On Maltsev digraphs[*]

Catarina Carvalho

School of Physics, Astronomy and Mathematics
University of Hertfordshire, U.K.
`c.carvalho2@herts.ac.uk`

László Egri[†]

School of Computer Science

McGill University, Canada
`laszlo.egri@mail.mcgill.ca`

Marcel Jackson   Todd Niven[‡]

School of Engineering and Mathematical Sciences
La Trobe University, Australia
`{m.g.jackson, t.niven}@latrobe.edu.au`

Abstract

We study digraphs preserved by a Maltsev operation: *Maltsev digraphs*. We show that these digraphs retract either onto a directed path or to the disjoint union of directed cycles, showing in this way that the constraint satisfaction problem for Maltsev digraphs is in logspace, L. We then generalize results from Kazda (2011) to show that a Maltsev digraph is preserved not only by a majority operation, but by a class of other operations (e.g., minority, Pixley) and obtain a $O(|V_G|^4)$-time algorithm to recognize Maltsev digraphs. We also prove analogous results for digraphs preserved by conservative Maltsev operations which we use to establish that the list homomorphism problem for Maltsev digraphs is in L. We then give a polynomial time characterisation of Maltsev digraphs admitting a conservative 2-semilattice operation. Finally, we give a simple inductive construction of directed acyclic digraphs preserved by a Maltsev operation, and relate them with series parallel digraphs.

1   Introduction

The study of relational structures and, in particular, digraphs preserved by certain operations from universal algebra became extremely important during the last decade. The

---

main driving force behind this is the algebraic constraint satisfaction problem (CSP) dichotomy conjecture, which states that a constraint satisfaction problem CSP($\mathbf{B}$) is tractable if the relational structure $\mathbf{B}$ is preserved by a weak-near-unanimity (weak-NU) operation, and is NP-complete otherwise [5, 6, 23]. Generalizing the dichotomy theorem of Hell and Nešetřil [19], the conjecture has been established for digraphs with no sinks and no sources by showing that the tractable digraphs are very structured, in fact, they retract onto a disjoint union of directed cycles [3]. Other results relating the complexity of CSPs on digraphs to the existence of operations that preserve the digraph can be found, for example, in [2, 1].

Once the tractability of a CSP is established, one also wishes to know the fine-grained complexity of that CSP, that is, for which structures $\mathbf{B}$ is CSP($\mathbf{B}$) in some subclass of P, such as L or NL? To establish the membership of a CSP in a complexity class inside P it is important to study the structure of relational structures (and digraphs) preserved by operations which are more "restrictive" than weak-NU operations, that is, operations that imply the presence of a weak-NU operation. Two important results in this direction, which we will invoke, are that if a relational structure $\mathbf{B}$ is preserved by a majority operation, then the corresponding CSP($\mathbf{B}$) is in the complexity class NL [11]; and if CSP($\mathbf{B}$) is definable in Datalog and $\mathbf{B}$ is preserved by a Maltsev operation, then CSP($\mathbf{B}$) is in L [12, 14].

We study the structure of finite digraphs preserved by a Maltsev operation; we call such graphs *Maltsev digraphs*. We show that Maltsev digraphs retract either onto the disjoint union of directed cycles or to a directed path. This gives a direct proof that the corresponding CSP can be expressed with a symmetric Datalog program of width $(2,3)$, and therefore the CSP is in L.[1] Kazda [22] showed that every Maltsev digraph also admits a majority operation. Subsequent work by Bulin, Delic, Jackson and Niven [8] showed that there is little opportunity for substantial extension of this result: the possible polymorphism behaviour of digraphs is mostly the same as for general CSPs. Nevertheless, we are able to revisit Kazda's argument, and use it to show that Maltsev digraphs also admit a minority polymorphism, and show how for Maltsev digraphs, a range of polymorphism properties are typically determined by the properties of an associated disjoint union of cycles. We also extend the results to the conservative setting, showing for example that a conservative Maltsev digraph is preserved by a class of conservative polymorphisms.

Kazda's article [22] concludes with an invitation to provide a more direct characterisation of Maltsev digraphs. We provide this by giving a graph-theoretic characterisation, and providing a $O(|V_G|^4)$ algorithm for recognising Maltsev digraphs. A similar characterisation is given for the conservative case; both conditions generalise the rectangularity

---

[1]Membership of these CSPs in symmetric Datalog was independently shown by Kazda [22]. However, his proof is rather indirect, and the width of his symmetric Datalog program is not explicitly stated. Nevertheless, it seems that his approach gives a width guarantee that grows with the size of the target structure: he uses the main result of [11], which gives a width guarantee only for the linear Datalog program that solves the CSP, and then he invokes [12] to argue that the CSP is actually in symmetric Datalog. The width in [12] also depends on the size of the target. The width of the symmetric Datalog program in the present paper is $(2,3)$.

property of [4].

We apply our results to the list homomorphism problem, LHOM, for directed graphs. We show that LHOM for Maltsev digraphs is in L. The complexity of LHOM for undirected graphs is completely understood [16], and for directed graphs there is a P vs. NP dichotomy is [20], and a recently announced L vs. NL-hard dichotomy [15].

Next we show that a digraph preserved by a Maltsev operation is also preserved by a conservative 2-semilattice operation if and only if the digraph satisfies a certain combinatorial property. We note that if a structure is preserved by a 2-semilattice operation, then CSP($\mathbf{B}$) is in Datalog and therefore in P [21]. We also characterise digraphs (without the assumption of a Maltsev polymorphism) preserved by a conservative 2-semilattice operation and show that these digraphs can be recognised in NL.

In Section 5 we relate directed acyclic graphs (DAGs) preserved by Maltsev, and conservative Maltsev polymorphisms with series parallel digraphs. We characterize these in the spirit of the existing characterizations of quasi-series parallel and complete bipartite composite digraphs. This characterization allow us to obtain a $O(|V_G|^3)$-time algorithm to recognize if an acyclic digraph is Maltsev. Finally, an inductive construction of directed acyclic graphs preserved by a Maltsev operation is given. The main motivation behind this construction is that we suspect that extending this construction to n-permutable digraphs (2-permutable digraphs are precisely the Maltsev digraphs [18]) could make progress towards identifying all list homomorphism problems for digraphs in L. We note that in [16], an inductive construction of conservative n-permutable digraphs is key to the identification of all graphs whose LHOM is in L.

2 Preliminaries

2.1 Algebra

We describe the algebraic definitions for digraphs, however, note that these definitions are straightforward to generalize to relational structures. Let $G = (V_G, E_G)$ and $H = (V_H, E_H)$ be digraphs. A *homomorphism* from G to H is a map f from the set of vertices of G, V_G, to the set of vertices of H, V_H, such that for every edge $(u, v) \in E_G$ we have that $(f(u), f(v))$ is an edge in H, that is $(f(u), f(v)) \in E_H$. A digraph G is called a *core* if every homomorphism from G to itself is an automorphism, that is, a permutation on V_G. Let G' be a subgraph of G. We say that G retracts onto G' if there is a homomorphism $h : G \to G'$ such that h is the identity map on G'. For a finite digraph H, we can then define CSP(H) as the class of all finite digraphs that admit a homomorphism to H

An n-ary operation on a set A is a function $f : A^n \to A$. Given a digraph G and an n-ary operation f on V_G, we say that f *preserves* G, or that f is a *polymorphism* of G, if for any n edges $(u_1, v_1), \ldots, (u_n, v_n) \in E_G$ (not necessarily distinct), the pair $(f(u_1, \ldots, u_n), f(v_1, \ldots, v_n)) \in E_G$. For an n-ary operation f, we write $f(x_1, \ldots, x_n) \approx f(y_1, \ldots, y_n)$ if $f(x_1, \ldots, x_n) = f(y_1, \ldots, y_n)$ for all possible values of the x_i, y_i, $i = 1, \ldots, r$. A ternary operation m is *Maltsev* if it satisfies $m(x, x, y) \approx m(y, x, x) \approx y$, *Pixley* if it satisfies $m(x, x, y) \approx m(y, x, x) \approx m(y, x, y) \approx y$, *majority* if it satis-

fies $m(x, x, y) \approx m(x, y, x) \approx m(y, x, x) \approx x$, and *minority* if it satisfies $m(x, y, x) \approx m(x, y, x) \approx m(y, x, x) \approx y$. A binary operation $*$ is *2-semilattice* if it satisfies $x * x \approx x$, $x * y \approx y * x \approx x * (x * y)$.

2.2   Graph Theory

Since all graphs in the paper are directed, we use the terms graph and digraph interchangeably. Throughout, all digraphs are assumed to be finite.

For a natural number n we write $[n] = \{1, 2, \ldots, n\}$. An *oriented path* is a sequence of, not necessarily distinct, vertices $v_1, \ldots, v_n$ such that for every $i \in [n-1]$, either (v_i, v_{i+1}) (*a forward edge*) or (v_{i+1}, v_i) (*a backward edge*) is an edge. We use the terms path and oriented path interchangeably. A *cycle* is an oriented path with starting point v_1 and endpoint v_m such that either (v_m, v_1) or (v_1, v_m) is an edge. The *net length* of a path P, $net(P)$, is the number of forward edges minus the number of backward edges in P. A *dipath* (*reverse dipath*) is a sequence of, not necessarily distinct, vertices $v_1, \ldots, v_n$ such that for every $i \in [n-1]$, (v_i, v_{i+1}) $((v_{i+1}, v_i))$ is an edge. A *directed cycle* is a dipath $v_1, \ldots, v_n$ such that (v_n, v_1) is also an edge. For a (reverse) dipath P, we let $len(P)$ denote the number of edges in P. We use the term *simple* dipath or (directed) cycle to indicate that all vertices of the dipath or (directed) cycle are distinct.

A *component* of digraph G is a maximal subgraph, H, of G such that for every pair of vertices $u, v \in V_H$, there is an oriented path from u to v. A digraph with one component is said to be *connected*. A digraph is a *directed acyclic graph* (DAG) if it contains no directed cycles. A DAG G is *layered* if there exists $q \in \mathbb{N}$ such that the vertices of G can be partitioned into q levels $L_0, \ldots, L_{q-1}$, such that any edge of G goes from L_i to L_{i+1}, for some $i = 0, \ldots, q-2$.

Let G be a digraph, and x a vertex of G. We define $x^{+1} = \{y \in V_G : (x, y) \in E_G\}$, and $x^{-1} = \{y \in V_G : (y, x) \in E_G\}$. We call a vertex v a *source* if $v^{-1} = \varnothing$, and a *sink* if $v^{+1} = \varnothing$. If u and v are vertices of G, $u \xrightarrow{k} v$ denotes the existence of a dipath from u to v of length k; $u \to v$ denotes $u \xrightarrow{1} v$. For each vertex x of G and each $k \geqslant 1$, we define $x^{+k} = \{y \in V_G : x \xrightarrow{k} y\}$, $x^{-k} = \{y \in V_G : y \xrightarrow{k} x\}$, and $x^{\pm k} = \{y \in V_G : x \xrightarrow{k} y \text{ or } y \xrightarrow{k} x\}$.

The kth *power* of a graph, G, is a graph, G^k, with the same set of vertices as G and an edge between two vertices u, v if and only if $u \xrightarrow{k} v$.

3   Retracts of Maltsev digraphs

Definition 3.1. [totally rectangular] A digraph G is *k-rectangular* if the following implication holds for all vertices x, y, u, v:

$$x \xrightarrow{k} u \ \& \ y \xrightarrow{k} u \ \& \ y \xrightarrow{k} v \Rightarrow x \xrightarrow{k} v.$$

A digraph is *rectangular* if it is 1-rectangular, and *totally rectangular* if it is k-rectangular for every $k \in \mathbb{N}$.

It is not hard to verify that a Maltsev digraph must be totally rectangular, and in Section 4 (see Corollary 4.11) we show that indeed the two properties are equivalent.

Example 3.2. *The digraph in Fig. 1 is rectangular but not 2-rectangular. While the digraph in Fig. 2 is totally rectangular.*
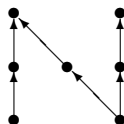

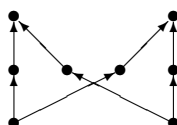Figure 1: A rectangular digraph that is not 2-rectangular.

Figure 2: A Maltsev digraph with no conservative Maltsev polymorphism.

We now state the main result of this section.

Theorem 3.3. *Let G be a totally rectangular digraph. If G is acyclic then G retracts onto a simple dipath. Otherwise G retracts onto the disjoint union of simple directed cycles.*

The proof of Theorem 3.3 is a direct consequence of Lemma 3.10 below. We begin with some definitions and simple observations.

Lemma 3.4. *Let G be a digraph. Then G is layered if and only if for every pair of vertices u, v in G, and any pair of oriented paths P and Q from u to v, it holds that $net(P) = net(Q)$.*

Definition 3.5. Let G be a digraph that contains a directed cycle. Let C be a shortest directed cycle in G and assume it has length m. We say that G is *inconsistent* if there exist two vertices u, v in G such that, there are two different oriented paths of net lengths ℓ_1 and ℓ_2 from u to v such that $\ell_1 \not\equiv \ell_2 \mod m$. Otherwise we say that G is *consistent*.

Proposition 3.6. *Let G be a connected digraph that contains a directed cycle. Let C be a shortest directed cycle in G. Then G retracts onto C if and only if G is consistent.*

Proof. Assume that C has length m. If G is consistent, then choose an arbitrary vertex of C and label it with 0. Call this vertex s. For each vertex $v \in G \backslash s$, find an oriented path P from s to v. Assume P has net length ℓ. Label v with the number $\ell \mod m$. Since G is consistent, this labeling is well defined. It is easy to see that the mapping that sends every vertex with label i to the vertex of C that has label i is a retraction of G onto C.

Conversely, assume that G retracts to C through a homomorphism h. Label every vertex v with $h(v)$. It is not difficult to see that for any two vertices u and v of G such that there are two different (not necessarily simple) oriented paths of net lengths ℓ_1 and ℓ_2 from u to v, it must be that $\ell_1 \equiv \ell_2 \mod m$. $\square$

The following lemma relates total rectangularity with lengths of paths in a digraph.

Lemma 3.7. *Let G be a totally rectangular digraph and u, v be vertices in G. Let P and Q be two dipaths in G from u to v, such that $len(P) > len(Q)$. Set $k = len(P)$, $\ell = len(Q)$, and $d = k - \ell$. Then one of the following two cases occurs:*

1. *If $2\ell > k$, then G contains vertices u', v' and dipaths P', Q' from u' to v' with the following property: $len(P') = \ell$, $len(Q') = 2\ell - k$, and $len(P') - len(Q') = d$;*

2. *If $2\ell \leqslant k$, then G contains a directed cycle of length d.*

Proof. See Fig. 3. In the first case, let u' be the vertex of P such that the subpath $P_{u'v}$ of P from u' to v has length ℓ. Let v' be the vertex of P such that the subpath $P_{uv'}$ of P from u to v' has length ℓ. Applying the ℓ-rectangularity of G to $P_{u'v}$, Q, and $P_{uv'}$, we obtain the desired dipath P' with $len(P') = \ell$. The other required dipath Q' is the subpath of P from u' to v'. Since $k = 2\ell - len(Q')$ we have that $len(Q') = 2\ell - k$, and $len(P') - len(Q') = \ell - (2\ell - k) = k - \ell = d$.

In the second case, the two paths P' and Q' form a cycle of length $\ell + (k - 2\ell) = d$ because $2\ell \leqslant k$. $\square$
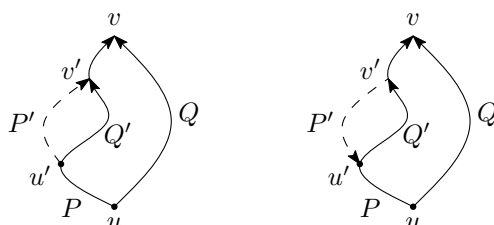

Figure 3: Case 1 (left) and Case 2 (right) of Lemma 3.7.

Definition 3.8. Let P be an oriented path. Let $P_1, \ldots, P_n$ be all maximal (with respect to length) directed subpaths of P, such that P_i and P_{i+1}, $i \in [n-1]$ share a common vertex. We call $P_1, \ldots, P_n$ a *directed path decomposition of P*. We can obtain a *directed path decomposition for oriented cycles* in a similar way. Note that if an oriented cycle has two maximal directed subpaths, then these two subpaths share two vertices.

Lemma 3.9. *Let G be a totally rectangular digraph and $u, v \in V_G$. Let P_1 and P_2 be oriented paths in G from u to v. Assume that $net(P_1) > net(P_2)$, and set $d = net(P_1) - net(P_2)$. Then there are vertices $s, t \in V_G$ and dipaths Q_1 and Q_2 in G from s to t, such that $len(Q_1) - len(Q_2) = d$.*

Proof. Consider the oriented cycle C formed by P_1 and P_2. Let $R_0, \ldots, R_{n-1}$ be a directed path decomposition of C, and assume that $n \geqslant 2$. Find a shortest segment R_i in the decomposition and assume it has length k. See Fig. 4. Then both R_{i-1} and R_{i+1} have length at least k, where $i - 1$ and $i + 1$ are considered modulo n. Assume without loss of generality that R_{i-1} and R_i have the same endpoint b, and R_i and R_{i+1} have the same
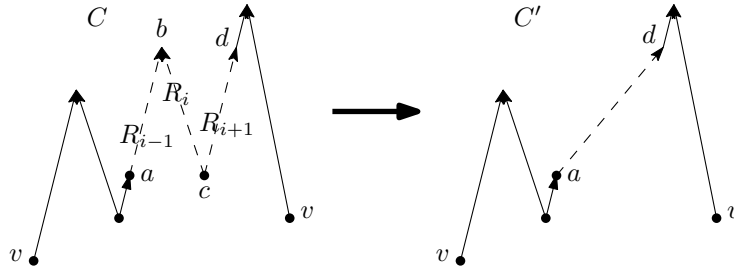
Figure 4: Making a shortcut.

starting point c. Let a be the vertex of R_{i-1} at distance k from b (going backward on R_{i-1}), and d be the vertex of R_{i+1} at distance k from c. Then using the k-rectangularity of G, we can make a *shortcut* from a to d to obtain a new oriented cycle C' in G. We repeat this procedure until we obtain a cycle which has a directed path decomposition consisting only of two directed paths, Q_1 and Q_2.

Assume without loss of generality that Q_1 is longer than Q_2. To see that $net(Q_1) - net(Q_2) = d$, observe that d is the difference of the forward and backward edges in C. Applying the above "shortcutting" procedure to C to obtain C' does not change the difference of the forward and backward edges in C. To see this, observe that we take out $2k$ forward edges and k backward edges from C, and then we add back k forward edges to obtain C'. $\qquad\square$

Lemma 3.10. *Let G be a connected totally rectangular digraph. If G is a DAG then G retracts onto a simple dipath. Otherwise G retracts onto a simple directed cycle.*

Proof. Assume first that G is a DAG. We claim that G must be layered. Assume, for a contradiction, that G is not layered. By Lemma 3.4, there exist $u, v \in V_G$ and oriented paths P and Q from u to v, such that $net(P) \neq net(Q)$. Using Lemma 3.9, we can assume that P and Q are dipaths of different length. Now we repeatedly apply Case 1 of Lemma 3.7 as long as it is possible, and then applying Case 2 yields a cycle, a contradiction. So G is layered.

Assume that G has levels $L_0, \ldots, L_{q-1}$. Fix vertices $s \in L_0$ and $t \in L_{q-1}$, and let O be any oriented path from s to t (such a path exists because G is connected). Applying the total rectangularity of G to appropriate subpaths of O, it is easy to see that there exists a dipath D of length $q-1$ from s to t in G. Clearly, G retracts onto D.

Suppose that G contains a directed cycle. By Proposition 3.6, it is enough to show that G is consistent. Assume this is not the case. Let C be a shortest directed cycle in G, and assume it has length m. Because G is inconsistent, we can find vertices $u, v \in V_G$ and oriented paths P_1 and P_2 from u to v, such that $net(P_1) \not\equiv net(P_2) \mod m$. Set $\ell_1 = net(P_1)$ and $\ell_2 = net(P_2)$. Assume w.l.o.g. that $\ell_1 > \ell_2$, and that u is a vertex of C. Note that if u is not a vertex of C, then we fix a vertex c of C and find any oriented path S from c to u. Then attaching S to P_1 and P_2 at vertex u gives us the desired oriented paths. Furthermore, we can assume that $\ell_1 - \ell_2 = d < m$, because if not, we can add

C-loops from u to u to P_2 to increase its length by a multiple of m, until $\ell_1 - \ell_2 < m$. Using Lemma 3.9 we obtain directed paths Q_1 and Q_2 such that $len(Q_1) - len(Q_2) = d$, and then, by applying Lemma 3.7, we obtain a cycle of length d in G, a contradiction. $\quad\square$

By Lemma 3.10, each connected component of G retracts either onto a simple dipath or to a simple directed cycle. The trivial observation that a dipath homomorphically maps to a cycle completes the proof of Theorem 3.3.

4 Characterisations, Polymorphisms and Algorithms

4.1 Rectangular Characterisations and Other Polymorphisms

In this section we generalise a technique of Kazda [22] to characterise digraphs that admit Maltsev and conservative Maltsev polymorphisms as those which are totally rectangular and universally rectangular respectively and to provide polynomial time algorithms for recognising the relevant properties. Furthermore, we show that Maltsev digraphs also admit many other polymorphisms, and under certain conditions, they also admit conservative 2-semilattice polymorphisms.

Definition 4.1. [conservatively k-rectangular, universally rectangular] We say that a graph is *conservatively k-rectangular* if it satisfies the following sentence:

$$\left.\begin{array}{lllll} x \to & x_1 \to & \cdots \to & x_{k-1} \to & u \\ y \to & y_1 \to & \cdots \to & y_{k-1} \to & u \\ y \to & z_1 \to & \cdots \to & z_{k-1} \to & v \end{array}\right\} \Rightarrow \begin{cases} \text{There is a path } x \to w_1 \to \\ \cdots \to w_{k-1} \to v \text{ with } w_i \in \\ \{x_i, y_i, z_i\} \text{ for each } i. \end{cases} \quad (1)$$

A graph that is conservatively k-rectangular for all $k \geqslant 1$ will be called *universally rectangular*.

Example 4.2. *The digraph in Fig. 2 is conservatively 1-rectangular but not conservatively 2-rectangular. The digraph in Fig. 5 is universally rectangular.*
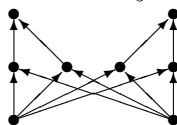
Figure 5: A universally rectangular digraph.

Definition 4.3. Let G be a digraph. Define the binary relation R^- on V_G by $x \, R^- \, y$ if $x^{-1} \cap y^{-1} \neq \varnothing$. The dual relation R^+ is defined by $x \, R^+ \, y$ if $x^{+1} \cap y^{+1} \neq \varnothing$.

In the presence of rectangularity, the relation R^+ is an equivalence relation on the set $\{x \in V_G : x^+ \neq \varnothing\}$, the set of vertices of G that are not sinks. Additionally, the relation R^- is an equivalence relation on the set $\{x \in V_G : x^- \neq \varnothing\}$, the set of vertices of G that are not sources. So it makes sense to consider the respective factor graphs, this was observed in [22]. We use the notation G/R^+ to denote the graph on the R^+-classes

of G. Given R^+-classes A, B, we write $A \to B$ if there is some $a \in A$ and $b \in B$ with $a \to b$. Similarly, G/R^- denotes the same construction, but using the relation R^-. Note that G/R^+ is not strictly an actual graph quotient of G, only a quotient of an induced subgraph of G. Nevertheless, we sometimes refer to it as "the quotient of G by R^+".

Lemma 4.4. *Let G be a rectangular digraph and $k > 1$.*

1. *G is ℓ-rectangular for all $\ell = 1, \ldots, k$ if and only if G/R^+ is ℓ-rectangular for all $\ell = 1, \ldots, k-1$.*

2. *If G is conservatively ℓ-rectangular for all $\ell = 1, \ldots, k$ then G/R^+ is conservatively ℓ-rectangular for all $\ell = 1, \ldots, k-1$.*

Proof. The rectangularity of G ensures that G/R^+ is well defined. Consider some $\ell \leqslant k - 1$. We show that G is $(\ell + 1)$-rectangular if and only if G/R^+ is ℓ-rectangular.

Consider vertices u, v, x, y of G such that $x/R^+ \xrightarrow{\ell} u/R^+$, $y/R^+ \xrightarrow{\ell} u/R^+$ and $y/R^+ \xrightarrow{\ell} v/R^+$ in G/R^+. This is equivalent to both the following properties holding:

(i) there are $u_x, u_y \in u/R^+$ and $v_y \in v/R^+$ such that $x \xrightarrow{\ell} u_x$, $y \xrightarrow{\ell} u_x$ and $y \xrightarrow{\ell} v_y$ in G,

(ii) there are u', v' with $u \to u'$ and $v \to v'$ in G.

Now (i) and (ii) combined are equivalent to $x \xrightarrow{\ell+1} u'$, $y \xrightarrow{\ell+1} u'$ and $y \xrightarrow{\ell+1} v'$ in G, a general instance of the premise of $(\ell + 1)$-rectangularity of G. If G is $(\ell + 1)$-rectangular then $x \xrightarrow{\ell+1} v'$ and as $v_x R^+ v$ we have that $x/R^+ \xrightarrow{\ell} v/R^+$ in G/R^+ showing that G/R^+ is ℓ-rectangular. Conversely, if G/R^+ is ℓ-rectangular then $x/R^+ \xrightarrow{\ell} v/R^+$ gives $x \xrightarrow{\ell+1} v'$, showing that G is $(\ell + 1)$-rectangular.

To prove the second statement, note that an instance of the premise in Definition 4.1 (1) in G/R^+ comes from an instance of the premise of conservative $(k+1)$-rectangularity in G. Then the path in G witnessing the conclusion of conservative $(k + 1)$-rectangularity provides a corresponding path in G/R^+ using the R^+-classes of the vertices in the path. $\square$

For a totally rectangular graph G, define $G_0 = G$ and $G_{i+1} = G_i/R^+$, $i \geqslant 1$. From Lemma 4.4 it follows that G_i is defined for all positive integers i, and eventually G_i will either be empty or a disjoint union of directed cycles (the only situations that R^+ can be trivial). We define $G_\infty = G_k$, where k is such that $G_k = G_{k+1}$ (that is, G_∞ is either empty or a disjoint union of directed cycles).

The next lemma is obtained by applying the Maltsev property to the columns of the premise of (1).

Lemma 4.5. *Let G be a digraph.*

1. *If G has a Maltsev polymorphism, then G is totally rectangular.*

2. *If G has a conservative Maltsev polymorphism, then G is universally rectangular.*

Proof. We prove the second statement, the first being similar but easier. Consider an instance of the premise of k-rectangularity in some graph G with a conservative Maltsev polymorphism p. As p is a polymorphism, applying p to the columns yields a path $x = p(x, y, y) \to p(x_1, y_1, z_1) \to \cdots \to p(x_{k-1}, y_{k-1}, z_{k-1}) \to p(u, u, v) = v$. But because p is conservative we have $p(x_i, y_i, z_i) \in \{x_i, y_i, z_i\}$ so that the path witnesses conservative k-rectangularity. $\square$

The following lemma is key to the proof of Theorem 4.10.

Lemma 4.6. *Let a, b be vertices in a totally rectangular digraph G satisfying conservative 2-rectangularity and assume that neither a nor b is a source or sink. If $a/R^+ \cap b/R^-$ is nonempty then either $b \in a/R^+$ or $a \in b/R^-$.*

Proof. Let $c \in a/R^+ \cap b/R^-$. There are vertices e, f, g, h such that $\{a, c\} \subseteq e^{-1}$, $b \in f^{-1}$, $a \in g^{+1}$ and $\{c, b\} \in h^{+1}$. However G is conservatively 2-rectangular so that there is either an edge from at least one of a, c to f or there is an edge from g to at least one of $\{b, c\}$. Then 1-rectangularity shows that either there is an edge from a to f or from g to b. $\square$

To give details of the proof of Theorem 4.10, it is useful to recall the following result, taken directly from [22].

Lemma 4.7. *Let G be a rectangular digraph. Then the following hold:*

1. *R^+ is an equivalence relation on $G \setminus S^+(G)$, where $S^+(G)$ is the set of sinks of G.*

2. *R^- is an equivalence relation on $G \setminus S^-(G)$, where $S^-(G)$ is the set of sources of G.*

3. *Whenever xR^+y, we have $x^{+1} = y^{+1}$ and x^{+1} is an equivalence class of R^-.*

4. *Whenever xR^-y, we have $x^{-1} = y^{-1}$ and x^{-1} is an equivalence class of R^+.*

5. *The mapping $\phi : X \mapsto X^{+1}$ is a bijection from the set of equivalence classes of R^+ to the set of equivalence classes of R^- and ϕ is an isomorphism from G/R^+ to G/R^-.*

We now give a technical lemma, present in [22]:

Lemma 4.8. *Let t^+ be a polymorphism on the graph G/R^+, and let t^- be the n-ary polymorphism defined by*

$$t^{-1}(x_1/R^-, \ldots, x_n/R^-) := \phi(t^+(\phi^{-1}(x_1/R^-), \ldots, \phi^{-1}(x_1/R^-))).$$

If t is an n-ary function on the vertices of G satisfies the following conditions then it is a polymorphism of G:

1. *$t_i^G(x_1, \ldots, x_n)/R^+ = t_i^+(x_1/R^+, \ldots, x_n/R^+)$, when $x_1, \ldots, x_n \in G$ are not sinks and*

2. *$t_i^G(x_1, \ldots, x_n)/R^- = t_i^-(x_1/R^-, \ldots, x_n/R^-)$, when $x_1, \ldots, x_n \in G$ are not sources.*

Proof. Note that t^- is a polymorphism by Lemma 4.7 (5). The proof of the present lemma is given by Kazda [22, pp. 394–395] in the particular case where t^+ is known to be a majority polymorphism, however the proof makes no use of the arity, nor of the majority equations. $\qquad\square$

Recall that an equation $s(x_1, \ldots, x_n) \approx t(y_1, \ldots, y_m)$ function is said to be *balanced* if $\{x_1, \ldots, x_n\} = \{y_1, \ldots, y_m\}$. Idempotence is an example of a balanced equation: $s(x, \ldots, x) \approx x$.

Remark 4.9. *The following theorem makes use of an assumption that a system of term equalities is nontrivial (that is, that $x \approx y$ is not a consequence). In general such assumptions are not algorithmically verifiable, but it will be explained in the course of the proof why the particular form used in the assumption can always be satisfied on any set.*

Theorem 4.10. *Consider a property C of digraphs defined either by a system of balanced equations, or by the existence of idempotent polymorphisms $t_1, t_2, \ldots, t_k$ (not necessarily distinct) satisfying a single equational sequence*

$$t_1(x_{1,1}, x_{1,2}, \ldots, x_{1,n_1}) \approx \cdots \approx t_k(x_{k,1}, x_{n,2}, \ldots, x_{k,n_k}) \approx x, \qquad (*)$$

where $\{x_{1,1}, \ldots, x_{1,n_1}\} = \cdots = \{x_{k,1}, \ldots, x_{k,n_k}\}$ and $x \in \{x_{1,1}, \ldots, x_{1,n_1}\}$. The following statements are true provided that the equation $x \approx y$ does not follow from C.

1. *Let G be a totally rectangular digraph. Then G has property C if and only if G_∞ has property C.*

2. *Let G be universally rectangular. Then G has property C with each t_i conservative if and only if G_∞ has property C with each of the t_i conservative.*

Proof. Our proof is very similar to the main proof in [22]. We focus primarily on the conservative case (not considered in [22] which also only explicitly applies to majority equations), as the non-conservative case is obtained by following this proof and missing some steps (we give some pointers of the required omissions during the proof). There are two possible kinds of equations mentioned as determining the property C in the theorem statement: systems of balanced equations, and those determined by $(*)$. In the proof to follow we refer to systems of balanced equations as equations of the first kind, while those of the form in $(*)$ (a single family of k of equalities, with $\approx x$) as equations of the second kind. We mention that in either case, C holds conservatively is the polymorphisms witnessing the defining equations are conservative.

We define a notion of *pattern* for n-tuples. Let $(x_1, \ldots, x_n)$ be an n-tuple of symbols, in which possibly some repeats occur. We say that an n-tuple $(a_1, \ldots, a_n)$ *agrees* with the n-tuple $(x_1, \ldots, x_n)$ if $x_i = x_j$ implies $a_i = a_j$. We will make frequent agreement comparisons to the tuples appearing as arguments of terms in equations.

It is easy to see that if G has conservative property C then so does G/R^+ by defining $t_i(x_1/R^+, \ldots, x_n/R^+) = t_i(x_1, \ldots, x_n)/R^+$. Thus, it suffices to show that if G/R^+ satisfies C conservatively then so does G. This will be shown by backward induction over successive

quotients by R^+. This will be verified using Lemma 4.4 rather than the assumption of the Maltsev property directly, moreover instead of explicit use of universal rectangularity, the argument uses only the fact that on each successive quotient by R^+, both rectangularity and the conclusion of Lemma 4.6 hold.

Assume that G/R^+ has conservative polymorphisms $t_1^+, \ldots, t_k^+$ witnessing C. Let $t_i^-(x_1, \ldots, x_n)$ denote the operation given by $\phi t_i^+(\phi^{-1}(x_1), \ldots, \phi^{-1}(x_n))$. We show that one can construct conservative functions $t_1^G, \ldots, t_k^G$ on G, each satisfying the conditions stated in Lemma 4.8 and such that the equation sequence defining C holds. By Lemma 4.8, these are the desired family of polymorphisms witnessing C on G.

In the second kind of equation system (those defined by $(*)$), the values of the t_i are forced at many tuples: if $(a_1, \ldots, a_n)$ agrees with the pattern $(x_{i,1}, \ldots, x_{i,n_i})$ (with $x_{i,p} = x$), then $t_i^G(a_1, \ldots, a_n)$ must equal a_p. Idempotence may be considered as one such case: it forces $t_i(a, \ldots, a) = a$. Note that if some tuple is asked to take two different values under the same term symbol by this process, then this facilitates a proof of $x \approx y$ from $(*)$, which we have assumed does not hold. Trivially, the reverse implication is true, which justifies the claim in Remark 4.9, for term conditions of this second kind. Note also that conditions (1) and (2) of Lemma 4.8 hold trivially for such tuples, because if $(a_1, \ldots, a_n)$ agrees with the pattern $(x_{i,1}, \ldots, x_{i,n_i})$, then $(a_1/R^+, \ldots, a_n/R^+)$ agrees with the pattern $(x_{i,1}, \ldots, x_{i,n_i})$ also, and so on. We observe that in this second kind of equation $(*)$, the equations defining C are already guaranteed to hold at this point (and what remains is just ensuring they can be extended as polymorphisms to other tuples).

We now fix some arbitrary linear order $\leqslant_G$ on the vertices V_G and proceed to choose values for $t^G(a_1, \ldots, a_n)$ which have not so far been forced. Note that this is where the definition of the t_i^G begins in the case of equations of the first kind. Assume that one of $a_1, \ldots, a_n \in V_G$ is a source. Choose the value of $t^G(a_1, \ldots, a_n)$ to be the smallest element a_ℓ (under $\leqslant_G$) of $\{a_1, \ldots, a_n\}$ such that $a_\ell/R^+ = t^+(a_1/R^+, \ldots, a_n/R^+)$. This implies that Condition (1) of Lemma 4.8 is satisfied for t^G at the tuple $(a_1, \ldots, a_n)$, while (2) is satisfied vacuously. (In the non-conservative case, we would simply take the smallest element of the R^+-equivalence class $t^+(a_1/R^+, \ldots, a_n/R^+)$.) Dual statements hold if one of $a_1, \ldots, a_n$ is a sink.

Note that, given a_ℓ/R^+, the choice of element in the class a_ℓ/R^+ is determined only by the set $\{a_1, \ldots, a_n\}$ and the fixed order $\leqslant_G$. Thus any balanced equation will hold at such tuples. This also justifies the claim in Remark 4.9 for term conditions of the first kind, as the construction just given can be applied to any set (treated as a graph with no edges, so that all tuples are sources) carrying a linear order, and clearly any set can be given a linear order.

Now consider the case where none of the $a_1, \ldots, a_n$ are sources or sinks. We will show that the R^+-class $t^+(a_1/R^+, \ldots, a_n/R^+)$ intersects the R^--class $t^-(a_1/R^-, \ldots, a_n/R^-)$ and that at least one of $a_1, \ldots, a_n$ lies in this intersection (this is the conservative part); this will be chosen as the value of $t^G(a_1, \ldots, a_n)$.

Since t^+ and t^- are conservative functions we have $t^+(a_1/R^+, \ldots, a_n/R^+) = a_l/R^+$ for some $l \leqslant n$ and $t^-(a_1/R^-, \ldots, a_n/R^-) = a_{l'}/R^-$ for some $l' \leqslant n$. We show that $a_l/R^+ \cap a_{l'}/R^- \neq \varnothing$. Because any vertex $u \in V_G$ that is neither a source nor sink has

$u \in u/R^+ \cap u/R^-$ it suffices to consider the case where $l \neq l'$.

Now, as $a_i/R^+ \cap a_i/R^- \neq \varnothing$, for each $i = 1, \ldots, n$, we have

$$\phi(\phi^{-1}(a_i/R^-)) \cap a_i/R^+ \neq \varnothing,$$

and therefore

$$(\phi^{-1}(a_i/R^-), a_i/R^+) \in E(G/R^+).$$

Thus

$$(t^+(\phi^{-1}(a_1/R^-), \ldots, \phi^{-1}(a_n/R^-)), t^+(a_1/R^+, \ldots, a_n/R^+)) \in E(G/R^+),$$

from which it follows that

$$t^-(a_1/R^-, \ldots, a_n/R^-) \cap t^+(a_1/R^+, \ldots, a_n/R^+) = a_l/R^+ \cap a_{l'}/R^- \neq \varnothing.$$

As $a_l/R^+ \cap a_{l'}/R^- \neq \varnothing$, we may employ Lemma 4.6 to show that there is $i \leqslant n$ with $a_i \in a_l/R^+ \cap a_{l'}/R^-$. Choose $t^G(a_1, \ldots, a_n)$ to be the $\leqslant_G$-earliest amongst the $a_i \in a_l/R^+ \cap a_{l'}/R^-$. It is now trivial that properties 1 and 2 hold, which as observed already, ensures that the t_i^G are polymorphisms of G. We must show that the equations $(*)$ hold for this case (or that balanced equations hold). The case labelled $(*)$ (where $\approx x$ is assumed), has been established earlier in the proof, so we are considering here only systems of balanced equations. Consider an individual instance of a balanced equation $t_i(x_{i,1}, \ldots, x_{i,n_i}) \approx t_j(x_{j,1}, \ldots, x_{j,n_j})$ in the list, where $\{x_{i,1}, \ldots, x_{i,n_i}\} = \{x_{j,1}, \ldots, x_{j,n_j}\}$. Now, the equations defining C hold at the level of t_i^+, t_j^+ and t_i^-, t_j^-, so that the process just performed for selecting the value of t_i and t_j at a tuple without sources or sinks will always at least be made within the intersection of the same block of R^+ and block of R^-. However the particular choice of element *within* this set was made only in terms of the order $\leqslant$ and the set $\{a_1, \ldots, a_n\}$, which is uniform for both t_i and t_j. Thus equality at such tuples is verified. $\qquad\square$

Some instances of polymorphisms satisfying the conditions of Theorem 4.10 are majority, Maltsev and Pixley. In these cases G_∞ always has the desired polymorphism, giving the following corollary.

Corollary 4.11. *Let G be a digraph.*

1. *G admits a (conservative) Maltsev polymorphism iff it admits a (conservative) Pixley operation iff it is totally (universally) rectangular.*

2. *If G is totally (universally) rectangular then G admits a (conservative) minority polymorphism and a (conservative) majority polymorphism.*

Remark 4.12. *The first part of Corollary 4.11 strengthens the result given in Lemma 4 of [13], for the case of digraphs. The relational clone $\langle \mathbf{B} \rangle$ of a structure $\mathbf{B}$ is the set of all relations that can be expressed with primitive positive first-order formulas (that is, only existential quantification, conjunction, and equality is allowed) from $\mathbf{B}$. When*

we restrict [13, Lemma 4] to digraphs, it can be stated as follows: A digraph G is pre-served by a Maltsev operation iff every binary relation in $\langle G \rangle$ is rectangular. It is easy to see and well known that every binary relation in $\langle G \rangle$ can be expressed as $B_G(S, a, b) = \{(h(a), h(b)) | h$ is a homomorphism from S to $G\}$ for a structure S with two distinguished vertices a and b. Then Corollary 4.11 implies that for a digraph G to be preserved by a Maltsev operation it is enough to require that only those binary relations in $\langle G \rangle$ that can be expressed as $B_G(S, a, b)$, where S is a directed path with initial vertex a and terminal vertex b, are rectangular.

The above corollary yields an algorithm for verifying if a graph has a Maltsev (or Pix-ley) polymorphism. Indeed, the rectangularity of a digraph is equivalent to the following property of its adjacency matrix: when two rows (or two columns) share a common 1 they are identical. On an n-vertex digraph this property may be verified in $O(n^3)$ steps. A digraph has a Maltsev polymorphism if and only if each (of at most n) successive quo-tient by R^+ is rectangular, with the process stopping once there are no R^+-classes of size more than 1 (which happens after at most n quotients). Overall this takes $O(n^4)$ steps (quadratic in terms of the size of the adjacency matrix).

Universal rectangularity (equivalently, the existence of a conservative Maltsev poly-morphism) can also be verified in polynomial time by verifying total rectangularity and conservative 2-rectangularity at each successive quotient by R^+. In fact, the proof of Theorem 4.10 is sufficiently constructive to construct the desired polymorphisms (when they exist): simply work backwards from their definition of G_∞.

4.2 Conservative 2-Semilattice Polymorphisms

A disjoint union of directed cycles admits a conservative commutative binary (ccb) poly-morphism (which coincide with conservative 2-semilattice operations) if and only if it contains no even cycles. This provides a case where Theorem 4.10 characterises a proper subclass of conservative Maltsev digraphs. In this section we classify Maltsev digraphs admitting a ccb polymorphism. A corollary of the result will be a sort of converse to Theorem 4.10: a Maltsev digraph with a ccb polymorphism is necessarily conservative Maltsev (Proposition 4.14 below).

Consider any digraph G and let $*$ be any conservative commutative binary operation on V_G. The operation $*$ has an easy interpretation as a colouring of the nondiagonal elements of the cartesian square $V_G^2 \setminus \{(v, v) \mid v \in V_G\}$: the pair (a, b) is coloured L if $a * b = a$ and R if $a * b = b$; commutativity is equivalent to (a, b) having different colour to (b, a). (The exclusion of the diagonal elements is only for convenience.) We now examine the consequences of $*$ being a ccb polymorphism.

For any digraph G we define a structure—the *ccb graph* of G—on the nondiagonal elements of the cartesian square $V_G^2 \setminus \{(v, v) \mid v \in V_G\}$. The ccb graph is a graph with two kinds of edges: "orienting" edges, which are directed, and "straight" edges which are considered as having no direction.

- A (directed) orienting edge is placed from (a_1, b_1) to (a_2, b_2) if there are parallel edges connecting each of the following pairs in G: a_1 and a_2; b_1 and b_2; and b_1 and

a_2 *but not* a_1 *and* b_2. The following diagram depicts two situations that an orienting edge pointing from (a_1, b_1) to (a_2, b_2) can arise:

- An (undirected) straight edge is placed from (a_1, b_1) to (a_2, b_2) if there are parallel edges connecting the following pairs: a_1 and a_2; b_1 and b_2; but not a_1 and b_2 or a_2 and b_1.

A directed path in the ccb graph is a path of orienting edges and straight edges, in which each orienting edge is traversed in a forward direction. We now adopt the notation $\leadsto$ to denote directed paths: note that $(a, b) \leadsto (c, d)$ if and only if $(d, c) \leadsto (b, a)$. An important observation is: if (a, b) is not connected to (c, d) in the ccb graph of G then the compatibility of the ccb operation with the edges of G does not fail at the pairs (a, b) and (c, d). Thus *the digraph G admits a ccb polymorphism iff the ccb graph of G can be coloured by L and R such that L and R are preserved across straight edges, L is preserved forward across orienting edges and R is preserved backward across orienting edges.* Expressed in terms of $\leadsto$ this becomes the following rules.

(L) If (u, v) is coloured L and $(u, v) \leadsto (x, y)$ then:

    (1) (x, y) is coloured L.

    (2) (y, x) is coloured R.

(R) If (u, v) is coloured R and $(x, y) \leadsto (u, v)$ then:

    (1) (x, y) is coloured R.

    (2) (y, x) is coloured L.

In order to show that the existence of a colouring of a ccb-graph can be decided in polynomial time, we next consider a kind of a Möbius band. As the following example illustrates, we may think of a directed path connecting (a, b) to (b, a) in the ccb-graph of G as a kind of Möbius band of parallel edges in G.

(Möbius band in G)

(directed path in the ccb-graph of G)
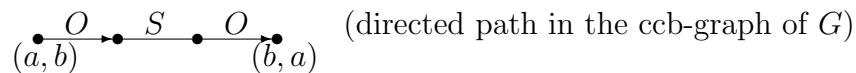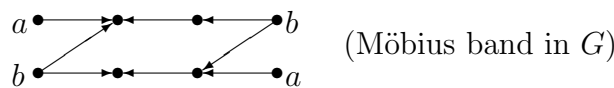
Theorem 4.13. *A digraph G admits a ccb polymorphism $*$ if and only if for every distinct $a, b \in G$, the ccb-graph of G does not contain a directed path both from (a, b) to (b, a) and from (b, a) to (a, b). When a ccb polymorphism exists, it can be constructed in a polynomial number of steps.*

Proof. The forward implication is simply the statement that the colouring of the ccb-graph must colour each pair (a, b) oppositely to its reverse (b, a), and the colouring rules (L)(1)–(R)(2) must be obeyed.

Now we show the converse: assume that for every $a, b \in G$, the ccb-graph of G does not contain a directed path from (a, b) to (b, a) and from (b, a) to (a, b). We construct (in a polynomial number of steps) a successful colouring by L and R, whence a ccb polymorphism.

First phase. Begin by finding any pairs (a, b) from which there is a directed path to the reverse pair (b, a) (necessarily it must contain an orienting edge, as otherwise there is a path from (b, a) to (a, b)). In every case, colour (a, b) by R and (b, a) by L. By assumption, no pair is coloured two different colours simultaneously. It might appear that the colouring rules might force the colouring of some further elements of G^2, but in fact this is not the case as all potential applications of rules (L)(1)–(R)(2) are between elements already consistently coloured by this first phase. For example, say that $(c, d) \rightsquigarrow (a, b)$, with (a, b) coloured R because of $(a, b) \rightsquigarrow (b, a)$. Then, while (R1) would force (c, d) to be coloured by R, we already have $(b, a) \rightsquigarrow (c, d)$, giving $(c, d) \rightsquigarrow (a, b) \rightsquigarrow (b, a) \rightsquigarrow (d, c)$. Hence (c, d) was already coloured by R. (The other cases are equivalent, up to reversing.)

Second phase. Now there may still be some uncoloured pairs. We are going to complete the colouring according to the following process which is to be repeated until all pairs are coloured (we show the process is well defined below): take any uncoloured pair, (a, b) say, and colour it arbitrarily; colour any other pairs for which rules (L)(1)–(R)(2) apply starting from (a, b).

We need to ensure that the second step is well defined, that is, that no pair (c, d) is forced to have been coloured two different colours simultaneously. Up to reversing everything, we can assume without loss of generality that at some stage a hitherto uncoloured pair (a, b) is given the colour L, and that $(a, b) \rightsquigarrow (c, d)$. There are two potential problems: (c, d) has already been coloured by R (that is, clashing with a previously assigned colour); or that the choice of colouring (a, b) by L itself also forces (c, d) to be coloured R (that is, $(a, b) \rightsquigarrow (d, c)$). In either case, we get a contradiction with the assumption that (a, b) was uncoloured. In the first case, rule (R)(1) already forces (a, b) to be coloured R (either at some earlier iteration of the second phase, or if (c, d) was coloured at the first phase, then so was (a, b) coloured at the first phase). In the second case, $(a, b) \rightsquigarrow (d, c)$ implies $(c, d) \rightsquigarrow (b, a)$, giving $(a, b) \rightsquigarrow (c, d) \rightsquigarrow (b, a)$, showing that (a, b) was coloured in the first phase.

Thus, eventually, the desired colouring (and ccb-polymorphism) is eventually achieved. □

Note that the property in Theorem 4.13 can be verified using Reachability in the ccb-graph of G, so is solvable in nondeterministic logarithmic space. Hence deciding if a digraph has a ccb polymorphism is in NL too.

Proposition 4.14. *The following are equivalent for a Maltsev graph G:*

1. G has a conservative commutative idempotent binary polymorphism;

2. G has a conservative commutative idempotent binary polymorphism and a conservative Maltsev polymorphism;

3. G has a conservative Maltsev polymorphism and G_∞ is empty or has no even length cycles.

Proof. (2)⇒(3)⇒(1) follow from Theorem 4.10, because a disjoint union of directed cycles has a ccb iff it has no even length cycles. For (1)⇒(2) we use Theorem 4.13 to prove a version of Lemma 4.6 with conservative 2-rectangularity replaced by ccb. Indeed we have:

Claim 4.15. *Let a, b be vertices in a totally rectangular digraph G admitting a ccb polymorphism and assume that neither a nor b is a source or sink. If there is $c \in a/R^+ \cap b/R^-$ then either $b \in a/R^+$ or $a \in b/R^-$.*

Proof. The assumptions on a and b lead to the existence of not-necessarily distinct points a, b, c, g, e, f, g with at least the edges given in Figure 2. Total rectangularity ensures the existence of d. If $c = d$ we are done. Otherwise, as there can be no path of straight edges connecting (c, d) to (d, c) in the ccb graph of G, there must be an edge between one of the following: $\{a, c\}$ to h; $\{b, d\}$ to e; g to $\{c, b\}$; or h to $\{a, d\}$. But then rectangularity forces either all edges from $\{a, b, c, d\}$ and $\{e, h\}$ or all edges from $\{g, f\}$ to $\{a, b, c, d\}$. □

It then follows, like in the proof of Theorem 4.10.

□

Note that the digraph in Fig. 2 is easily seen to admit a conservative majority polymorphism, a Maltsev polymorphism, but no conservative Maltsev polymorphism. So conservative majority cannot replace ccb in Proposition 4.14.

5 Maltsev DAGs

In this section all graphs will be acyclic, that is, DAGs. Recall that all DAGs preserved by a Maltsev polymorphism are layered and retract onto a simple dipath, see Theorem 3.3.

5.1 A simple inductive construction of Maltsev DAGs

In this section we provide a simple inductive characterisation of totally rectangular DAG. We note that in [22, Corollary 16] Kazda gives an inductive construction of Maltsev digraphs, however, this construction is not fully satisfying in the sense that it is non-deterministic, that is, it does not specify how to obtain the desired preimages, and it is not clear if it can be made deterministic. The construction described below, that only works for DAGs, consists of repeated applications of two straightforward steps (and their reverse versions) which clearly specify how to obtain a new Maltsev DAG from an already constructed one by a certain copying process. We need the following definitions.

Definition 5.1. [(Reverse) arborescence] An *(reverse) arborescence* is a directed tree with root r such that every edge points away from (towards) r. An (reverse) arborescence is balanced if every (reverse) dipath from the root to a leaf has the same length.

Definition 5.2. [$\nabla(r, h)$ and $\Delta(r, h)$] Let G be a digraph, $r \in V_G$, and $h \in \mathbb{N}$. $\nabla(r, h)$ ($\Delta(r, h)$) is defined to be the subgraph of G whose vertices and edges are the vertices and edges of all (reverse) sub-dipaths of G which have initial vertex r and length h. Let v be a vertex of $\nabla(r, h)$ ($\Delta(r, h)$) such that $v \neq r$. Then v is called an *endpoint* of $\nabla(r, h)$ ($\Delta(r, h)$) if there is a (reverse) dipath of length h from r to v. Otherwise v is called an *inner vertex* of $\nabla(r, h)$ ($\Delta(r, h)$).

Definition 5.3. [isolated $\nabla(r, h)$] Let G be a digraph. Consider $\nabla(r, h)$ ($\Delta(r, h)$) for some $r \in V_G$ and $h \in \mathbb{N}$. We say that $\nabla(r, h)$ ($\Delta(r, h)$) is *isolated in G* if for every inner vertex v of $\nabla(r, h)$, both the in-neighborhood and the out-neighborhood of v belong to $\nabla(r, h)$ ($\Delta(r, h)$).

We are ready to define the construction formally in Fig. 6. This construction can be used, for example, to define a minority operation for a totally rectangular DAG.

---

1. $\mathcal{C}$ contains dipaths of all possible lengths $n \in \mathbb{N}_0$;

2. $\mathcal{C}$ is closed under applying the following operations:

   (a) Given a digraph G, let $r \in V_G$ and $h \in \mathbb{N}$ such that $\nabla(r, h)$ is an arborescence. Let ∇' be a copy of $\nabla(r, h)$. Join ∇' to G by identifying the corresponding endpoints of ∇' and $\nabla(r, h)$. Let the resulting graph be G';

   (b) Given a digraph G, let $r \in V_G$ be such that r has exactly one incoming edge (p, r), and $h \in \mathbb{N}$ such that $\nabla(r, h)$ is an isolated arborescence. Let ∇' be a copy of $\nabla(r, h)$ with root r'. Join ∇' to C by identifying the corresponding endpoints of ∇' and $\nabla(r, h)$, and adding the edge (p, r'). Let the resulting graph be G';

   (c) The reverse version of Step 2a (defined in the natural way);

   (d) The reverse version of Step 2b (defined in the natural way).

3. $\mathcal{M}$ is the set of digraphs that can be obtained by taking disjoint unions of digraphs in $\mathcal{C}$.

Figure 6: Inductive construction of the set $\mathcal{M}$ of totally rectangular DAGs.

Example 5.4. *Consider the totally rectangular DAG G'' in Fig. 7. To construct it using the method in Fig. 6, we start with the dipath G and first apply Step 2a to G to obtain G'. Next we apply Step 2d to G' obtain G''. The thick edges indicate the subgraphs $\nabla(r, 2)$ and $\Delta(r, 2)$, which are the subgraphs to be copied and attached appropriately.*

Reversing the argument of Lemma 5.10 reveals that if G is a totally rectangular DAG, then each component of G can be built using the construction in Figure 6 in a very structured manner. We make this restricted construction explicit in Figure 8.
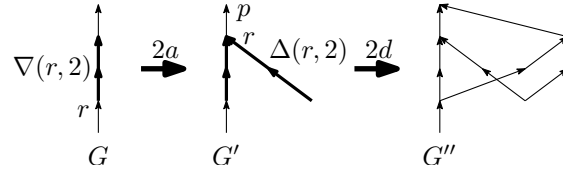
Figure 7: Construction of a totally rectangular DAG.

Construct every component C of G as follows:

1. Initially, C is a directed path of arbitrary length;

    (a) Using Step 2c of the construction in Figure 6, build an arborescence A.

    (b) Choose a vertex layer L_m such that each connected component of the subgraph between L_0 and L_m is a balanced arborescence (notice that these arborescences must have their endpoints in L_m).

        i. Choose an arbitrary $q \geqslant 0$, and repeat the following q times. Choose a vertex $v \in L_k$, where $0 \leqslant k < m$. Apply Step 2a of the construction in Figure 6 with $\nabla(v, m - k)$ to obtain a new graph.

        ii. Choose an arbitrary $q \geqslant 0$, and repeat the following q times. Choose a vertex $v \in L_k$, where $0 < k < m$. Apply Step 2b of the construction in Figure 6 with $\nabla(v, m - k)$.

        iii. Choose an arbitrary $q \geqslant 0$, and repeat the following q times. Choose a vertex $v \in L_k$, where $0 < k < m$. Apply Step 2c of the construction in Figure 6 with $\Delta(v, \ell)$, where $1 \leqslant \ell \leqslant k$.

    (c) Either go back to Step 1b, choose L_m with a larger m than previously and repeat, or finish the construction.

Figure 8: Strict construction of totally rectangular DAGs.

We observe the following properties of this restricted construction. In Steps 1(b)i and 1(b)ii, $\nabla(v, m - k)$ is always isolated. In Step 1(b)iii $F = \Delta(v, \cdot)$ is always a reverse directed path, and every vertex of F has precisely one in-neighbour.

We can now describe the class of totally rectangular DAGs.

Theorem 5.5. *The class of totally rectangular DAGs is the set of digraphs $\mathcal{M}$ defined in Fig. 6.*

Proof. Let $G = (V_G, E_G)$ be a layered digraph with levels $L_0, L_1, \ldots, L_n$. We define a natural order on the levels of G as follows: $L_0 < L_1 < \cdots < L_n$. Given $u \in V_G$ we denote by $level(u)$ the level of G that the vertex u belongs to. The result then follows from Lemma 5.6 and Lemma 5.10 below. $\qquad \square$

Before stating Lemma 5.10 we need some auxiliary results.

Lemma 5.6. *Let G be a digraph built using the construction in Fig. 6. Then G is totally rectangular.*

Proof. We show the result assuming that G is connected, and the result follows. We show that applying a construction step to a totally rectangular digraph G yields a totally rectangular digraph G'. Since a directed path is totally rectangular, this will prove the lemma.

Consider the case when G' was obtained from G by an application of Step 2a attaching ∇' to G. Take any 4 distinct vertices a, b, c, d such that there are directed paths P_{ab} from a to b, P_{cb} from c to b, and P_{cd} from c to d, each of some length ℓ. We show that there must be a directed path of length ℓ from d to a. Observe first that the construction has the property that there is no directed path from a vertex of G to a vertex of ∇', and vice versa, except for the endpoints of ∇' (which are identified with the corresponding vertices of G). Let I be the common vertices of G and ∇'. We consider four cases:

1. *Vertices a and c are in G.* Then P_{ab}, P_{cd}, and P_{cb} are subpaths of G, so the total rectangularity of G provides a directed path of length ℓ from a to d.

2. *Vertex $a \in G$ and $c \in \nabla'$.* See Figure 9 for an illustration. Note that this is possible only if each path P_{ab}, P_{cd}, and P_{cb} has a vertex u, v and w, respectively, in I. Let P_{cv} denote the subpath of P_{cb} from c to v, and P_{vb} the subpath of P_{cb} from v to b. Similarly define the subpaths P_{cw} and P_{wd} of P_{cd}. Observe that by construction, P_{cv} is a copy of a dipath $P_{c'v}$ from a vertex c' to v in G. Similarly, P_{cw} is a copy of a dipath $P_{c'w}$ from a vertex c' to w in G. We have dipaths P_{ab} from a to b, $P_{c'v}P_{vb}$ from c' to b, and $P_{c'w}P_{wd}$ from c' to d, each of length ℓ and fully inside G. Therefore we have a dipath of length ℓ from a to d.

3. *Vertex $c \in G$ and $a \in \nabla'$.*
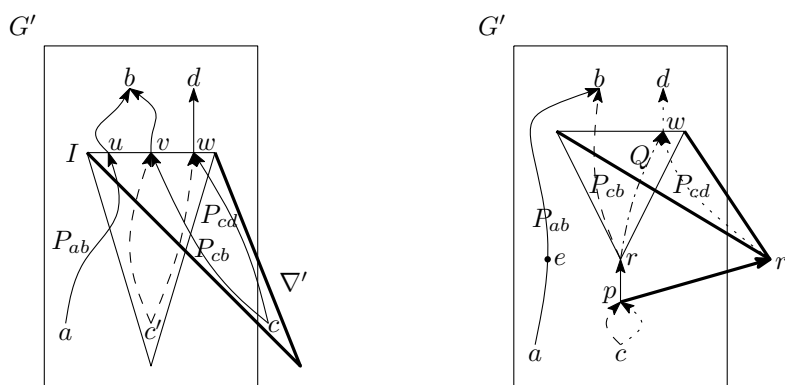
4. *Both vertices a and c are in ∇'.*

Figure 9: Left: Vertex $a \in G$ and $c \in \nabla'$ in analysis of Step 2a. Right: The case when P_{cb} and P_{cd} meet at p in the analysis of Step 2b.

For the last two cases, an analysis fairly similar to the analysis of the second case can be done.

Suppose now that G' was obtained from G by an application of Step 2b of the construction. Let r, r', h, p, $\nabla(r, h)$, and ∇' be as in the definition of Step 2b. Take any 4 distinct vertices a, b, c, d such that there are directed paths P_{ab} from a to b, P_{cb} from c to b, and P_{cd} from c to d, each of the same length ℓ. We show that $d \xrightarrow{\ell} a$. If $level(a) = level(c) \geqslant level(r)$, then obviously, the same analysis works as for Step 2a. In addition, $level(b) = level(d) > level(p)$ because otherwise all paths are fully inside G and the total rectangularity of G proves the claim.

Furthermore, there must be two (distinct) edges e_1 and e_2 among the edges of P_{ab}, P_{cb} P_{cd}, such that e_1 is an edge in $A = \{\nabla(r, h) \cup (p, r)\}$ and $e_2 \in B = \{\nabla' \cup (p, r')\}$. To see this, remove A from G' to obtain the graph G, and observe that F is isomorphic to G. Assume then that there is no edge of P_{ab}, P_{cb} P_{cd} in A. The total rectangularity of F yields the desired path from a to d. Because $\nabla(r, h)$ and ∇' are isolated and p is the only incoming edge of r and r', the only way that at least two of P_{ab}, P_{cb} P_{cd} can be in A and B is if at least two of P_{ab}, P_{cb} P_{cd} enters A and B through p.

If P_{ab} and P_{cd} both meet p, then we obviously have a dipath from a to d. If P_{ab} and P_{cb} meet at p then let P_{ap} be the subpath of P_{ab} from a to b, and P_{cp} the subpath of P_{cb} from c to p. Assume $k = len(P_{ap})(= len(P_{cp}))$, and let Q be the subpath of P_{cd} starting at c and taking the first k edges. Suppose the endpoint of Q is e. Then by the total rectangularity of G, there is a path of length k from a to e. This path from a to e and the subpath of P_{cd} from e to d provides the desired path from a to d.

If P_{cb} and P_{cd} meet at p, then using the fact that G and F are isomorphic, we can assume w.l.o.g. that P_{cb} passes through r and P_{cd} passes through r'. Let $P_{r'w}$ be the subpath of P_{cd} from r' to an endpoint w. Because $P_{r'w}$ is entirely in B, $P_{r'w}$ is a copy of some dipath Q from r to w entirely in A and therefore entirely in G. Let P_{wd} the subpath of P_{cd} from w to d. Let P_{eb} the subpath of P_{ab} from e to b, where e is the vertex of P_{ab} such that $level(e) = level(r)$. Let P_{rb} be the subpath of P_{cb} from r to b. Then using the dipaths P_{eb}, P_{rb}, QP_{wd}, we obtain a dipath R from e to d using the total rectangularity of G. Then following P_{ab} up to e and then using R we obtain the desired dipath from a to d. $\qquad\square$

Lemma 5.7. *Assume that G is a totally rectangular DAG, r is a vertex of G, and $h \in \mathbb{N}$ such that $A = \nabla(r, h)$ is an arborescence in G. Suppose that G contains a dipath $P = v_0, \ldots, v_h$ of length h such that v_h is an endpoint of A, but P and A do not share any other vertices. Then $A' = \nabla(v_0, h)$ is isomorphic to A through an isomorphism φ, such that φ is the identity map on the endpoints of A'. Furthermore, A' and A share only the endpoints.*

Proof. The proof is by induction on the height h of A. When $h = 1$ the result trivially follows from the assumption that G is totally rectangular. Assume that $h > 1$. Remove the root r of A to obtain sub-arborescences $A_1, \ldots, A_s$, each of height $h - 1$, and with roots $r_1, \ldots, r_s$, respectively. Fix an arbitrary endpoint e_i for each A_i. Once the following

Claim is established, it is easy to see that the inductive hypothesis, applied on paths $Q_1, \ldots, Q_n$, yields the result.

Claim: There are dipaths $Q_1, \ldots, Q_s$ of length h from v_0 to e_i, for each $i \in [s]$, such that the only vertex any pair of these paths share is v_0. Furthermore for any $i \in [s]$, Q_i does not contain any inner vertices of A.

If $s = 1$, that is, r has only one out-neighbour, then the claim obviously follows. So assume that $s \geqslant 2$, and fix a dipath Q_i from v_0 to e_i for each $i \in [s]$ as follows. Let Q be a path from r to e_i. By assumption v_h is an endpoint of A, so there is a path Q' from r to v_h. By the total rectangularity of G, there must be a path of length h from v_0 to e_i. Let this path be Q_i.

Assume, for a contradiction, that there is pair of paths Q_i and Q_j that share the vertex x, and $x \neq v_0$. Suppose that x is as far as possible from v_0 (note that if the distance from v_0 to x is different on Q_i and Q_j, then G cannot be layered). Let $Q_{i,x}$ and $Q_{j,x}$ be the subpath of Q_i and Q_j from x to e_i and e_j, respectively. Suppose that $len(Q_{i,x}) = len(Q_{j,x}) = \ell$. Let R be the reverse directed path of length ℓ in S_i starting at e_i and ending at some vertex t. Then $Q_{i,x}$, $Q_{j,x}$ and R together with the total rectangularity of G guarantee the existence of a path of length ℓ from t to e_j. This provides two distinct dipaths of length h from r to e_j, contradicting the fact that A is an arborescence.

Suppose, for a contradiction, that for some $i \in [s]$, Q_i contains an inner vertex f of A. Let $Q_{i,f}$ be the subpath of Q_i from v_0 to f, and set $k = len(Q_{i,f})$. Let P_{rf} be a dipath from r to f of length k. Then by the rectangularity of G, the existence of the paths $Q_{i,f}$, P_{rf} and $v_0, \ldots, v_k$ guarantees the existence of a path of length k from r to v_k, and therefore v_k is an inner vertex of A, a contradiction. $\qquad\square$

The following proposition is an easy observation.

Proposition 5.8. *Let G be a digraph, r a vertex of G, and $h \in \mathbb{N}$. Then $\nabla(r, h)$ is an arborescence if, and only if for every pair of distinct vertices v, u in $\nabla(r, h)$, there is at most one dipath from v to u.*

Lemma 5.9. *Let G be a totally rectangular DAG. Let r be a vertex of G and $h \in \mathbb{N}$. Assume that $\nabla(r, h)$ is isolated in G. Then $\nabla(r, h)$ is either an arborescence, or it can be obtained from an arborescence by repeated applications of Step 2b of the construction.*

Proof. If $\nabla(r, h)$ is an arborescence there is nothing to do. If $\nabla(r, h)$ is not an arborescence, then by Proposition 5.8, there is at least one pair of vertices u, v such that there are at least two distinct dipaths from u to v. It is sufficient to show that in this scenario, we can produce a totally rectangular subgraph G_s of G in which there are fewer directed paths from u to v, and an application of Step 2b of the construction yields G from G_s.

Take the largest $g = 1, \ldots, h$ such that $\nabla(r, g)$ is still an arborescence (note that $\nabla(r, 1)$ is always an arborescence). Then we can find an endpoint v of $\nabla(r, g + 1)$ and a vertex u in $\nabla(r, g)$ (a subgraph of $\nabla(r, g + 1)$) such that there are two disjoint dipaths from u to v. Choose such a pair of vertices u and v such that these two distinct dipaths $P = u, w_1, \ldots, w_m, v$ and $Q = u, w'_1, \ldots, w'_m, v$ are shortest possible.

By the choice of u and v, there is no vertex pair u', v' in $\nabla(w_1, m)$ such that there is more than one dipath from u' to v'. Therefore $\nabla(w_1, m)$ is an arborescence. To see that $\nabla(w_1, m)$ is in fact an isolated arborescence, assume for contradiction that there is an edge (s, t) such that t is an inner vertex of $\nabla(w_1, m)$ and $s \notin \nabla(w_1, m)$. Then t is also an inner vertex of $\nabla(r, h)$, and because $\nabla(r, h)$ is isolated, s must be in $\nabla(r, h)$. This implies that $\nabla(r, h)$ contains some vertex s' such that there are at least two directed paths from s' to t. By Proposition 5.8, this would contradict the choice of g because $\nabla(r, g)$ is an arborescence. Similarly, $\nabla(w_1', m)$ is an isolated arborescence in G. By a similar argument, w_1 or w_2 has only the in-neighbour u.

Using the subpath $w_1', \ldots, w_m', v$ of Q and Lemma 5.7, we obtain that $\nabla(w_1', m)$ is isomorphic to $\nabla(w_1, m)$. Let G_s be the graph obtained from G be removing $\nabla(w_1', m)$ (but not the endpoints) and the edge (u, w_1') (but not u). It is clear now that we can use Step 2b of the construction to obtain G from G_s.

To complete the proof, we argue that G_s is totally rectangular. Assume that there are directed paths of length k from a to c, a to d and b to d in G_s (and therefore in G), but not from a to d in G_s. Let the path from a to d in G be P. Clearly, some part of P must be removed in G_s. If there is a directed path from a to u or $a = u$, and there is a directed path from an endpoint of $\nabla(w_1', m)$ to d or if d is such an endpoint, then going through $\nabla(w_1, m)$ instead of $\nabla(w_1', m)$ gives a path of length k from a to d. Suppose therefore that there is no path from a to u and $a \neq u$ (the other case is similar). Since $\nabla(w_1', m)$ is isolated, and the only in-neighbor of w_1' is u, if a is not a vertex of $\nabla(w_1', m)$ (except for the endpoints), then P does not go through any removed part of the graph, and therefore P is unaffected by the removal. Thus, a is must be a vertex of $\nabla(w_1', m)$ that is not an endpoint, and therefore a is not in G_s, a contradiction. $\qquad\square$

Lemma 5.10. *Let G be a totally rectangular DAG. Then each component of G can be built using the construction in Fig. 6.*

Proof. Without loss of generality, we assume that G has only one component. By Theorem 3.3, G retracts onto a directed path of some length n. We show that if G is not a dipath, then we can always define a totally rectangular DAG (which has one component) such that G' is smaller than G and an application of a step of the construction gives G from G'. Let L_0 be the bottommost vertex level of G (i.e., the vertex level that contains only sources). $(\star)$ Let L_i be the first level above L_0 such that it contains a vertex w which has at least two in-neighbours. Consider the induced subgraph of G between L_0 and L_i and call it H.

Case 1: H contains a sink that is not in L_i. Let t be such a sink with the smallest possible level. By the choice of w, t has precisely one in-neighbour. Find a shortest reverse directed path P from t to a vertex c such that c has at least two out-neighbours. To see that such a vertex exists, observe that there is an oriented path Q from t to w because G is connected. Since t is a sink, the first edge of Q is a backward edge. As we follow the vertices of Q, eventually we must meet a forward edge because $level(t) < level(w)$.

Set $\ell = len(P)$. Since c has at least two out-neighbours, there is a dipath R starting at c and distinct from P. We can choose R with length ℓ because t is a sink closest to

L_0. Every vertex of R must have in-degree one by the choice of w. Furthermore, P and R shares only the vertex c because otherwise the choice of w would be violated. Let the endpoint of R be t'. Remove P from G to obtain G'. Then clearly, G' is connected, totally rectangular, and we can obtain G back from G' by applying Step 2c with $\Delta(t', \ell)$. Repeating this argument if necessary, we can assume that H contains no sinks except in L_i.

Case 2: H contains no sinks except in L_i. Suppose that there is a source s in H such that $level(s) = L_{i-k}$ and $\nabla(s, k)$ is not an arborescence. Observe that by the choice of L_i, $\nabla(s, k)$ is isolated, and therefore using Lemma 5.9, we can obtain a smaller graph G' such that we can get back G from G' by applying Step 2b. Repeating this argument if necessary, we can assume that for any source s of H, $\nabla(s, \cdot)$ is an arborescence. (Also note that we can still assume that H contains no sinks except in L_i.)

Therefore if there is a connected component C of H that is not an (balanced) arborescence, then C must have more than one sources. Let s and s' be such sources, and assume WLOG that $level(s) \leqslant level(s')$. Since C is connected, there is a directed path P of some length ℓ from s' to an endpoint e of the arborescence A_s rooted at s. Let r be the unique vertex of A_s such that there is a directed path of length ℓ from r to e. Then if G' is the graph obtained from G by removing $\nabla(s', \ell)$ (except for the endpoints), then we can get back G from G' by applying Step 2a of the construction with $\nabla(r, \ell)$ of G'. Repeating this argument if necessary, we can assume that any connected component of H is an (balanced) arborescence such that all endpoints are in L_i.

If G contains a vertex that has indegree at least 2, then we jump back to $(\star)$ and repeat the whole argument. Notice that when we choose L_i, i will increase. If all vertices of G have indegree 1, then G is an arborescence, and it can be easily constructed starting with a directed path and repeatedly applying Step 2c. $\qquad\square$

5.2 Connections with series parallel digraphs

The class of series parallel digraphs, or series parallel orders, and several modifications of these, such as generalized series parallel (GSP) digraphs, have been widely studied since the 70s. They not only can be used to model series and parallel circuits but also can be applied to scheduling problems, and are interesting from a computational point of view, since several problems that are NP-hard in general, can be solved efficiently when restricted to this class of digraphs. We refer the reader to [24] for a survey on some computational and structural properties of these digraphs. We will focus on generalisations of series parallel digraphs, and relate DAGs admitting Maltsev and conservative Maltsev polymorphisms with these classes of graphs. We start with some definitions.

We say that a digraph G is *N-free* if for any four distinct vertices of G, x, y, u, v, we have

$$x \to u \ \& \ y \to u \ \& \ y \to v \Rightarrow x \to v.$$

A digraph G^t is said to be a *transitive reduction* of the digraph G if

(i) G^t has a directed path from vertex u to vertex v if and only if G has a directed path from vertex u to vertex v, and

(ii) there is no graph with fewer arcs than G^t satisfying condition (i).

The class of *Minimal Series Parallel (MSP)* is the class of DAGs, defined recursively as follows:

- the DAG having a single vertex and no edges is MSP;

- if $\mathbf{H}_1 = (V_1, E_1)$ and $\mathbf{H}_2 = (V_2, E_2)$ are MSP, then so are the DAGs obtained from $\mathbf{H}_1$ and $\mathbf{H}_2$, by applying the following operations:

  - *parallel composition* $\mathbf{H} = (V_1 \cup V_2, E_1 \cup E_2)$, and
  - *series composition* $\mathbf{H} = (V_1 \cup V_2, E_1 \cup E_2 \cup (N_1 \times R_2))$, where N_1 is the set of sinks of $\mathbf{H}_1$ and R_2 is the set of sources of $\mathbf{H}_2$.

A DAG is *Generalized Series Parallel (GSP)* if its transitive reduction is MSP.

The class of *Quasi-Series-Parallel (QSP)* is the class of DAGs obtained recursively as follows:

- the DAG having a single vertex and no edges is QSP;

- if $\mathbf{H}_1 = (V_1, E_1)$ and $\mathbf{H}_2 = (V_2, E_2)$ are QSP, then so are the DAGs obtained from $\mathbf{H}_1$ and $\mathbf{H}_2$, by applying the following operations:

  - parallel composition, and
  - *quasi-series composition* $\mathbf{H} = (V_1 \cup V_2, E_1 \cup E_1 \cup (M \times N))$, with M a subset of the set of sinks of $\mathbf{H}_1$ and N a subset of the set of sources of $\mathbf{H}_2$.

Quasi-Series-Parallel graphs have been characterized as follows:

Theorem 5.11. [25] Given a DAG $\mathbf{H} = (V, E)$, the following are equivalent:

1. $\mathbf{H}$ is QSP;

2. for all $u, v \in V$ either $u^{-1} = v^{-1}$ or $u^{-1} \cap v^{-1} = \emptyset$;

3. the transitive closure of $\mathbf{H}$ contains no induced N.

It is not hard to see that any DAG admitting a Maltsev polymorphism is QSP, but the reverse is not true, indeed the DAG in Figure 1 is QSP but does not admit a Maltsev polymorphism.

A DAG $\mathbf{H}$ is *complete bipartite composite (CBC)* if there exists a set of complete bipartite subgraphs of $\mathbf{H}$: $B_1, B_2, \ldots, B_k$ called *bipartite components* of $\mathbf{H}$, such that: each edge of $\mathbf{H}$ belongs to exactly one bipartite component; for every non-sink vertex v, all edges leaving v belong to the same bipartite component; for each non-source vertex v, all edges entering v belong to the same bipartite component. CBC digraphs admit the following characterization.

Theorem 5.12. [24] Given a DAG $\mathbf{H} = (V, E)$, the following are equivalent:

1. $\mathbf{H}$ is CBC;

2. for all $u, v \in V$ either $u^{\pm 1} = v^{\pm 1}$ or $u^{\pm 1} \cap v^{\pm 1} = \emptyset$;

3. $\mathbf{H}$ is the transitive reduction of an N-free partial order;

4. $\mathbf{H}$ is rectangular.

Clearly all Maltsev DAGs are CBC but the reverse is not true, the DAG in Fig. 1 is an example of a CBC that is not Maltsev. In fact a characterization in terms of series parallel constructions for Maltsev DAGs eluded us, even though the characterization of Maltsev DAGs in terms of neighbourhoods is a natural generalization of the characterizations given above. Also, we can think of a characterization of Maltsev DAGs in terms of CBC, in the sense that a DAG is Maltsev if and only if every power of it is CBC, indeed we have:

Theorem 5.13. Given a DAG $\mathbf{H} = (V, E)$, the following are equivalent:

1. $\mathbf{H}$ admits a Maltsev polymorphism;

2. for every $n \geqslant 1$, the DAG $\mathbf{H}^n$ is CBC, that is, $\mathbf{H}$ is totally rectangular;

3. $\mathbf{H}$ is layered and for any $u, v \in V$ on the same level, we have, for all $k \geqslant 1$ either $u^{\pm k} = v^{\pm k}$ or $u^{\pm k} \cap v^{\pm k} = \emptyset$;

4. $\mathbf{H}$ can be obtained from a dipath as described in Fig. 6.

Proof. The fact that (1) $\Leftrightarrow$ (2) was seen in Lemma 4.5, and Theorem 5.5 says that (1) $\Leftrightarrow$ (4). We now show that (2) $\Leftrightarrow$ (3). Assume that (3) holds, and let $x, y, u, v \in V$ be such that there exists $k > 0$ such that $x \overset{k}{\to} u$, $y \overset{k}{\to} u$ and $y \overset{k}{\to} v$. It follows that $x^k \cap y^k \neq \emptyset$, so we must have $x^k = y^k$. Hence $x \overset{k}{\to} v$, and so $\mathbf{H}$ is totally rectangular. Assume now that $\mathbf{H}$ is totally rectangular. From Theorem 3.3 we know that $\mathbf{H}$ is layered. Let $k > 0$ be arbitrary, and $u, v \in V$ be such that $u^k \cap v^k \neq \emptyset$. Let $b \in u^k \cap v^k$, and let a be such that $a \in u^k$. We have $u \overset{k}{\to} a$, $u \overset{k}{\to} b$ and $v \overset{k}{\to} b$. It follows that, since $\mathbf{H}$ is totally rectangular, $v \overset{k}{\to} a$, and so $u^k \subseteq v^k$. Similarly we can show that $v^k \subseteq u^k$, and so $u^k = v^k$. The case when $k < 0$ can be dealt with similarly, and we can conclude that (3) holds. □
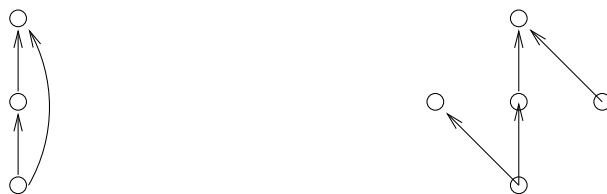

Figure 10: A DAG that is not Maltsev nor QSP, and a conservative Maltsev DAG that is not MSP.

If we restrict ourselves to DAGs that admit a conservative Maltsev polymorphism, then a simpler characterization of these can be obtained. It is not hard to see that all MSP DAGs admit a conservative Maltsev, but the reverse is not true, the DAG on the right-hand side of Fig. 10 admits a conservative Maltsev but is not MSP. In fact, to obtain a characterization of DAGs with conservative Maltsev polymorphism we just need to allow a small generalisation of the construction of MSPs.

Definition 5.14. The class of *Quasi-Minimal Series Parallel DAGs* (QMSP) is the class of DAGs obtained recursively as follows:

- the DAG having a single vertex and no edges is QMSP;

- if $\mathbf{H}_1 = (V_1, E_1)$ and $\mathbf{H}_2 = (V_2, E_2)$ and QMSP, then so are the DAGs obtained from $\mathbf{H}_1$ and $\mathbf{H}_2$, by applying the following operations:

  - parallel composition, and
  - *End-quasi-series composition*: $\mathbf{H} = (V_1 \cup V_2, E_1 \cup E_2 \cup (L_1 \times F_2))$, where L_1 is a subset of the set of sinks on the last level of $\mathbf{H}_1$, and F_2 is a subset of the set of sources on the first level of $\mathbf{H}_2$.

We finally give a characterization, similar to the ones presented earlier, of DAGs admitting a conservative Maltsev polymorphism.

Theorem 5.15. Let $\mathbf{H} = (V, E)$ be a DAG. The following are equivalent:

1. $\mathbf{H}$ admits a conservative Maltsev polymorphism;

2. $\mathbf{H}$ is layered, and for any two vertices u, v of $\mathbf{H}$, on the same level, we have, for all $i \geqslant 1$, either $u^{\pm i} = v^{\pm i}$ or $u^{\pm i} \cap v^{\pm i} = \emptyset$, and

   a) if $u^{+(i-1)} \cap v^{+(i-1)} = \emptyset$ and $u^{+i} = v^{+i}$ then $(u^{+(i-1)} \cup v^{+(i-1)}) \times u^{+i} \subseteq E$, for all $i > 1$,

   b) if $u^{-(j+1)} \cap v^{-(j+1)} = \emptyset$ and $u^{-j} = v^{-j}$ then $u^{-j} \times (u^{-(j+1)} \cup v^{-(j+1)}) \subseteq E$, for all $j > 1$;

3. $\mathbf{H}$ is QMSP;

4. $\mathbf{H}$ is universally rectangular.

Before proving the theorem, we need the following two lemmas.

Lemma 5.16. *Let $\mathbf{H}$ be a connected DAG that admits a conservative Maltsev polymorphism. There exists $i \geqslant 0$ such that all vertices of $\mathbf{H}$ on levels $i, i + 1$ that are neither sources on level $i + 1$, nor sinks on level i, form a complete bipartite graph.*

Proof. Recall that $\mathbf{H}$ is layered, and assume that it has levels $0, \ldots, n$. The proof is by induction on n. If $n = 0$ then the result trivially holds. Assume now that the statement is true for all connected DAGs admitting a conservative Maltsev polymorphism, and with at most $n - 1$ levels. Since $\mathbf{H}$ is preserved by a conservative Maltsev polymorphism, the subgraph $\mathbf{H}'$ induced by the vertices on levels $0, 1, \ldots, n - 1$ of $\mathbf{H}$, also admits a conservative Maltsev polymorphism. If $\mathbf{H}'$ is connected then the result follows by the inductive hypothesis. So, assume now that $\mathbf{H}'$ is not connected. Let $\mathbf{H}_1, \mathbf{H}_2$ be two connected components of $\mathbf{H}'$. These components must be connected by an element on level n in $\mathbf{H}$, since $\mathbf{H}$ is a connected DAG. Let a be an element of level n of $\mathbf{H}$ that connects these components. If $\mathbf{H}$ only has one element on level n, then it follows that the vertices of $\mathbf{H}$ on levels $n - 1$ and n form a complete bipartite DAG, so we assume now that there is at least another element, say b, on level n of $\mathbf{H}$ and, for a contradiction, assume that $a^{-1} \cap b^{-1} = \emptyset$. We can assume, without loss of generality, that b is connected to $\mathbf{H}_2$. So we have that there exist distinct elements v, v' belonging to $\mathbf{H}_2$ and on level $n - 1$ of $\mathbf{H}$ (that is, on the top level of $\mathbf{H}_2$) such that $(v, a), (v', b)$ are arcs of $\mathbf{H}$. Since $\mathbf{H}_2$ is connected, there exists an element p in $\mathbf{H}_2$, say on level $k < n - 1$ of $\mathbf{H}$, such that $p \xrightarrow{n-1-k} v$ and $p \xrightarrow{n-1-k} v'$. Choose p such that these paths are distinct (that is, p is there only common vertex) and as short as possible. This means that v, v' meet in p and do not meet earlier, and consequently the levels of $\mathbf{H}_2$ that form the complete bipartite DAG (which we know exist by the inductive hypothesis) appear below level k of $\mathbf{H}$. If $\mathbf{H}$ has no elements on level k that belong to the connected component $\mathbf{H}_1$ then the DAG $\mathbf{H}_1 \cup \mathbf{H}_2 \cup \{a, b\}$ has a complete bipartite induced subgraph on the same levels as $\mathbf{H}_2$, indeed the same complete bipartite component. If this is the case we just need to deal with other remaining connected components and elements on level n of $\mathbf{H}$ in the same way, which proves the claim.

So, assume now that $\mathbf{H}$ has an element on level k, say q, that belongs to $\mathbf{H}_1$, so there is a vertex u on $\mathbf{H}_1$ such that (u, a) is an arc of $\mathbf{H}$ and $q \xrightarrow{n-1-k} u$. So we have $q \xrightarrow{n-k} a$, $p \xrightarrow{n-k} a$ and $p \xrightarrow{n-k} b$. Since $\mathbf{H}$ admits a conservative Maltsev polymorphism there exists a path of length $n - k$ from q to b using only the elements on the paths from q to a, p to a and p to b. Since $\mathbf{H}_1$ and $\mathbf{H}_2$ are disconnected, we must have that this path is all in $\mathbf{H}_1$ up to level $n - 1$ and so (u, b) must be an arc of $\mathbf{H}$, which contradicts the fact that $a^{-1} \cap b^{-1} = \emptyset$.

Thus it follows that $a^{-1} = b^{-1}$, that is, $a^{-1} \times \{a, b\}$ forms a complete bipartite DAG, on levels $n - 1$ and n. This can clearly be extended to more connected components of $\mathbf{H}$ and other elements on level n, which proves the claim. $\square$

Lemma 5.17. *If a DAG $\mathbf{H}$ admits a conservative Maltsev polymorphism then it is QMSP.*

Proof. Let $\mathbf{H}$ be a DAG, and assume that it has a conservative Maltsev polymorphism. We then know that it retracts to a directed path of some length, say $n + 1$, and so we know that the vertices of $\mathbf{H}$ can be placed into levels $0, 1, \ldots, n$, in such a way that edges are placed between consecutive levels.

We show that $\mathbf{H}$ can be obtained from the one vertex and no edges graph, by a sequence of parallel and quasi-series constructions. Clearly all connected components are

joined together by parallel construction, so let $\mathbf{G}$ be a (maximal) connected component of $\mathbf{H}$. By Lemma 5.16 we know that each connected component of $\mathbf{H}$ can be obtained by a quasi-series construction from two smaller DAGs, $\mathbf{H}_1$ and $\mathbf{H}_2$. Clearly $\mathbf{H}_1$ and $\mathbf{H}_2$ are DAGs that admit a conservative Maltsev polymorphism, since they are induced subgraphs of $\mathbf{H}$. So, we can now look at the connected components of each of these graphs, and repeat the process. We can then conclude that $\mathbf{H}$ can be obtained by means of parallel and quasi-series constructions from the DAG with a single vertex. $\qquad \square$

Proof of Theorem 5.15. In Lemma 5.17, we have seen that $(1) \Rightarrow (3)$. We now prove that $(3) \Rightarrow (2)$. Assume that $\mathbf{H}$ is QMSP. We show by induction on the number of steps of the quasi-series parallel construction of $\mathbf{H}$, that condition (2) is satisfied. The base case trivially holds, and parallel constructions do not interfere with neighbourhoods. It is also easy to see that quasi-series construction also preserves condition (2).

Finally, assume that condition (2) holds, and let us see that this implies (1), that is, that $\mathbf{H}$ (satisfying (2)) admits a conservative Maltsev polymorphism. From condition 2 a) it follows immediately that $\mathbf{H}$ admits a Maltsev polymorphism, that is, it is totally rectangular. We now show that $\mathbf{H}$ is universally rectangular. Suppose this is not the case, then we have distinct paths $x \to x_1 \to \cdots \to x_{k-1} \to u$, $y \to y_1 \to \cdots \to y_{k-1} \to u$ and $y \to z_1 \to \cdots \to z_{k-1} \to v$, and a path $x \to w_1 \to \cdots \to w_{k-1} \to v$, and at least one of the $w_i's$ is different from x_i, y_i, z_i. We can assume, without loss of generality, that this happens with w_1, and that x_1, y_1, z_1 are all different. We have $u^{-k} = v^{-k}$, and $u^{-k+1} \cap v^{-k+1} = \emptyset$, since $x_1 \in u^{-k+1}$ and $x_1 \notin v^{-k+1}$. Note that, if $x_1 \in v^{-k+1}$ then we could have set $w_1 = x_1$. It then follows, from (2), that $u^{-k} \times (u^{-k+1} \cup v^{-k+1}) \subseteq E(\mathbf{H})$. In particular $x_1 \to y_1$, so we could have set $w_1 = y_1$, contradiction. This proves the theorem. $\qquad \square$

It is known that GSP digraphs can be recognized in linear time. In [25] it is given an algorithm that given a DAG $\mathbf{H}$, with n vertices and m edges, recognizes in $O(n + m)$ if $\mathbf{H}$ is GSP. We leave as an open question if Maltsev DAGs (or conservative Maltsev DAGs) can also be recognized in linear time. In [25] it is also mentioned an algorithm that checks in time linear to the size of the input, if $\mathbf{H}$ is CBC. This allows us to check if a given DAG, with n vertices and m edges, is preserved by a Maltsev polymorphism in $O(l(n + m))$-time, where l is the number of levels of the digraph.

6   Some Applications to the Constraint Satisfaction Problem

The logic programming language Datalog is one of the main tools to solve CSPs in P. The fragments of Datalog called linear and symmetric Datalog are conjectured to contain all CSPs in NL and L, respectively, see [7, 10, 17]. A minor technicality is that it is actually the complement of a CSP that can be defined in Datalog and its fragments, not the actual CSP.

By Theorem 3.3, the core of a Maltsev digraph is either a directed path or a disjoint union of cycles, and for such digraphs, Corollary 6.1 is not difficult to show. For the proof below, note that the relevant definitions can be found, for example, in [17]. Furthermore,

if $0 \leqslant j \leqslant k$, then a symmetric Datalog program is said to have *width* (j, k) if each rule of the program has at most k variables, and at most j variables in the head.

Corollary 6.1. *Let H be a Maltsev digraph. Then the complement of $\mathrm{CSP}(H)$ can be defined in symmetric Datalog of width $(2, 3)$, and therefore $\mathrm{CSP}(H)$ is in* L.

Proof Sketch. By Theorem 3.3, we can assume that H is a dipath or a disjoint union of cycles. Assume first that H is the disjoint union of directed cycles. Then a digraph G does not map to H if and only if G contains a connected component C that does not map to any directed cycles in H. It is easy to see that C does not map to a directed cycle of length ℓ if and only if C contains a closed walk $v_1, v_2, \ldots, v_1$ such that its net length is not $0 \bmod \ell$. This is easy to detect with a symmetric Datalog program of width $(2, 3)$. For example, the program below works for $\mathrm{CSP}(H)$, where H is the disjoint union of a directed 2-cycle and 3-cycle. It is straightforward to generalize this program to any set of cycles and to give a formal proof of correctness by induction. (For brevity, we did not include the symmetric pairs of the rules in the program below. Observe that those rules are easily seen to be harmless.)

$$
\begin{array}{llll}
I_1(s, x) \leftarrow E(s, x) & (1) & J_0(s, s) \leftarrow K(s) & (10) \\
I_1(s, x) \leftarrow E(x, s) & (2) & J_1(s, x) \leftarrow J_0(s, y) \wedge E(y, x) & (11) \\
I_0(s, x) \leftarrow I_1(s, y) \wedge E(y, x) & (3) & J_2(s, x) \leftarrow J_0(s, y) \wedge E(x, y) & (12) \\
I_0(s, x) \leftarrow I_1(s, y) \wedge E(x, y) & (4) & J_2(s, x) \leftarrow J_1(s, y) \wedge E(y, x) & (13) \\
I_1(s, x) \leftarrow I_0(s, y) \wedge E(y, x) & (5) & J_0(s, x) \leftarrow J_1(s, y) \wedge E(x, y) & (14) \\
I_1(s, x) \leftarrow I_0(s, y) \wedge E(x, y) & (6) & J_0(s, x) \leftarrow J_2(s, y) \wedge E(y, x) & (15) \\
K(x) \leftarrow I_1(x, x) & (7) & J_1(s, x) \leftarrow J_2(s, y) \wedge E(x, y) & (16) \\
K(x) \leftarrow K(y) \wedge E(y, x) & (8) & Goal \leftarrow J_1(x, x) & (17) \\
K(x) \leftarrow K(y) \wedge E(x, y) & (9) & Goal \leftarrow J_2(x, x) & (18)
\end{array}
$$

In the program above, EDB E is the edge relation of the input graph, and *Goal* is the goal predicate. IDB $I_1(a, b)$ ($I_0(a, b)$) is derived if and only if there is a walk of odd (even) net length from a to b. Not that rule (7) derives $K(c)$ if and only if there is a walk of odd net length starting at a and ending at c. At this stage, we know that the connected component C of the input digraph that contains c has an odd cycle, so C cannot map to the 2-cycle. Using rules (8) and (9), the program can walk to any vertex of C. In the second part of the program, for example, the role of rule (12) is the following. Assume that the variables of rule (12) are instantiated to vertives of G as $J_2(d, f) \leftarrow J_0(d, e) \wedge E(f, e)$. Then since $J_0(d, e)$ is derived, we know that there is a walk from d to e of net length $0 \bmod 3$. Since $E(f, e)$ is an edge of G, we know that there is a walk from d to f of net length $2 \bmod 3$, so the head $J_2(d, f)$ is derived.

Similar ideas work when H is a directed path. For example, if H is a dipath of length 2, then we have to detect if G contains a walk of net length at least 3. If so, our program should accept, indicating that there is no homomorphism from G to H. This can be done

with a program of width $(1, 2)$. If no such walk exists, then we still have to check that every walk from a vertex to itself has net length 0. This can be easily done with a program of width $(2, 3)$, similar in spirit to the program above. $\qquad\square$

The *list homomorphism problem* for a digraph H, LHOM(H), is the following decision problem. Given an input digraph G and for each vertex $v \in V_G$ a list $L_v \subseteq V_H$, determine if there is a homomorphism h from G to H such that for each $v \in V_G$, $h(v) \in L_v$. This problem is exactly CSP(H_u) where H_u is the structure obtained by expanding the digraph H with unary relations U, where U runs through all non-empty subsets of V_H . Using Corollary 4.11, the following corollary is easy to deduce.

Corollary 6.2. *The complement of* LHOM(H) *for a conservative Maltsev digraph H can be defined in symmetric Datalog, and therefore* LHOM(H) *is in* L.

Proof. Let H be the digraph in the corollary. Assume the conservative Maltsev polymorphism is mal. Then by Corollary 4.11, H also has a conservative majority polymorphism maj. Since mal and maj are conservative, they preserve any unary relations. Construct **B** from H as above. Then mal and maj is also a polymorphism of **B**. If a relational structure has a Maltsev and a majority polymorphism, then it is in symmetric Datalog by results of [12, 11] (see also [14]). $\qquad\square$

References

[1] L. Barto and J. Bulin. CSP dichotomy for special polyads. *International Journal of Algebra and Computation 23/05*, pages 1151–1174, 2013.

[2] L. Barto, M. Kozik, M. Maróti, and T. Niven. CSP dichotomy for special triads. *Proceedings of the AMS*, 137:2921–2934, 2009.

[3] L. Barto, M. Kozik, and T. Niven. Graphs, polymorphisms and the complexity of homomorphism problems. In *Proceedings of the 40th annual ACM symposium on Theory of Computing*, STOC '08, pages 789–796, 2008.

[4] A. Bulatov and V. Dalmau. A simple algorithm for Mal'tsev constraints. *SIAM Journal on Computing*, 36(1):16–27, 2006.

[5] A. Bulatov, P. Jeavons, and A. Krokhin. Constraint satisfaction problems and finite algebras. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, ICALP '00, pages 272–282, 2000.

[6] A. Bulatov, P. Jeavons, and A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.

[7] A. Bulatov, A. Krokhin, and B. Larose. Dualities for constraint satisfaction problems. In *LNCS Surveys on Complexity of Constraints*, volume 5250, pages 93–124. 2008.

[8] J. Bulin, D. Delic, M. Jackson, and T. Niven. On the reduction of CSP dichotomy to digraphs. In *Theory and Practice of Constraint Programming*, CP2013, pages 184–199, 2013.

[9] C. Carvalho, L. Egri, M. Jackson, and T. Niven. On maltsev digraphs. In *Computer Science–Theory and Applications*, pages 181–194. Springer, 2011.

[10] V. Dalmau. Linear Datalog and bounded path duality of relational structures. *Logical Methods in Computer Science*, 1:1–32, 2005.

[11] V. Dalmau and A. Krokhin. Majority constraints have bounded pathwidth duality. *European Journal of Combinatorics*, 29(4):821–837, 2008.

[12] V. Dalmau and B. Larose. Maltsev + Datalog ⇒ Symmetric Datalog. In *Proceedings of the 23rd IEEE Symposium on Logic in Computer Science*, LICS '08, pages 297–306, 2008.

[13] M. E. Dyer and D. Richerby. On the complexity of #CSP. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, STOC '10, pages 725–734, 2010.

[14] L. Egri. On constraint satisfaction problems below P. *Journal of Logic and Computation*, 2014.

[15] L. Egri, P. Hell, B. Larose, and A. Rafiey. Space complexity of list H-colouring: a dichotomy. SODA 2014, pages 349–365.

[16] L. Egri, A. Krokhin, B. Larose, and P. Tesson. The complexity of the list homomorphism problem for graphs. In *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science*, STACS '10, pages 335–346, 2010.

[17] L. Egri, B. Larose, and P. Tesson. Symmetric Datalog and constraint satisfaction problems in logspace. In *Proceedings of the 22nd Annual IEEE Symposium on Logic in Computer Science*, LICS '07, pages 193–202, 2007.

[18] J. Hagemann and A. Mitschke. On n-permutable congruences. *Algebra Universalis*, 3:8–12, 1973.

[19] P. Hell and J. Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990.

[20] P. Hell and A. Rafiey. The dichotomy of list homomorphisms for digraphs. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 1703–1713, 2011.

[21] P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997.

[22] A. Kazda. Maltsev digraphs have a majority polymorphism. *European Journal of Combinatorics*, 32:390–397, 2011.

[23] M. Maróti and R. McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 59(3-4):463–489, 2008.

[24] R. H. Möhring. Computationally tractable classes of ordered sets. In *Algorithms and order (Ottawa, ON, 1987)*, pages 105–193. Kluwer Acad. Publ., Dordrecht, 1989.

[25] J. Valdes, R. E. Tarjan, and E. L. Lawler. The recognition of series parallel digraphs. *SIAM J. Comput.*, 11(2):298–313, 1982.