mathbb{H}_u$; for this it clearly suffices to prove that $N(N^2(\alpha)) = N(\alpha)$. Since $\alpha \subseteq N^2(\alpha)$ we have that $N(\alpha) \supseteq N(N^2(\alpha))$. Now let $y \in N(\alpha)$ and $z \in N^2(\alpha)$: by definition of N^2, z is a neighbour of y so $y \in N(N^2(\alpha))$ and we're done.

Next we prove that F obeys the near-unanimity identities. Fix some $1 \leq j \leq k$ and let $A_i \times B_i = A \times B$ for all $i \neq j$ and let $A_j \times B_j = A' \times B'$. If there is some $d \in D$ such that $(A_1 \times B_1, \ldots, A_k \times B_k) \in I_d$ or some $u \in U$ such that $(A_1 \times B_1, \ldots, A_k \times B_k) \in J_u$ then by definition we have that $F(A_1 \times B_1, \ldots) = A \times B$. Otherwise, we have that $|A|, |B| \geq 2$ and $N(A) = B$ and $N(B) = A$. We claim that

$$f(A, \ldots, A, A', A, \ldots, A) = A.$$

Indeed, if $a \in A$ then for any $x \in A'$ we have

$$a = f(a, \ldots, a, x, a, \ldots, a) \in f(A, \ldots, A, A', A, \ldots, A)$$

so one inclusion is taken care of. Now let $a_i \in A$ for $i \neq j$ and let $x \in A'$. Since $A = N(B)$, for every $b \in B$ we have $(a_i, b) \in E(\mathbb{H})$ for all $i \neq j$. Since $\mathbb{H}$ is connected, x has some neighbour $y \in U$. Thus

$$(f(a_1, \ldots, a_{j-1}, x, a_{j+1}, \ldots, a_k), f(b, \ldots, b, y, b, \ldots, b)) \in E(\mathbb{H})$$

and since $f(b, \ldots, b, y, b, \ldots, b) = b$ it follows that

$$f(a_1, \ldots, a_{j-1}, x, a_{j+1}, \ldots, a_k) \in N(b).$$

Since this holds for all $b \in B$, we conclude that

$$f(A, \ldots, A, A', A, \ldots, A) \subseteq N(B) = A$$

which proves our claim. Then

$$\begin{aligned} F(A_1 \times B_1, \ldots, A_k \times B_k) &= N^2(f(A, \ldots, A, A', A, \ldots)) \times N(f(A, \ldots, A, A', A, \ldots)) \\ &= N^2(A) \times N(A) \\ &= N(B) \times B \\ &= A \times B. \end{aligned}$$

Finally we show that F is order-preserving. Let

$$\alpha = (A_1 \times B_1, \ldots, A_k \times B_k) \leq (C_1 \times D_1, \ldots, C_k \times D_k) = \delta,$$

i.e. $A_i \subseteq C_i$ and $B_i \supseteq D_i$ for all $1 \leq i \leq k$. If $\alpha \in J_u$ for some $u \in U$ or if $\delta \in I_d$ for some $d \in D$ then $F(\alpha) = F(\delta)$. If $\alpha \in I_d$ for some $d \in D$ and $\delta \in J_u$ for some $u \in U$ then $F(\alpha) \leq F(\delta)$ since $k \geq 3$. Next suppose that α is in some I_d and that δ lies in no I_d and no J_u. Then $F(\alpha) = \{d\} \times N(d)$ and $F(\delta) = N^2(\epsilon) \times N(\epsilon)$ where $\epsilon = f(C_1, \ldots, C_k)$. Since $d \in C_i$ for all but one index i we have that $d = f(d, \ldots, d, x, d, \ldots) \in \epsilon \subseteq N^2(\epsilon)$ and $N(d) \supseteq N(\epsilon)$, thus $F(\alpha) \leq F(\delta)$. The case where δ is in some J_u and α lies in no I_d and no J_u is similar. So now assume that α and δ lie in no I_d and in no J_u. Then $F(\alpha) = N^2(\beta) \times N(\beta)$ and $F(\delta) = N^2(\epsilon) \times N(\epsilon)$ where $\beta = f(A_1, \ldots, A_k)$ and $\epsilon = f(C_1, \ldots, C_k)$. Since $A_i \subseteq C_i$ for all i it is immediate that $\beta \subseteq \epsilon$, and hence $N(\beta) \supseteq N(\epsilon)$ and $N^2(\beta) \subseteq N^2(\epsilon)$ so $F(\alpha) \leq F(\delta)$.

□

Proof of Theorem 3.1. Since $\mathbb{H}^c$ is a core, if $CSP(\mathbb{H}^c)$ has finite duality then $\mathbb{H}^c$, and hence $\mathbb{H}$, admits an NU polymorphism by Corollary 4.5 of [10]. Conversely, suppose that $\mathbb{H}$ admits an NU polymorphism. Then by Lemma 3.5 $\mathbb{P} = \mathbb{P}_{\mathbb{H}}$ also admits an NU polymorphism, so by Theorem 3.2 $CSP_p(\mathbb{P}^c)$ has finite duality and by Lemma 3.4 CSP($\mathbb{H}^c$) has finite duality. □

4. Strongly bipartite k-NU digraphs

We now state the main result of this paper, whose proof will be split into several lemmas. The theorem asserts that the following digraphs constitute a generating set for the class of k-NU strongly bipartite digraphs.

Definition 4.1. *Let $\mathbb{T}$ be a non-trivial tree and let D and U be its colour classes. Define a strongly bipartite digraph $G(\mathbb{T})$ as follows*[1]*: its vertices are of the form $(0, X)$ where X is a set of edges of $\mathbb{T}$ satisfying the condition that for every $d \in D$ of degree greater than 1 there exists a unique $e \in X$ incident to d, and of the form $(1, Y)$ where Y is a set of edges of $\mathbb{T}$ satisfying the condition that for every $u \in U$ of degree greater than 1 there exists a unique $e \in Y$ incident to u. There is an arc from (i, X) to (j, Y) if and only if $(i, j) = (0, 1)$ and $X \cap Y = \emptyset$.*

Example. Let $\mathbb{T}$ be the tree in Figure 1, a star with three leaves; let $D = \{d\}$ be the central vertex and let $U = \{a, b, c\}$ be $\mathbb{T}$'s leaves. The digraph $G(\mathbb{T})$ is pictured in Figure 1.

Theorem 4.2. *Let $k \geq 3$. Let $\mathbb{H}$ be a connected, strongly bipartite digraph. Then the following are equivalent:*

[1]**Note.** Strictly speaking, the digraph $G(\mathbb{T})$ depends on the choice of D and U, however if we exchange their role we simply obtain $G(\mathbb{T})$ with all arcs reversed. To simplify the discussion we shall not consider this technicality, as this should not create any confusion.

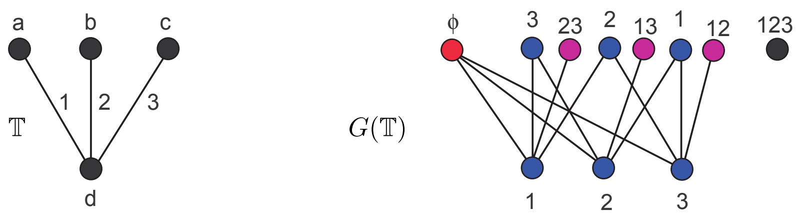

FIGURE 1. The tree $\mathbb{T}$ and its digraph $G(\mathbb{T})$; all edges are implicitly directed bottom to top. In $G(\mathbb{T})$'s diagram, the bottom vertices are those of the form $(0, X)$, the top ones of the form $(1, Y)$, and the labels indicate the corresponding set of edges, for instance the vertex $(1, \{1, 2\})$ is labelled simply 12.

(1) $\mathbb{H}$ *is a k-NU digraph;*
(2) $\mathbb{H}$ *is a retract of a product of finitely many digraphs of the form* $G(\mathbb{T})$ *where* $\mathbb{T}$ *is a tree with at most* $k-1$ *leaves.*

4.1. **Finite duality and duals of τ-trees.** For the remainder of this subsection let τ be an arbitrary signature. Let $\mathbb{T}$ and $\mathbb{H}$ be τ-structures. We say that $\mathbb{T}$ is an *obstruction* of $\mathbb{H}$ if there is no homomorphism from $\mathbb{T}$ to $\mathbb{H}$; furthermore if every proper substructure of $\mathbb{T}$ admits a homomorphism to $\mathbb{H}$ we say $\mathbb{T}$ is a *critical* obstruction of $\mathbb{H}$. The following is a very slight adaptation of a result of Zádori [18], its proof can be found in [5]:

Lemma 4.3. *Let* $\mathbb{H}$ *be a* τ*-structure. Then the following are equivalent:*

(1) $\mathbb{H}$ *is a k-NU structure;*
(2) $\mathbb{H}^c$ *is a k-NU structure;*
(3) every critical obstruction of $\mathbb{H}^c$ *has at most* $k-1$ *coloured elements.*

We shall require the notion of a τ-tree [10]. Let $\mathbb{T}$ be a τ-structure. We define the *incidence multigraph* $\mathrm{Inc}(\mathbb{T})$ of $\mathbb{T}$ as the bipartite multigraph with parts T and $\mathrm{Block}(\mathbb{T})$ which consists of all pairs (R, r) such that $R \in \tau$ and $r \in R(\mathbb{T})$, and with edges $e_{a,i,B}$ joining $a \in T$ to $B = (R, (x_1, \ldots, x_r)) \in \mathrm{Block}(\mathbb{T})$ when $x_i = a$. Roughly speaking, one colour class consists of all vertices of $\mathbb{T}$, the other ($\mathrm{Block}(\mathbb{T})$) consists of all tuples that appear in the relations $R(\mathbb{T})$ (with repetitions: if a tuple appears in several relations it will appear as many times in the multigraph); a vertex t is adjacent to a tuple if it appears in it. The structure $\mathbb{T}$ is a τ*-tree* if its associated multigraph is a tree, i.e. it is connected and acyclic.

Theorem 4.4 ([15])**.**

(1) Let $\mathbb{H}$ *be a* τ*-structure with finite duality. Then there exists a duality for* $\mathbb{H}$ *consisting of finitely many* τ*-trees;*
(2) (Existence of duals) Let $\mathcal{T}$ *be a finite family of* τ*-trees. Then there exists a* τ*-structure* D *such that* $\mathcal{T}$ *is a duality for* D*.*

4.2. **Digraph duals.** We now give an explicit description, in the case where τ is the signature of digraphs, of a structure $D(\mathbb{T})$ which is a dual to the $\tau_{\mathbb{H}}$-tree $\mathbb{T}$. In Lemma 4.14 we will see how the structures $D(\mathbb{T})$ are related to the digraphs described in Definition 4.1. It is not difficult to see that if $\mathbb{T}$ is a $\tau_{\mathbb{H}}$-tree then $(\mathbb{T}^{\tau})^u$ is a tree and each of its coloured vertices is a leaf.

Definition 4.5 ([16])**.** *Let τ be the signature of digraphs. Let $\mathbb{H}$ be a τ-structure and let $\mathbb{T}$ be a $\tau_{\mathbb{H}}$-tree. The $\tau_{\mathbb{H}}$-structure $D(\mathbb{T})$ has universe*

$$\{f : T \to \mathrm{Block}(\mathbb{T}) : [t, f(t)] \in E(\mathrm{Inc}(\mathbb{T})) \text{ for all } t \in T\},$$

and (f, g) is not *an arc of $D(\mathbb{T})$ if there exists an arc $e = (s, t)$ of $\mathbb{T}$ such that $f(s) = e = g(t)$; the map f does* not *belong to $S_h(D(\mathbb{T}))$ if there exists some $t \in S_h$ such that $f(t) = (t)$.*

Theorem 4.6 ([16])**.** *Let $\mathbb{H}$ be a τ-structure and let $\mathbb{T}$ be a $\tau_{\mathbb{H}}$-tree. Then $\{\mathbb{T}\}$ is a duality for $D(\mathbb{T})$.*

More generally, it is immediate by the definition of product that if $\mathbb{T}_1, \ldots, \mathbb{T}_s$ are $\tau_{\mathbb{H}}$-trees, then $\{\mathbb{T}_1, \ldots, \mathbb{T}_s\}$ is a duality for $\prod_{i=1}^{s} D(\mathbb{T}_i)$.

Definition 4.7. *Let $\mathbb{H}$ be any digraph. Let $\mathbb{P}_1(\mathbb{H})$ denote the $\tau_{\mathbb{H}}$-structure consisting of a single arc $(0, 1)$ and such that both 0 and 1 belong to every S_h, $h \in H$.*

The structure $\mathbb{P}_1(\mathbb{H})$ is pictured in the lower right corner of Figure 3, it is the dual of the path of length 2. We shall just write $\mathbb{P}_1$ when there is no confusion possible.

Lemma 4.8. *Let s be a positive integer. Let $\mathbb{H}$ be a connected strongly bipartite digraph. Then the core of the structure $\mathbb{P}_1^s \times \mathbb{H}^c$ is isomorphic to $\mathbb{H}^c$.*

Proof. Let D and U be the colour classes of $\mathbb{H}$. Notice first that the substructure of $\mathbb{P}_1^s \times \mathbb{H}^c$ induced by the set of vertices

$$\{(0, \ldots, 0, d) : d \in D\} \cup \{(1, \ldots, 1, u) : u \in U\}$$

is isomorphic to $\mathbb{H}^c$ via the projection on the last coordinate, and that the map $(x_1, \ldots, x_s, h) \mapsto h$ is a retraction of $\mathbb{P}_1^s \times \mathbb{H}^c$ onto $\mathbb{H}^c$. □

We distinguish a handful of basic critical obstructions (and their duals) that will be useful in what follows. We'll say a structure is *non-trivial* if it has at least 2 vertices.

Lemma 4.9. *Let $\mathbb{H}$ be a non-trivial, connected, strongly bipartite digraph with colour classes D and U. Then the following are critical obstructions of $\mathbb{H}^c$ (see Figure 3):*

(A) for all $a, b \in H$ distinct, a single vertex coloured with a and b;
(B) for every $d \in D$, a single arc with sink coloured by d;
(C) for every $u \in U$, a single arc with source coloured by u;
(D) the directed path of length 2 (with no colours).

Furthermore, for each $\mathbb{S}$ in the above list, the associated dual $D(\mathbb{S})$, the digraph $\mathbb{P}_1 \times D(\mathbb{S})$, and its core are pictured in Figure 3.

Proof. This is a simple exercise. □

Definition 4.10. *Let $\mathbb{H}$ be a non-trivial, connected, strongly bipartite digraph, and let $\mathbb{T}$ be a $\tau_{\mathbb{H}}$-tree. If $\mathbb{T}$ is a consistently $\mathbb{H}$-coloured strongly bipartite digraph whose coloured vertices are exactly the leaves of $(\mathbb{T}^{\tau})^u$, we'll say that $\mathbb{T}$ is a* good *$\tau_{\mathbb{H}}$-tree (see Figure 4).*

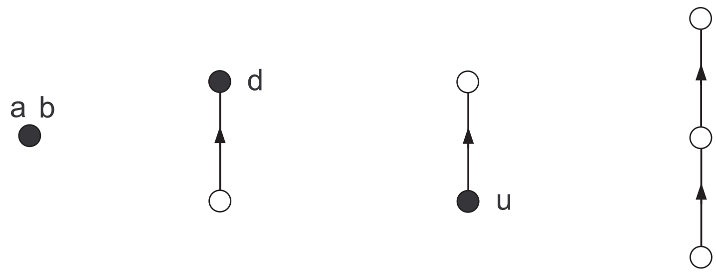

FIGURE 2. Basic critical obstructions.

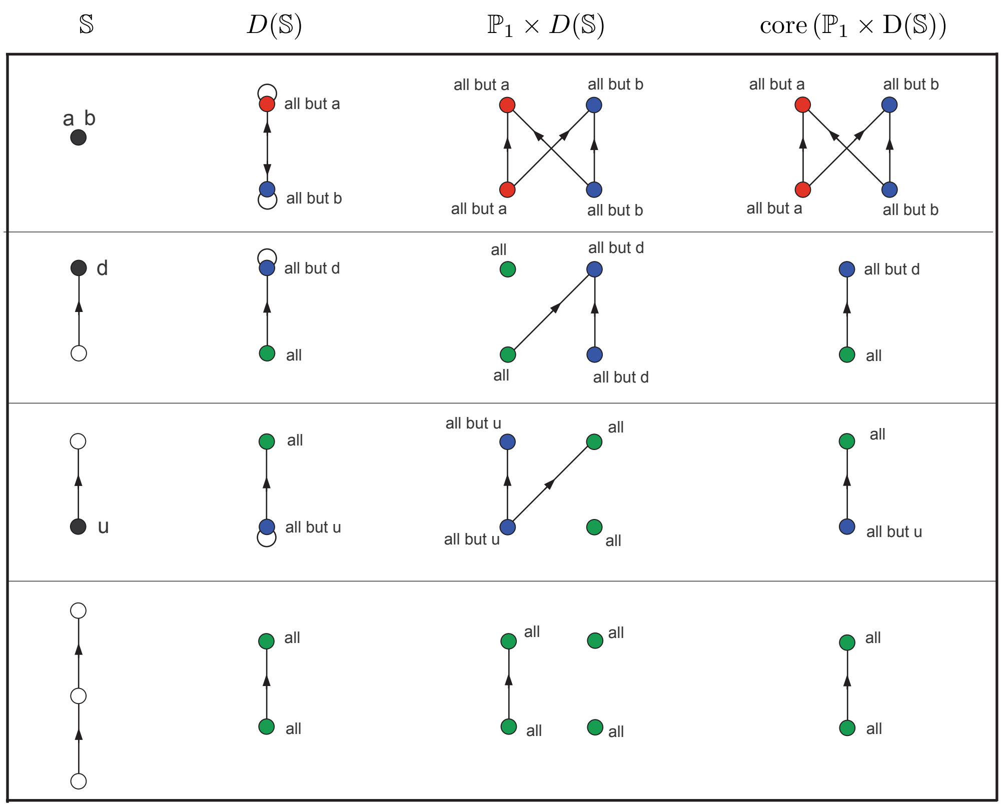

FIGURE 3. For each $\tau_{\mathbb{H}}$-tree $\mathbb{S}$ of Lemma 4.9, the associated dual $D(\mathbb{S})$, the digraph $\mathbb{P}_1 \times D(\mathbb{S})$ and its core.

Lemma 4.11. *Let $\mathbb{H}$ be a non-trivial, connected, strongly bipartite digraph. Then every $\tau_{\mathbb{H}}$-tree which is a critical obstruction of $\mathbb{H}^c$ other than those mentioned in Lemma 4.9 is a good $\tau_{\mathbb{H}}$-tree.*

Proof. Let $\mathbb{T}$ be a $\tau_{\mathbb{H}}$-tree which is a critical obstruction of $\mathbb{H}^c$ not of the form (A)-(D) in Lemma 4.9; it is easy to see that in fact, because $\mathbb{T}$ is critical, no obstruction

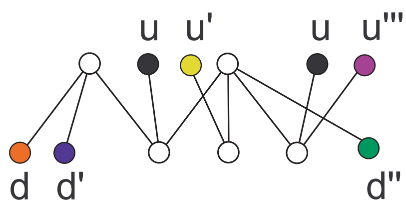

FIGURE 4. A good $\tau_{\mathbb{H}}$-tree. All edges are implicitly directed bottom to top.

of type (A)-(D) can admit a homomorphism to $\mathbb{T}$. In particular, by (A) we get that no vertex of $\mathbb{T}$ can have more than one colour, and by (D) every vertex of $\mathbb{T}$ is a source or a sink but not both, and by (B) and (C) $\mathbb{T}$ is consistently coloured, i.e. $\mathbb{T}$ is a consistently coloured strongly bipartite digraph. It is clear that no leaf of $\mathbb{T}$ is uncoloured since $\mathbb{T}$ is critical (one can retract this vertex). Finally, suppose that some vertex x of $\mathbb{T}$ is coloured and has (without loss of generality) out degree 2 or more; removal of the vertex x creates several subtrees of $\mathbb{T}$; to each of these we add back a copy of the vertex x with its original colour. By minimality of $\mathbb{T}$ there exist homomorphisms from each of these structures to $\mathbb{H}^c$, and since they all agree at x we can "glue" these maps back together to get a full homomorphism from $\mathbb{T}$ to $\mathbb{H}^c$, a contradiction. □

For our purposes, $\mathbb{H}$ is a *strongly bipartite path* if it is strongly bipartite and $\mathbb{H}^u$ is a path.

Lemma 4.12. *Let $\mathbb{T}$ be a path of length s. Then $G(\mathbb{T})$ is the disjoint union of a strongly bipartite path of length $s+2$ and isolated vertices.*

Proof. Let $\mathbb{T}$ be a path of length $2s$, $s \geq 1$; the case where $\mathbb{T}$ has odd length is similar. Assume that $|D| = |U| + 1$ and let the edges of the path $\mathbb{T}$ be labelled simply by $\{1, 2, \dots, 2s\}$ (see Figure 5). By definition, the vertices of $G(\mathbb{T})$ consist of pairs $(0, X)$ and pairs $(1, Y)$ where X and Y have the following form. A set $Y \subseteq \{1, \dots, 2s\}$ contains exactly one element of each pair $\{2j-1, 2j\}$ for $1 \leq j \leq s$. A set $X \subseteq \{1, \dots, 2s\}$ contains exactly one element of each pair $\{2i, 2i+1\}$ for $1 \leq i \leq s-1$, and may also contain 1 and/or $2s$.

Notice that if X contains a pair $\{2j-1, 2j\}$ where $1 \leq j \leq s$ then $(0, X)$ is an isolated vertex; similarly if Y contains $\{2i, 2i+1\}$ where $1 \leq i \leq s-1$, then $(1, Y)$ is isolated. We may thus restrict our attention to the vertices $(0, X_i)$ where $X_0 = \{1, 3, \dots, 2s-1\}$, $X_1 = \{3, 5, \dots, 2s-1\}$, $X_s = \{2, 4, \dots, 2s-2\}$, $X_{s+1} = \{2, 4, \dots, 2s\}$, and

$$X_i = \{2, \dots, 2i-2, 2i+1, \dots, 2s-1\},\ 2 \leq i \leq s-1;$$

and $(1, Y_j)$ where $Y_1 = \{2, 4, \dots, 2s\}$, $Y_{s+1} = \{1, 3, \dots, 2s-1\}$, and

$$Y_j = \{1, 3, \dots, 2j-3, 2j, \dots, 2s\},\ 2 \leq j \leq s.$$

Now observe that Y_j is disjoint from X_i only if $i = j-1$ or $i = j$. So the vertices $(0, X_0), (1, Y_1), (0, X_1), (1, Y_2), \dots, (1, Y_{s+1}), (0, X_{s+1})$ make up the path of length $2s+2$ that we needed to exhibit (see Figure 5).

□

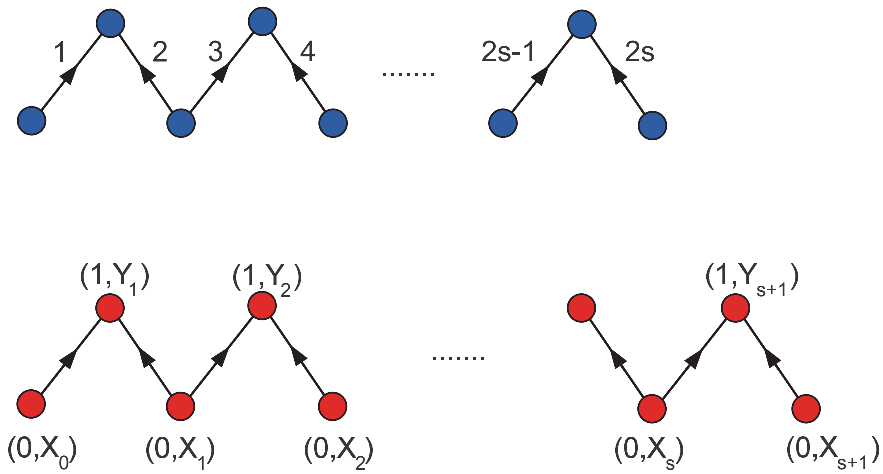

FIGURE 5. Above, the labelling of the edges of $\mathbb{T}$ used in the proof of Lemma 4.12. Below, the resulting path of length $2s+2$ in $G(\mathbb{T})$.

We now associate to every good $\tau_{\mathbb{H}}$-tree $\mathbb{T}$ a $\tau_{\mathbb{H}}$-structure $R(\mathbb{T})$. It turns out this structure is homomorphically equivalent to the product of $\mathbb{P}_1(\mathbb{H})$ and the dual of $\mathbb{T}$. Secondly, the underlying digraphs of these structures are precisely the building blocks we seek. Notice that $R(\mathbb{T})$ is similar to (the product of the arc with) the dual construction, except that functions are restricted to the colour classes of the tree.

Definition 4.13. *Let $\mathbb{H}$ be a strongly bipartite digraph, let $\mathbb{T}$ be a good $\tau_{\mathbb{H}}$-tree and let D and U denote the colour classes of the tree $(\mathbb{T}^\tau)^u$. Define the $\tau_{\mathbb{H}}$-structure $R(\mathbb{T})$ as follows: its universe consists of*

$$\left\{(0,f)\middle| f: D \to \mathrm{Block}(\mathbb{T}) \textit{ such that } [d, f(d)] \in E(\mathrm{Inc}(\mathbb{T})) \textit{ for all } d \in D\right\} \bigcup$$
$$\left\{(1,g)\middle| g: U \to \mathrm{Block}(\mathbb{T}) \textit{ such that } [u, g(u)] \in E(\mathrm{Inc}(\mathbb{T})) \textit{ for all } u \in U\right\};$$

the pair $((i,f),(j,g))$ is not *an arc of $D(\mathbb{T})$ if either $(i,j) \neq (0,1)$ or there exists an arc $e = (d,u)$ of $\mathbb{T}$ such that $f(d) = e = g(u)$; (i,f) does* not *belong to $S_h(R(\mathbb{T}))$ if there exists some $t \in S_h$ such that $f(t) = (t)$.*

Example. Let $\mathbb{H}$ be a strongly bipartite digraph with colour classes D and U. Let c_1, c_2, and c_3 in U have a common neighbour in D. Let $\mathbb{T}$ be the $\tau_{\mathbb{H}}$-tree consisting of the arcs $1 = (d, u_1), 2 = (d, u_2)$, and $3 = (d, u_3)$ and coloured vertices $u_i \in S_{c_i}(\mathbb{T})$ for $i = 1, 2, 3$. Observe that the reduct $\mathbb{T}^\tau$ is the first graph, $\mathbb{T}$, shown in Figure 1, (with $(u_1, u_2, u_3) = (a, b, c)$). We now describe $R(\mathbb{T})$. The reduct $R(\mathbb{T})^\tau$ is the graph $G(\mathbb{T})$ of Figure 1 where vertex labels in the figure are now interpreted as follows. For vertices on the bottom, the label i represents the vertex $(0, f_i)$ where f_i is the function mapping d to i. For vertices on the top, the label $S \subset [3]$ represents the vertex $(1, g_S)$ where $g_S : \{a, b, c\} \to \mathrm{Block}(\mathbb{T})$ is defined by $g_S(u_i) = i$ if $i \in S$ and $g_S(u_i) = u_i \in S_{c_i}(\mathbb{T})$ otherwise. Further $R(\mathbb{T})$ contains the following coloured vertices: $(0, f_i)$ is in $S_h(R(\mathbb{T}))$ for all $h \in \mathbb{H}$ and $(1, g_S)$ is in $S_h(R(\mathbb{T}))$ for all $h \in \mathbb{H}$ except $h = c_i$ for $i \in [3] \setminus S$.

Lemma 4.14. *Let $\mathbb{H}$ be a connected, non-trivial strongly bipartite digraph, and let $\mathbb{T}$ be a good $\tau_{\mathbb{H}}$-tree. Let $\mathbb{T}'$ denote the tree $(\mathbb{T}^\tau)^u$. Then*

(1) The structures $R(\mathbb{T})$ and $\mathbb{P}_1 \times D(\mathbb{T})$ are homomorphically equivalent;
(2) $R(\mathbb{T})^\tau$, the underlying digraph of $R(\mathbb{T})$, is isomorphic to $G(\mathbb{T}')$.

Proof. Let D and U denote the colour classes of $\mathbb{T}'$. (1) First define a map $\phi : \mathbb{P}_1 \times D(\mathbb{T}) \to R(\mathbb{T})$ as follows:

$$\phi((i, f)) = \begin{cases} (0, f|_D), & \text{if } i = 0, \\ (1, f|_U), & \text{if } i = 1. \end{cases}$$

It is immediate that this map preserves arcs and colours.

Now we define a homomorphism $\psi : R(\mathbb{T}) \to \mathbb{P}_1 \times D(\mathbb{T})$. For every $u \in U$ let $e(u)$ be some fixed arc of $\mathbb{T}$ incident to u, and for each $d \in D$ let $e(d)$ be some fixed arc of $\mathbb{T}$ incident to d. Let $\psi((i, f)) = (i, f_i)$ where $f_i(t) = f(t)$ if either $t \in D$ and $i = 0$ or $t \in U$ and $i = 1$, and $f_i(t) = e(t)$ otherwise. It is easy to verify that ψ is a structure homomorphism. Indeed, for any arc $((0, f), (1, g))$ of $R(\mathbb{T})$, $((0, f_0), (1, g_1))$ is an arc of $\mathbb{P}_1 \times D(\mathbb{T})$, as for any arc (d, u) of $\mathbb{T}$, f and f_0 agree on d while g and g_1 agree on u. Further, for $(i, f) \in R(\mathbb{T})$, (i, f_i) belongs to all relations S_h that (i, f) does, as f is extended to f_i by mapping everything to arcs of $\mathbb{T}$.

(2) Define a map $\beta : R(\mathbb{T}) \to G(\mathbb{T}')$ as follows: $\beta((i, f)) = (i, E_f)$ where

$$E_f = \{e \in E(\mathbb{T}) : \exists t \in T, f(t) = e\}.$$

First we show that this map is well-defined, i.e. that the sets E_f satisfy the conditions of the definition of $G(\mathbb{T}')$. Let $i = 0$, and let $d \in D$ have degree greater than 1; since $\mathbb{T}$ is a good tree d is not coloured and hence $f(d)$ is an edge incident to d, and it is obviously the only such edge in E_f. The argument for $i = 1$ is identical. This last argument also shows that β is a one-to-one map. To show it is onto, let $(0, E) \in G(\mathbb{T}')$, and define a map f as follows: for $d \in D$, if there exists an edge $e \in E$ incident to d then it is unique, and let $f(d) = e$; if no edge incident to d belongs to E, then d is a leaf and hence has a colour h so we define $f(d) = \{h\}$. It is clear that $E_f = E$. We proceed similarly for $(1, g) \in G(\mathbb{T}')$. Now let $\beta((0, f)) = (0, E_f)$, $\beta((1, g)) = (1, E_g)$. Then by definition of β, there is an edge $e = (d, u)$ that lies in $E_f \cap E_g$ if and only if $f(d) = e = g(u)$, i.e. if and only if there is no arc from $(0, f)$ to $(1, g)$, thus β is a digraph isomorphism. □

4.3. **Proof of the existence of the NU.**

The following description of the building blocks for k-NU reflexive graphs is from [5].

Definition 4.15 ([5])**.** *Let $\mathbb{T}$ be a tree with colour classes D and U and edges $e_1, \ldots, e_m$. Define a graph $\mathbb{K}(\mathbb{T})$ as follows: its vertices are the tuples $(x_1, \ldots, x_m)$ such that*

(1) $x_i \in \{0, 1, 2\}$ for every $1 \le i \le m$;

(2) for each $d \in D$ of degree greater than 1, $x_i = 0$ for exactly one edge i incident to d;
(3) for each $u \in U$ of degree greater than 1, $x_i = 2$ for exactly one edge i incident to u.

Tuples $(x_1, \dots, x_m)$ and $(y_1, \dots, y_m)$ are adjacent if $|x_i - y_i| \leq 1$ for all i.

Definition 4.16. *Define two equivalence relations β and θ on the vertex set of $\mathbb{K}(\mathbb{T})$ as follows:*
let $(x_1, \dots, x_m)\, \beta\, (y_1, \dots, y_m)$ when both tuples have 0 in exactly the same coordinates, and $(x_1, \dots, x_m)\, \theta\, (y_1, \dots, y_m)$ if they have 2 in exactly the same coordinates. Define a digraph $\mathbb{Q}$ whose vertices are pairs $(0, U)$ with U a β-block and $(1, V)$ with V a θ-block; there is an arc from (i, U) to (j, V) if $(i, j) = (0, 1)$ and $U \cap V \neq \emptyset$. Notice that if $U \cap V \neq \emptyset$ the intersection consists of exactly one tuple.

Lemma 4.17. *$\mathbb{Q}$ is isomorphic to the subdigraph of $G(\mathbb{T})$ induced by non-isolated vertices.*

Proof. Indeed, let U be a β-block, let I be the set of positions i such that $x_i = 0$ for all tuples in U and let $X_U = \{e_i : i \in I\}$. Similarly if V is a θ-block let J be the set of positions j such that $x_j = 2$ for all tuples in T and let $Y_V = \{e_j : j \in J\}$. Define a map from the vertex set of $\mathbb{Q}$ to $G(\mathbb{T})$ that sends, for each β-block U, $(0, U)$ to $(0, X_U)$ and for each θ-block V, sends $(1, V)$ to $(1, Y_V)$. Clearly this a well-defined map. If $\alpha = ((0, X), (1, Y))$ is an arc of $G(\mathbb{T})$, define the tuple $(x_1, \dots, x_m)$ where $x_i = 0$ if $e_i \in X$, $x_i = 2$ if $e_i \in Y$ and $x_i = 1$ otherwise. It is clear that the β-block of this tuple maps to $(0, X)$ and that its θ-block maps to $(1, Y)$: notice in fact that this correspondence is a one-to-one correspondence between pairs (U, V) where $U \cap V \neq \emptyset$ and arcs of $G(\mathbb{T})$, and the result follows. $\square$

Corollary 4.18. *Let $k \geq 3$ and let $\mathbb{T}$ be a tree with at most $k - 1$ leaves. Then $G(\mathbb{T})$ is a k-NU digraph.*

Proof. In the proof of Lemma 4.2 of [5], we have the following k-NU polymorphism of $K(\mathbb{T})$: let $e_1, \dots, e_m$ denote the edges of $\mathbb{T}$. For each $1 \leq i \leq m$ define an integer c_i as follows: remove the edge e_i from $\mathbb{T}$ to obtain two connected components, exactly one of which contains a vertex $u \in U$ incident to e_i. Let c_i denote the number of leaves of $\mathbb{T}$ that are in this component. Clearly $1 \leq c_i \leq k - 2$. Viewing elements of $\mathbb{K}$ as columns for convenience of notation, define $f : (\mathbb{K}(\mathbb{T}))^k \to \mathbb{K}(\mathbb{T})$ by

$$f\left(\begin{bmatrix} x_{1,1} \\ \vdots \\ x_{m,1} \end{bmatrix} \cdots \begin{bmatrix} x_{1,k} \\ \vdots \\ x_{m,k} \end{bmatrix} \right) = r \begin{bmatrix} f_1(x_{1,1}, \dots, x_{1,k}) \\ \vdots \\ f_m(x_{m,1}, \dots, x_{m,k}) \end{bmatrix}$$

where f_i returns the $(c_i + 1)$-th smallest entry in row i, i.e. if $\{x_{i,1}, \dots, x_{i,k}\} = \{u_1, \dots, u_k\}$ where $u_i \leq u_j$ when $i \leq j$ then $f_i(x_{i,1}, \dots, x_{i,k}) = u_{c_i+1}$, and

$$r \begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix} = \begin{bmatrix} y_1 \\ \vdots \\ y_k \end{bmatrix}$$

where y_i is defined as follows: if $x_i = 0$, and e_i is incident to a vertex $d \in D$ such that there exists an edge e_j incident to d with $j < i$ and $x_j = 0$, or if $x_i = 2$, and

e_i is incident to a vertex $u \in U$ such that there exists an edge e_j incident to u with $j < i$ and $x_j = 2$, then let $y_i = 1$; otherwise let $y_i = x_i$.

Notice first that f preserves both equivalences β and θ since these depend only on the set of positions where the tuples are 0 and 2 respectively. Thus we may define a partial map on the digraph $\mathbb{Q}$ define above by setting $\phi((0, B_1), \ldots, (0, B_k)) = (0, B)$ where B is the β-block containing $f(x_1, \ldots, x_k)$ for any choice of $x_i \in B_i$; similarly let $\phi((1, T_1), \ldots, (1, T_k)) = (1, T)$ where T is the θ-block containing $f(x_1, \ldots, x_k)$ for any choice of $x_i \in T_i$. By Lemma 2.3 (3) and the preceding discussion, ϕ trivially extends to a full map from $G(\mathbb{T})^k$ to $G(\mathbb{T})$. Since f obeys the k-NU identities so does ϕ, and it remains only to verify that ϕ is arc-preserving, but this is immediate: if $x_i \in B_i \cap T_i$ for all $1 \leq i \leq k$ then $f(x_1, \ldots, x_k)$ belongs to $B \cap T$ where $\phi((0, B_1), \ldots, (0, B_k)) = (0, B)$ and $\phi((1, T_1), \ldots, (1, T_k)) = (1, T)$. $\square$

4.4. **Proof of Theorem 4.2.**

Proof of Theorem 4.2.

($\Leftarrow$) By Corollary 4.18 every digraph $G(\mathbb{T})$ admits a k-ary NU polymorphism, and it follows that their product and any retract of it also admits a k-NU polymorphism, see the proof of (2) $\Rightarrow$ (1) of Theorem 3.9 of [5].

($\Rightarrow$) If the strongly bipartite digraph $\mathbb{H}$ is k-NU then so is $\mathbb{H}^c$ by Lemma 4.3; by Lemma 3.1 it follows that $\mathbb{H}^c$ has finite duality, and hence has finitely many obstructions $\mathbb{T}_1, \ldots, \mathbb{T}_s$. We may assume these obstructions are critical $\tau_{\mathbb{H}}$-tree obstructions (this is easy, see for instance the proof of Theorem 3.9 in [5]). By Lemma 4.3 again, each of these $\tau_{\mathbb{H}}$-trees has at most $k-1$ coloured vertices. By Theorem 4.6 and the remark following it, $\mathbb{H}^c$ and the product of the duals of the $\mathbb{T}_i$ are homomorphically equivalent:

$$\mathbb{H}^c \longleftrightarrow \prod_{i=1}^{s} D(\mathbb{T}_i).$$

By taking the product on each side with s copies of the structure $\mathbb{P}_1 = \mathbb{P}_1(\mathbb{H})$ and applying Lemma 4.8, we obtain

$$\mathbb{H}^c \longleftrightarrow \mathbb{P}_1^s \times \mathbb{H}^c \longleftrightarrow \prod_{i=1}^{s} (\mathbb{P}_1 \times D(\mathbb{T}_i)).$$

Let $\mathbb{T}_{r+1}, \ldots, \mathbb{T}_s$ denote the good $\tau_{\mathbb{H}}$-trees in our set of obstructions (notice that, by Lemma 4.11, it follows that $\mathbb{T}_1, \ldots, \mathbb{T}_r$ are basic obstructions of one of the forms (A)-(D) in Lemma 4.9.) Let F_i denote the core of $\mathbb{P}_1 \times D(\mathbb{T}_i)$ for $1 \leq i \leq r$. By Lemma 4.14 (1) we have that

$$\mathbb{H}^c \longleftrightarrow \prod_{i=1}^{r} F_i \times \prod_{i=r+1}^{s} R(\mathbb{T}_i)$$

and since $\mathbb{H}^c$ is a core, we conclude that

$$\mathbb{H}^c \trianglelefteq \prod_{i=1}^{r} F_i \times \prod_{i=r+1}^{s} R(\mathbb{T}_i).$$

Now consider only the digraph part of these structures: since taking reducts commutes with products, it follows from Lemma 4.14 (2) that

$$\mathbb{H} \trianglelefteq \prod_{i=1}^{r} (F_i)^{\tau} \times \prod_{i=r+1}^{s} G(\mathbb{T}_i')$$

where $\mathbb{T}_i' = (\mathbb{T}_i^{\tau})^u$. It is immediate by Lemma 4.9 (and see Figures 3 and 6) that each digraph $(F_i)^{\tau}$ is a retract of a product of strongly bipartite paths; by Lemma 4.12 each such path is a connected component of a digraph $G(\mathbb{T})$ for some path $\mathbb{T}$; and thus we conclude that $\mathbb{H}$ is indeed a retract of a product of digraphs $G(\mathbb{T})$ where the $\mathbb{T}$ are trees with at most $k-1$ leaves.

□

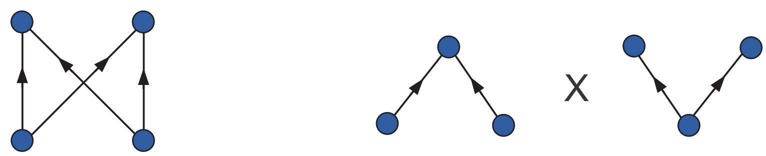

FIGURE 6. The digraph on the left is a retract of the product of the two paths.

5. IRREFLEXIVE k-NU GRAPHS

In this section, when we say *graph*, we mean an undirected, irreflexive graph. We now apply our work on strongly bipartite digraphs to the case of graphs. Recall that if $\mathbb{G}$ is a (irreflexive) digraph then $\mathbb{G}^u$ denotes the underlying (undirected) graph. Let $\mathbb{H}$ be a connected graph; if $\mathbb{H}$ is bipartite with colour classes D and U, let $\overrightarrow{\mathbb{H}}$ denote the digraph obtained from $\mathbb{H}$ by orienting every edge from D to U, i.e. $E(\overrightarrow{\mathbb{H}})$ consists of all pairs $(d, u) \in D \times U$ such that $(d, u) \in E(\mathbb{H})$. Clearly the digraph obtained this way is strongly bipartite.[2]

Lemma 5.1. *Let $\mathbb{H}$ be a connected graph.*

(1) If $\mathbb{H}$ admits an NU polymorphism then it is bipartite;

(2) if $\mathbb{H}$ is bipartite then $\mathbb{H}$ is k-NU if and only if $\overrightarrow{\mathbb{H}}$ is k-NU.

Proof.

(1) The result follows from the fact that any graph that admits a so-called *Taylor polymorphism* must be bipartite, see [3] (in fact, if a graph admits an NU polymorphism it dismantles to an edge and hence it must be bipartite [8]).

(2) Let D and U denote the colour classes of $\mathbb{H}$ such that $x \to y$ in $\overrightarrow{\mathbb{H}}$ implies that $x \in D$ and $y \in U$.

[2]As we remarked earlier, strictly speaking the digraph obtained is dependent on the choice of D and U.

Let f be a k-NU polymorphism of $\mathbb{H}$. Let $d \in D$ and $u \in U$; since $\mathbb{H}$ is connected there exists a large enough (even) integer N such that there exists a path of length N from any $x \in D$ to d, and similarly for every $y \in U$ to u. Since f is idempotent it follows that $f(D^k) \subseteq D$ and $f(U^k) \subseteq U$. It is then clear that f preserves the edge structure of $\overrightarrow{\mathbb{H}}$ also.

The converse is very similar to the proof of Proposition 2.4. Let f be a k-NU polymorphism of $\overrightarrow{\mathbb{H}}$. We define a map $F : H^k \to H$ as follows: let $x = (x_1, \dots, x_k) \in H^k$, and let $\Delta = \{i : x_i \in D\}$. If $|\Delta| \in \{0, k\}$ let $F(x) = f(x)$. If $|\Delta| \in \{1, k-1\}$ let $F(x) = x_i$ where i is the smallest index in the repeated block, i.e. if $x \in W \times \cdots \times W \times Z \times W \times \cdots \times W$ where $\{W, Z\} = \{D, U\}$ and Z appears in the j-th position, then i is the smallest index different from j. Otherwise let $F(x) = x_1$. It is easy to see, using Lemma 2.3, that F is a graph homomorphism, and it is clearly NU. □

Theorem 5.2. *Let $k \geq 3$ and $\mathbb{H}$ be a connected graph. Then the following are equivalent:*

(1) $\mathbb{H}$ is a k-NU graph;
(2) $\mathbb{H}$ is a retract of a product of finitely many graphs of the form $G(\mathbb{T})^u$ with $\mathbb{T}$ a tree with at most $k-1$ leaves.

Proof. ($\Leftarrow$) By Corollary 4.18 and Lemma 5.1, every graph $G(\mathbb{T})^u$ admits a k-ary NU polymorphism, and it follows that their product and any retract of it also admits a k-NU polymorphism.

($\Rightarrow$) We may assume that $\mathbb{H}$ is non-trivial. By Theorem 4.2 and Lemma 5.1 the digraph $\overrightarrow{\mathbb{H}}$ is a retract of a product of digraphs $G(\mathbb{T}_i)$, $i = 1, \dots, s$ where each $\mathbb{T}_i$ is a tree with at most $k-1$ leaves, i.e.

$$\overrightarrow{\mathbb{H}} \trianglelefteq \prod_{i=1}^{s} G(\mathbb{T}_i),$$

which easily implies that

$$\mathbb{H} = (\overrightarrow{\mathbb{H}})^u \trianglelefteq \left(\prod_{i=1}^{s} G(\mathbb{T}_i) \right)^u .$$

Claim. Let $\mathbb{C}$ be a non-trivial, connected graph and let $\mathbb{A}, \mathbb{B}$ be strongly bipartite digraphs. If $\mathbb{C} \trianglelefteq (\mathbb{A} \times \mathbb{B})^u$ then $\mathbb{C} \trianglelefteq \mathbb{A}^u \times \mathbb{B}^u$.

Proof of Claim. For $X \in \{\mathbb{A}, \mathbb{B}\}$ let D_X and U_X denote the colour classes of X such that $(v, w) \in E(X)$ implies $v \in D_X$ and $w \in U_X$. It is easy to see that (i) $\mathbb{A}^u \times \mathbb{B}^u$ is the disjoint union of its subgraphs $\mathbb{V}$ and $\mathbb{W}$ induced respectively by $(D_\mathbb{A} \times D_\mathbb{B}) \cup (U_\mathbb{A} \times U_\mathbb{B})$ and $(D_\mathbb{A} \times U_\mathbb{B}) \cup (U_\mathbb{A} \times D_\mathbb{B})$; (ii) the subgraph of $(\mathbb{A} \times \mathbb{B})^u$ induced by $(D_\mathbb{A} \times U_\mathbb{B}) \cup (U_\mathbb{A} \times D_\mathbb{B})$ consists of isolated vertices only; (iii) the subgraph $\mathbb{V}'$ of $(\mathbb{A} \times \mathbb{B})^u$ induced by $(D_\mathbb{A} \times D_\mathbb{B}) \cup (U_\mathbb{A} \times U_\mathbb{B})$ is isomorphic to $\mathbb{V}$. Since $\mathbb{C}$ is non-trivial and connected, its image $\mathbb{R}$ in $(\mathbb{A} \times \mathbb{B})^u$ under the coretraction must lie in $\mathbb{V}'$; we may then retract $\mathbb{V} \cong \mathbb{V}'$ onto $\mathbb{R}$ and $\mathbb{W}$ onto any edge of $\mathbb{R}$ and we're done.

It follows by the claim that

$$\mathbb{H} \trianglelefteq \prod_{i=1}^{s} G(\mathbb{T}_i)^u$$

and this concludes the proof. □

The following corollary, one of the motivations for our work, can be proved by combining results from [1], [3] and [7], but is an easy consequence of our results. A *majority* operation is a 3-ary NU operation.

Corollary 5.3. *Let $\mathbb{H}$ be a connected graph. Then the following are equivalent:*

(1) $\mathbb{H}$ admits a compatible majority operation;
(2) $\mathbb{H}$ is a retract of a product of paths.

Proof. Immediate by Lemma 4.12 and Theorem 5.2. □

6. The Decision Problem

We conclude with a few remarks concerning the problem of recognizing NU graphs.

Definition 6.1. *Let u, and v be vertices of a graph $\mathbb{H}$. We say u* dominates *v in $\mathbb{H}$ if every neighbour of v is also a neighbour of u. Let $\mathbb{H}$ be a bipartite graph with colour classes D and U, and let $\Delta = \{(x,x) : x \in H\}$. We'll say that $\mathbb{H}^2$* dismantles to the diagonal *if the following holds: there is a sequence of pairs $\{(x_i, y_i) : i = 1, \dots, m\}$ such that*

(1) $\{(x_1, y_1), \dots, (x_m, y_m)\} = D^2 \cup U^2 \setminus \Delta$;
(2) for each $1 \leq i \leq m$ the vertex (x_i, y_i) is dominated in the subgraph of H^2 induced by $\{(x_i, y_i), \dots, (x_m, y_m)\} \cup \Delta$.

In other words, $\mathbb{H}^2$ dismantles to the diagonal if, in the subgraph induced by $D^2 \cup U^2$, we can obtain the diagonal by successively removing dominated vertices. One can show that if $\mathbb{H}^2$ dismantles to the diagonal, the removal can be done greedily, in any order. In particular, one can determine in polynomial time if $\mathbb{H}^2$ dismantles to the diagonal.

Theorem 6.2. *There is a poly-time algorithm to recognize NU graphs. More precisely, the bipartite graph $\mathbb{H}$ is an NU graph if and only if $\mathbb{H}^2$ dismantles to the diagonal.*

Proof. As we noted in Lemma 5.1, the graph $\mathbb{H}$ is NU if and only if it is bipartite and the digraph $\overrightarrow{\mathbb{H}}$ is NU. By Theorem 3.1, this is equivalent to the problem $CSP(\overrightarrow{\mathbb{H}}^c)$ having finite duality, and since the structure $\overrightarrow{\mathbb{H}}^c$ is a core, we can determine if it has finite duality using the dismantling algorithm from Theorem 6.2 in [10]. Since pairs $(x, y) \in (D \times U) \cup (U \times D)$ have neither in- nor out-neighbours in $\mathbb{H}^2$ and hence are dominated by any other pair, it is easy to see that the algorithm will succeed precisely when $\mathbb{H}^2$ dismantles to the diagonal, as defined above. □

The following is a special case of an algorithm due independently to Mároti and to Barto and Kozik.

Theorem 6.3 ([2],[14])**.** *For each fixed $k \geq 3$, there is a poly-time algorithm to recognize k-NU graphs. In fact, the algorithm produces a k-NU polymorphism if there is one.*

Proof. Let $\mathbb{H}$ be a graph, and as before let $\tau_{\mathbb{H}}$ denote the signature of the structure $\mathbb{H}^c$. Let $\mathbb{G}$ be the following $\tau_{\mathbb{H}}$-structure: its underlying τ-structure is the graph $\mathbb{H}^k$, and we colour each vertex of the form $(x, \dots, x, y, x, \dots, x)$ with the colour x. Obviously $\mathbb{H}$ is a k-NU graph precisely if $\mathbb{G}$ admits a homomorphism to $\mathbb{H}^c$.

It is known that if $\mathbb{H}$ admits an NU polymorphism, then the problem $CSP(\mathbb{H}^c)$ has bounded width, i.e. can be solved (in polynomial time) by local consistency checking [6]. This algorithm has the property that, whatever the nature of the target, if it returns a NO answer, it is always correct. So proceed as follows: run the algorithm on instance $\mathbb{G}$; if it returns NO then $\mathbb{H}$ is not k-NU and we stop. Otherwise, choose an uncoloured vertex of $\mathbb{G}$, colour it and run the algorithm. If the algorithm returns NO, repeat with a different colour until either (a) all colours give NO and hence $\mathbb{H}$ is not k-NU or (b) some colour returns YES. In that case keep the colour, choose another uncoloured vertex and repeat. Note that the total number of calls we make to the bounded width algorithm is $\mathcal{O}(|H|^{k+1})$. Unless we get a NO answer for every colour assigned to a vertex, in which case $\mathbb{H}$ is not k-NU and we are done, we shall obtain a fully defined NU map from H^k to H. It now suffices to check that the map is a homomorphism; if it is, $\mathbb{H}^k$ is k-NU, and otherwise it is not (if $\mathbb{H}$ were NU the algorithm could not return a false positive.) $\square$

References

[1] H.-J. Bandelt, Graphs with edge-preserving majority functions, *Discrete Math.* **103** (1992) no. 1, 1-5.
[2] L. Barto, M. Kozik, personal communication.
[3] A. Bulatov, H-coloring dichotomy revisited. *Theoret. Comput. Sci.* **349**, no. 1, 31-39, 2005.
[4] A. Bulatov, A. Krokhin, B. Larose. Dualities for constraint satisfaction problems. In *Complexity of Constraints*, volume 5250 of *LNCS*, pages 93–124. 2008. Discrete Math. 306 (2006), no. 15, 17581764.
[5] T. Feder, P. Hell, B. Larose, C. Loten, M. Siggers, C. Tardif, Graphs admitting k-NU operations. Part 1: the reflexive case. 26 pages, submitted, 2012.
[6] T. Feder, M. Y. Vardi, Monotone monadic SNP and constraint satisfaction, in *Proceedings of the 25rd Annual ACM Symposium on Theory of Computing (STOC)*, San Diego, California, (1993), 612–622.
[7] P. Hell, *Rétractions de graphes*, Ph.D. thesis, Université de Montréal (1972).
[8] B. Larose, Graphs admitting an NU polymorphism, manuscript, 1998.
[9] B. Larose A. Lemaître, List-Homomorphism Problems and Arc-Consistency, 20 pages, submitted, 2012.
[10] B. Larose, C. Loten, C. Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3(4), 2007.
[11] B. Larose, C. Loten, L. Zádori, A polynomial-time algorithm to recognise near-unanimity graphs, *J. Algorithms*, **55** no. 2, 177–191, 2005.
[12] B. Larose, M. Siggers, Absolute retracts and NU polymorphisms, in preparation.
[13] B. Larose, L. Zádori, Algebraic properties and dismantlability of finite posets, *Discrete Math.*, **163**, 89–99, 1997.
[14] M. Mároti, personal communication.
[15] J. Nešetřil and C. Tardif. Duality theorems for finite structures (characterising gaps and good characterisations). *Journal of Combininatorial Theory, Ser.B*, 80:80–97, 2000.
[16] J. Nešetřil and C. Tardif. Short answers to exponentially long questions: extremal aspects of homomorphism duality. *SIAM Journal on Discrete Mathematics*, 19(4):914–920, 2005.
[17] A. Szendrei. *Clones in Universal Algebra*, volume 99 of *Seminaires de Mathematiques Superieures*. Université de Montréal, 1986.
[18] L. Zádori, Relational sets and categorical equivalence of algebras, *Int. J. Algebra Comput.* **7**, No.5, 561–576, 1997.

E-mail address: `tomas@theory.stanford.edu`

School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada V5A 1S6
E-mail address: pavol@cs.sfu.ca

Department of Mathematics and Statistics, Concordia University, 1455 de Maisonneuve West, Montréal, Qc, Canada, H3G 1M8
E-mail address: larose@mathstat.concordia.ca

College of Natural Sciences, Kyungpook National University, Daegu 702-701, South Korea
E-mail address: mhsiggers@knu.ac.kr

Department of Mathematics and Computer Science, Royal Military College of Canada, PO Box 17000 Station "Forces", Kingston, Ontario, Canada, K7K 7B4
E-mail address: Claude.Tardif@rmc.ca