COMBINATORICA
Bolyai Society – Springer-Verlag

COMBINATORICA **42** Supplement Issue (1) (2022) 997–1010
DOI: 10.1007/s00493-022-4918-1

MAXIMAL DIGRAPHS WITH RESPECT TO PRIMITIVE POSITIVE CONSTRUCTABILITY

MANUEL BODIRSKY, FLORIAN STARKE*

Received April 15, 2021
Revised December 1, 2021
Online First October 10, 2022

We study the class of all finite directed graphs (digraphs) up to primitive positive constructibility. The resulting order has a unique maximal element, namely the digraph P_1 with one vertex and no edges. The digraph P_1 has a unique maximal lower bound, namely the digraph P_2 with two vertices and one directed edge. Our main result is a complete description of the maximal lower bounds of P_2; we call these digraphs *submaximal*. We show that every digraph that is not equivalent to P_1 and P_2 is below one of the submaximal digraphs.

1. Introduction

A *homomorphism* from a directed graph G to a directed graph H is a map from the vertices of G to the vertices of H which maps each edge of G to an edge of H. Two directed graphs G and H are called *homomorphically equivalent* if there is a homomorphism from G to H and from H to G. The study of the *homomorphism order* on the class of all finite directed graphs (or short: *digraphs*), factored by homomorphic equivalence, has a long history in graph theory. It is known to have a quite complicated structure; we refer to Nešetřil and Tardif [1] and the references therein.

A classical topic in graph homomorphisms is the H-coloring problem, which is the computational problem of deciding whether a given finite digraph G maps homomorphically to H. The computational complexity of this problem has been classified for finite undirected graphs H by Hell and

Mathematics Subject Classification (2010): 05C15, 05C20

* The second author is supported by DFG Graduiertenkolleg 1763 (QuantLA).

Nešetřil [2] in 1990: they are either in the complexity class L (i.e., they can be decided deterministically with logarithmic work space) or NP-complete. Feder and Vardi [3] proved that every finite-domain CSP is polynomial-time equivalent to an H-coloring problem for a finite *directed* graph H[1], and they conjectured that each of these problems are either in P or NP-complete. This dichotomy conjecture was eventually solved in 2017 by Bulatov and, independently, by Zhuk [5,6]. However, other long-standing open problems about the complexity of H-coloring for finite digraphs H remain open, for example, the characterisation of when this problem is in L, or in NL [7,8,9].

The border between polynomial-time tractable and NP-complete H-colouring problems can be described in terms of *primitive positive (pp) constructions*, which is a concept that has been introduced by Barto, Opršal, and Pinsker [10] in the setting of general relational structures. The idea is that if G has a pp construction in H, then, intuitively, *'G can be simulated by H'*, and the G-coloring problem reduces (in logarithmic space) to the H-coloring problem. In particular, H-coloring is NP-hard if K_3 has a pp construction in H, where K_3 is the clique with three vertices, by reduction from the NP-hard three-colorability problem. It follows from the dichotomy theorem of Bulatov and Zhuk that otherwise H-coloring is in P. Note that pp constructibility can also be used to study the question of which H-coloring problems are in L or in NL. The surprising power of pp constructions is the motivation for studying pp constructions on finite digraphs more systematically.

For digraphs G and H that have at least one edge, the definition of pp constructions takes the following elegant combinatorial form: G has a pp construction in H if there exists a digraph K and $a,b \in V(K)^d$ for some $d \in \mathbb{N}$ such that G is homomorphically equivalent to the digraph with vertices $V(H)^d$ and where (u,v) forms an edge if there is a homomorphism from K to H that maps $a_1,\dots,a_d,b_1,\dots,b_d$ to $u_1,\dots,u_d,v_1,\dots,v_d$, respectively. We write $H \leq G$ if G has a pp construction in H; we deliberately chose the symbol $\leq$ rather than $\geq$; the motivation will become clear in Section 2. It can be shown that $\leq$ is transitive (Corollary 3.10 in [10]) and so it gives rise to a partial order $\mathfrak{P}_{\text{Digraphs}}$ on the class of all finite digraphs (where we take the liberty to identify two digraphs G and H if they pp construct each other). Since all finite digraphs have a pp construction in K_3 (see, e.g. [11]), it is the smallest element of the poset $\mathfrak{P}_{\text{Digraphs}}$. For $n \geq 1$ the directed path of length n is the digraph $P_n := (\mathbb{Z}_n, \{(u,u+1) \mid 0 \leq u < n-1\})$. The poset $\mathfrak{P}_{\text{Digraphs}}$ also has a greatest element, namely the digraph P_1. The digraph P_1 has a unique maximal lower bound, namely the digraph P_2, which is, in

[1] This result has been sharpened in [4].

$\mathfrak{P}_{\text{Digraphs}}$, equivalent to P_n for any $n \geq 2$; this is not hard to see and will be shown in Section 3.

In this article, we present a complete description of the maximal lower bounds of P_2 in $\mathfrak{P}_{\text{Digraphs}}$; we call these digraphs *submaximal.* We also prove that every finite digraph which does not pp constructs P_2 is smaller than one of the submaximal digraphs (Theorem 3.5; also see Figure 1). The submaximal digraphs are:

- The directed cycles C_p for p prime. (For $k \in \mathbb{N}^+$, the directed cycle C_k is defined to be the digraph $(\mathbb{Z}_k, \{(u, u+1 \mod k) \mid u \in \mathbb{Z}_k\})$.)
- $T_3 := (\{0,1,2\}, <)$, the transitive tournament with three vertices.

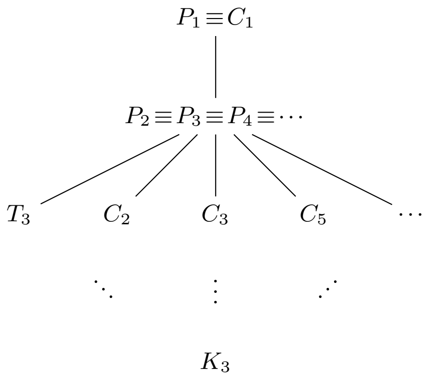

Figure 1. The pp constructibility poset on finite digraphs

Related work

The pp constructibility poset for smooth digraphs, i.e., digraphs where every vertex has indegree at least one and outdegree at least one (digraphs without sources and sinks), has been described in [11]. The pp constructibility poset on general relational structures over a two-element set has been described in [12].

2. Minor conditions

Primitive positive constructibility has a universal algebraic characterisation; this characterisation plays a role in our proof, so we present it here. If $H =$

(V, E) is a digraph, then H^k denotes the k-th direct power of H, which is the digraph with vertex set V^k and edges set

$$\{((u_1, \ldots, u_k), (v_1, \ldots, v_k)) \mid (u_1, v_1) \in E, \ldots, (u_k, v_k) \in E\}.$$

A *polymorphism* of H is a homomorphism f from H^k to H, for some $k \in \mathbb{N}$, which is called the *arity* of f. We write $\mathrm{Pol}(H)$ for the set of all polymorphisms of H. This set contains the projections and is closed under composition.[2] An operation f is called *idempotent* if $f(x,\ldots,x) = x$ for all $x \in V$.

A central topic in universal algebra are *minor conditions*. If $f\colon V^k \to V$ is an operation and $\sigma\colon \{1,\ldots,k\} \to \{1,\ldots,n\}$ is a function, then f_σ denotes the operation

$$(x_1, \ldots, x_n) \mapsto f(x_{\sigma(1)}, \ldots, x_{\sigma(k)}),$$

and f_σ is called a *minor* of f. A *minor condition* is a set Σ of expressions of the form $f_\sigma = g_\tau$ where f and g are function symbols (f and g might be the same symbol) and $\sigma\colon \{1,\ldots,k\} \to \{1,\ldots,n\}$, $\tau\colon \{1,\ldots,\ell\} \to \{1,\ldots,n\}$ are functions.

Example 2.1. An operation $f\colon V^n \to V$ is called *cyclic* if for all $x_1,\ldots,x_n \in V$

$$f(x_1, x_2, \ldots, x_n) = f(x_2, \ldots, x_n, x_1).$$

This condition can be expressed by the minor condition

$$\Sigma_n := \{f_{\mathrm{id}} = f_\tau\},$$

where id denotes the identity function on $\{1, 2, \ldots, n\}$ and τ denotes the cyclic permutation $(1, 2, \ldots, n)$ on $\{1, \ldots, n\}$.

If a minor condition Σ contains several expressions, then different expressions in Σ might share the same function symbols.

Example 2.2. An idempotent operation f is called a *Maltsev operation* if for all $x, y \in V$

$$f(y, y, x) = f(x, x, x) = f(x, y, y).$$

This condition can be expressed by the minor condition

$$\Sigma_M := \{f_\sigma = f_\tau, f_\tau = f_\rho\},$$

where $\sigma, \tau, \rho\colon \{1, 2, 3\} \to \{1, 2\}$ are given by $\sigma(1, 2, 3) = (2, 2, 1)$, $\tau(1, 2, 3) = (1, 1, 1)$, and $\rho(1, 2, 3) = (1, 2, 2)$.

[2] Sets of operations with these properties are called *clones* in universal algebra.

A set of operations F *satisfies* a minor condition Σ if the function symbols in Σ can be replaced by operations from F so that all the expressions in Σ hold; in this case we write $F \models \Sigma$. If H is a digraph, then $\Sigma(H)$ denotes the class of all minor conditions that are satisfied in $\mathrm{Pol}(H)$.

Theorem 2.3 (Barto, Opršal, and Pinsker [10]). *Let G and H be finite digraphs. Then*

$$H \text{ pp constructs } G \qquad \text{if and only if} \qquad \Sigma(H) \subseteq \Sigma(G).$$

3. The pp construction poset

We have already defined pp constructibility for digraphs in the introduction, but present an equivalent description here which is convenient when specifying pp constructions, and which is also closer to the presentation of Barto, Opršal, and Pinsker [10]. A *primitive positive formula* is a formula $\phi(x_1,\dots,x_k)$ of the form

$$\exists y_1,\dots,y_n(\psi_1 \wedge \cdots \wedge \psi_m),$$

where each of the formulas $\psi_1,\dots,\psi_m$ is of the form $\perp$ (for *false*), of the form $z_1 = z_2$, or of the form $E(z_1,z_2)$ where z_1,z_2 are variables from $\{x_1,\dots,x_k,y_1,\dots,y_n\}$.

Definition 3.1. Let $H=(V,E)$ be a digraph. A digraph G with vertex set V^d is called a *pp power of H of dimension d* if there exists a primitive positive formula $\phi(x_1,\dots,x_d,y_1,\dots,y_d)$ such that the edge set of G equals

$$\{((u_1,\dots,u_d),(v_1,\dots,v_d)) \mid \phi(u_1,\dots,u_d,v_1,\dots,v_d) \text{ holds in } H\}.$$

It follows from the definitions that $H \leq G$ if and only if G is homomorphically equivalent to a pp power of H. We write

- $H \equiv G$ if $H \leq G$ and $G \leq H$;
- $H < G$ if $H \leq G$ and not $G \leq H$.

Lemma 3.2. *P_1 is the greatest element of $\mathfrak{P}_{\text{Digraphs}}$. Moreover, $P_1 \equiv C_1$.*

Proof. Let G be a finite digraph. Consider the pp power of G of dimension one given by the formula $\phi(x,y) := \perp$. The resulting digraph has no edges and is therefore homomorphically equivalent to P_1. Now consider the pp power of G of dimension one given by the formula $\phi(x,y) := (x=y)$. The resulting digraph is homomorphically equivalent to the digraph C_1 with one vertex and a loop, which implies the statement. ∎

In the proof of the following lemma we need the fundamental concept of *cores* from the theory of graph homomorphisms (see, e.g., [13]). A digraph $H=(V,E)$ is called a *core* if every *endomorphism* of H (i.e., every homomorphism from H to H) is an *embedding* (i.e., an isomorphism between H and an induced subgraph of H; for background, see, e.g., [14]). It is easy to see that every finite digraph H is homomorphically equivalent to a core digraph, and that all finite core digraphs G that are homomorphically equivalent to H are isomorphic to each other; we therefore call G *the* core of H. When studying $\mathfrak{P}_{\text{Digraphs}}$ we may therefore restrict our attention to core digraphs; the big advantage of cores is the following useful lemma.

Lemma 3.3 (follows from Lemma 3.9 in [10]). *Let $H = (V,E)$ be a finite core digraph. Then $H \leq G$ if and only if G is homomorphically equivalent to a pp power of H where the primitive positive formula might additionally contain conjuncts of the form $x=c$ where x is a variable and $c\in V$ is a constant.*

Lemma 3.4. *We have $P_2 < P_1$. Moreover, P_2 is the only coatom of $\mathfrak{P}_{\text{Digraphs}}$, i.e., P_2 is the unique maximal lower bound of P_1 in $\mathfrak{P}_{\text{Digraphs}}$.*

Proof. We have already seen that $P_2 \leq P_1$. To prove that $P_2 \not\geq P_1$, first observe that P_1 has constant polymorphisms, while P_2 does not. Let $\Sigma_c :=$ $\{f_\rho = f_\sigma\}$ where f is a unary function symbol, $\rho\colon \{1\} \to \{1,2\}$, $1 \mapsto 1$ and $\sigma\colon \{1\} \to \{1,2\}$, $1,\mapsto 2$. Then $\mathrm{Pol}(P_1) \models \Sigma_c$, but $\mathrm{Pol}(P_2) \not\models \Sigma_c$. Then (the easy direction of) Theorem 2.3 implies that $P_1 \leq P_2$ does not hold.

For the second statement, let G be a finite digraph such that $G<P_1$. We have to show that $G\leq P_2$. Without loss of generality we may assume that G is a core. Hence, by Lemma 3.3, we can use constants in pp constructions. Note that G must have at least two different vertices u and v. The pp power of G of dimension one given by the formula $\phi(x,y) := (x=u)\wedge(y=v)$ is a digraph that has exactly one edge, and this edge is not a loop; therefore the graph is homomorphically equivalent to P_2. ∎

The following theorem is our main result and will be shown in the remainder of the article; see Figure 1.

Theorem 3.5. *The submaximal elements of $\mathfrak{P}_{\text{Digraphs}}$ are precisely T_3, C_2, C_3, C_5, ... If G is a finite digraph that does not have a pp construction in P_2, then $G\leq T_3$ or $G\leq C_p$ for some prime p.*

4. Submaximal digraphs and minor conditions

We first discuss which of the minor conditions that we have encountered are satisfied by the polymorphisms of the digraphs that appear in Theorem 3.5. The facts presented in this section are well-known; we present the proof for the convenience of the reader.

Lemma 4.1. *Let p and q be primes. Then* $\mathrm{Pol}(C_p) \models \Sigma_q$ *(introduced in Example 2.1) if and only if* $p \neq q$.

Proof. If $p \neq q$, then there is an $n \in \mathbb{N}^+$ such that $q \cdot n = 1 \pmod p$. The map

$$(x_1, \ldots, x_q) \mapsto n \cdot (x_1 + \ldots + x_q) \pmod p$$

is a polymorphism of C_p satisfying Σ_q.

Now suppose that $p = q$ and assume for contradiction that f is a polymorphism of C_p satisfying Σ_p. Then

$$f(0, \ldots, p-2, p-1) = a = f(1, \ldots, p-1, 0)$$

and hence $(a, a) \in E$, which is impossible since C_p has a loop only if $p = 1$. ∎

Lemma 4.2. $\mathrm{Pol}(C_n) \models \Sigma_M$ *for every* $n \in \mathbb{N}$.

Proof. The ternary operation $(x_1, x_2, x_3) \mapsto x_1 - x_2 + x_3 \pmod n$ is a Maltsev polymorphism of C_n. ∎

Let $H = (V, E)$ be a finite digraph, $u, v \in V$, and $k \in \mathbb{N}$. A *directed walk* of length k from u to v is a k-tuple $(v_0, \ldots, v_{k-1}) \in V^k$ such that $v_0 = u$, $v_{k-1} = v$, and $(v_i, v_{i+1}) \in E$ for all $i \in \{0, \ldots, k-2\}$. The digraph H is called *k-rectangular* if whenever H has directed walks of length k from a to b, from c to b, and from c to d, then also from a to d. See Figure 2. A digraph H is called *totally rectangular* if it is k-rectangular for all $k \geq 1$. The following well-known lemma connects total rectangularity with Σ_M.

Lemma 4.3. *A finite digraph H is totally rectangular if and only if it has a Maltsev polymorphism. A finite core digraph H has a Maltsev polymorphism if and only if* $\mathrm{Pol}(H) \models \Sigma_M$.

Proof. The first part of the statement is Corollary 4.11 in [15]. For the second statement, let $H = (V, E)$ be a core digraph which has a polymorphism f that satisfies $f(x, y, y) = f(x, x, x) = f(y, y, x)$ for all $x, y \in V$; we have to find a polymorphism that is additionally idempotent. Note that the function $x \mapsto f(x, x, x)$ is an endomorphism; since H is a core, the endomorphism is injective. Since H is finite the endomorphism must in fact be an automorphism, and has an inverse i which is an endomorphism as well. Then the operation $(x_1,x_2,x_3) \mapsto i(f(x_1,x_2,x_3))$ is idempotent and a Maltsev operation. ∎

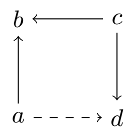

Figure 2. Rectangularity in digraphs

Lemma 4.4. $\mathrm{Pol}(T_3) \models \Sigma_n$ *for every* $n \in \mathbb{N}$, *but* $\mathrm{Pol}(T_3) \not\models \Sigma_M$.

Proof. The operation $(x_1,\dots,x_n) \mapsto \max(x_1,\dots,x_n)$ is a polymorphism of T_3 that satisfies Σ_n. On the other hand, $T_3 = (\{0,1,2\},E)$ is not 1-rectangular, witnessed by $(1,2),(0,2),(0,1) \in E$ but $(1,1) \notin E$; the second statement therefore follows from Lemma 4.3. ∎

The following theorem states that the digraph P_2 is the unique smallest element of $\mathfrak{P}_{\mathrm{Digraphs}}$ that satisfies Σ_M and Σ_p for all p prime.

Theorem 4.5. *Let G be a finite digraph that satisfies Σ_M and Σ_p for all primes p. Then $P_2 \leq G$.*

In the proof of Theorem 4.5 we make use the following result of Carvalho, Egri, Jackson, and Niven [15], which guides us in our further proof steps.

Theorem 4.6 (Lemma 3.10 in [15]). *If G is totally rectangular, then G is homomorphically equivalent to either a directed path or a disjoint union of directed cycles.*

Before we come to the proof of Theorem 5.4 we show that P_2 can pp construct all other directed paths.

Lemma 4.7. *The digraph P_2 pp constructs P_k for all $k \in \mathbb{N}^+$.*

Proof. Clearly, $P_2 \leq P_1$ and $P_2 \leq P_2$. Let $k \geq 3$ and consider the pp power G of P_2 of dimension $k-1$ given by the following formula

$\phi(x_1,\ldots x_{k-1},y_1,\ldots,y_{k-1})$

$$(x_1 = y_2) \wedge (x_2 = y_3) \wedge \ldots \wedge (x_{k-2} = y_{k-1}) \wedge E(x_{k-1}, y_1).$$

Then G contains the following path with k vertices

$$(0,0,\ldots,0) \to (1,0,\ldots,0) \to (1,1,\ldots,0) \to \cdots \to (1,1,\ldots,1),$$

which shows that there exists a homomorphism from P_k to G. Note that whenever there is an edge from u to v in G, then the tuple v contains exactly one 1 more than the tuple u. Therefore, the function $V(G) \to \{0,\ldots,k-1\}$ that maps v to the number of 1's in v is a homomorphism from G to P_k. Hence $P_2 \leq P_k$. ∎

Proof of Theorem 4.5. Let G be a finite digraph satisfying Σ_M and Σ_p for every prime p. By Lemma 4.3 and Theorem 4.6 there are two cases to consider: the first is that G is homomorphically equivalent to P_k for some k. Then $P_2 \leq G$ by Lemma 4.7.

The second case is that G is homomorphically equivalent to a disjoint union of directed cycles. Without loss of generality we may assume that G is a disjoint union of directed cycles. Let $(a_0,\ldots,a_{\ell-1})$ be a shortest cycle in G. Let p be a prime and $k \in \mathbb{N}^+$ such that $p \cdot k = \ell$, and let $f \in \mathrm{Pol}(G)$ be a function that witnesses that $\mathrm{Pol}(G) \models \Sigma_p$. Then

$$f(a_0, a_k, \ldots, a_{(p-1)\cdot k}) = a = f(a_k, a_{2k}, \ldots, a_0).$$

Note that there are directed walks of length k from $a_{(p-1)\cdot k}$ to a_0 and from $a_{i\cdot k}$ to $a_{(i+1)\cdot k}$ for $i \in \{0,\ldots,p-2\}$. Since f is a polymorphism of G there is a directed walk of length k from a to a. Thus, G contains a directed cycle whose length divides k, which contradicts the assumption that ℓ is the length of the shortest directed cycle in G. Therefore, ℓ has no prime divisors, and $\ell = 1$. So G contains a loop and hence is homomorphically equivalent to C_1; it follows that $P_2 \leq G$. ∎

5. Proof of the main result

We use the following general result about when a finite digraph can pp construct a finite disjoint union of cycles. If C is a finite disjoint union of cycles and $c \in \mathbb{N}$, then $C \dot{-} c$ denotes the union of cycles which contains for every cycle of length n in C a cycle of length $n / \gcd(n,c)$. If G is any directed graph with vertices $u_1,\ldots,u_n$ and edges $e_1,\ldots,e_m$, then Σ_G denotes the minor condition $f_\sigma = f_\tau$, where $\sigma,\tau \colon \{1,\ldots,m\} \to \{1,\ldots,n\}$ are such that if $e_i = (u_p,u_q)$, then $\sigma(i) = p$ and $\tau(i) = q$. Note that whether $\mathrm{Pol}(H)$ satisfies

Σ_G does not depend on the choice of the above enumerations of the edges and vertices of G. In particular, the condition Σ_{C_k} is equivalent to the condition Σ_k.

Lemma 5.1 (Lemma 6.8 in [11]). *Let C be a finite disjoint union of cycles and let G be a finite digraph. Then*

$$G \leq C \quad \textit{iff} \quad \mathrm{Pol}(G) \models \Sigma_{C \dot{-} c} \textit{ implies } \mathrm{Pol}(C) \models \Sigma_{C \dot{-} c} \textit{ for all } c \in \mathbb{N}^+.$$

For the special case that $C = C_p$, there are only two conditions of the form $\Sigma_{C \dot{-} c}$, namely Σ_1, which is trivial and hence satisfied by both $\mathrm{Pol}(G)$ and $\mathrm{Pol}(C)$, and Σ_p, which is not satisfied by C_p. Hence, we obtain the following result.

Theorem 5.2. *Let G be a finite digraph. If p is a prime number such that $\mathrm{Pol}(G) \not\models \Sigma_p$, then $G \leq C_p$.*

We also need a similar result for Σ_M instead of Σ_p.

Lemma 5.3. *Let G be a finite digraph. If $\mathrm{Pol}(G) \not\models \Sigma_M$, then $G \leq T_3$.*

Proof. Since $\leq$ is transitive we may assume without loss of generality that $H = (V, E)$ is a core. By Lemma 4.3, H is not totally rectangular. Hence, there are vertices $a, b, c, d \in V$ such that in G there are directed walks of length k from a to b, from c to b, from c to d, and there is no directed walk of length k from a to d. Note that by Lemma 3.3 we are allowed to use constants in pp constructions. We write $x \xrightarrow{k} y$ as a shortcut for the primitive positive formula $\exists u_1, \dots, u_{k-1}(E(x,u_1) \wedge E(u_1,u_2) \wedge \cdots \wedge E(u_{k-1}, y))$. Consider the pp power of G of dimension two given by the formula

$$\phi(x_1, x_2, y_1, y_2) \coloneqq x_1 \xrightarrow{k} y_2 \wedge (x_2 = d) \wedge (y_1 = a).$$

Let H be the resulting digraph. Consider the vertices $v_0 = (c, d)$, $v_1 = (a, d)$, and $v_2 = (a, b)$ of H. Note that the only vertex of H that can have incoming and outgoing edges is v_1. Since there is no directed walk of length k from a to d the vertex v_1 does not have a loop. Furthermore, H has the edges $(v_0, v_1), (v_1, v_2)$, and (v_0, v_2) (see Figure 3). Hence, $i \mapsto v_i$ is an embedding of T_3 into H. Let V_0 be the set of all vertices in H that have outgoing edges and V_2 be the set of all vertices in H that have incoming edges. Let V_1 denote the set $(V_0 \cap V_2) \cup (V(H) \setminus (V_0 \cup V_2))$. Note that V_1 consists of v_1 and all isolated vertices. Clearly, $V_0 \backslash V_2$, V_1, and $V_2 \backslash V_0$ form a partition of $V(H)$ and the map

$$v \mapsto \begin{cases} 2 & \text{if } v \in V_2 \setminus V_0 \\ 1 & \text{if } v \in V_1 \\ 0 & \text{if } v \in V_0 \setminus V_2 \end{cases}$$

is a homomorphism from H to T_3. Hence $G \leq T_3$. ∎

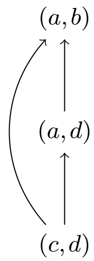

Figure 3. The primitive positive construction of T_3 in the proof of Lemma 5.3

Proof of Theorem 3.5. Let G be a digraph such that $P_2 \not\leq G$. Theorem 4.5 implies that either $\mathrm{Pol}(G)$ does not satisfy Σ_M or that it does not satisfy Σ_p for some prime p. In the first case $G \leq T_3$, by Lemma 5.3. In the second case $G \leq C_p$, by Theorem 5.2. Hence, all submaximal elements of $\mathfrak{P}_{\text{Digraphs}}$ are contained in $\{T_3, C_2, C_3, C_5, \dots\}$. Lemma 4.1, Lemma 4.2, and Lemma 4.4 in combination with Theorem 2.3 imply that these digraphs form an antichain in $\mathfrak{P}_{\text{Digraphs}}$, and hence each of these digraphs is submaximal. ∎

Note that our result implies the following.

Corollary 5.4. *If a finite digraph G satisfies Σ_M, Σ_2, Σ_3, Σ_5, ..., then any minor condition satisfied by* $\mathrm{Pol}(P_2)$ *is also satisfied by* $\mathrm{Pol}(G)$.

The statement of Corollary 5.4 may also be phrased as

$$\{\Sigma_M, \Sigma_2, \Sigma_3, \Sigma_5, \dots\} \subseteq \Sigma(G) \quad \Rightarrow \quad \Sigma(P_2) \subseteq \Sigma(G).$$

Remark 5.5. We do not know whether Corollary 5.4 holds for arbitrary clones of operations on a finite set, instead of just clones of the form $\mathrm{Pol}(G)$ for a finite digraph G. However, the statement is false for clones of operations on an infinite set, as illustrated by the clone of operations on $\mathbb{Q}$ of the form

$(x_1,\dots,x_n)\mapsto a_1x_1+\dots+a_nx_n$ for $a_1,\dots,a_n\in\mathbb{Q}$ such that $a_1+\dots+a_n=1$. This clone satisfies Σ_n for every $n\in\mathbb{N}$, and contains the function $(x_1,x_2,x_3)\mapsto x_1-x_2+x_3$, so it also satisfies Σ_M. However, it is easy to see that it does not contain an operation f that satisfies

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

for all $x,y\in\mathbb{Q}$; however, this minor condition is satisfied by $\mathrm{Pol}(P_2)$ (for example, by $f=\max$).

Remark 5.6. Many, but not all the statements that we have shown also apply to *infinite* digraphs. Clearly, P_1 is still the greatest element in the respective poset. In Theorem 2.3, only the forward direction holds if G and H are infinite; however, in this text we only used (e.g., in Lemma 3.4) the forward direction of this theorem. In the proof that P_2 is the unique lower bound of P_1 we used the fact that every finite graph has a core, which is no longer true for infinite digraphs [16].

For the maximal lower bounds of P_2, the situation looks as follows. In the proof that digraphs G such that $\mathrm{Pol}(G)\not\models\Sigma_M$ pp construct T_3, we needed to work with an expansion of G by constants; expansions by constants are pp constructible in G if G is countably infinite and an ω-categorical model-complete core; see [10,14]. Every digraph with a Maltsev polymorphism is totally rectangular even if the digraph is infinite. The proof of Theorem 4.6 of Carvalho, Egri, Jackson, and Niven can be generalised to show that every infinite digraph which is totally rectangular is homomorphically equivalent to an infinite disjoint union of cycles or to one of the infinite paths $P^\infty:=(\mathbb{N},\{(u,u+1)\mid u\in\mathbb{N}\})$, $P_\infty:=(\mathbb{N},\{(u+1,u)\mid u\in\mathbb{N}\})$, the disjoint union $P_\infty+P^\infty$ of P_∞ and P^∞, and $P_\infty^\infty:=(\mathbb{Z},\{(u,u+1)\mid u\in\mathbb{Z}\})$. (All of these graphs have a Maltsev polymorphism.)

An infinite disjoint union of cycles C is not maximal below P_2: to see this, let k be the length of a shortest cycle in C. Observe that the pp power of C of dimension one given by the formula $\phi(x,y):=x\to y\wedge x\xrightarrow{k}x$ is homomorphically equivalent to C_k. If $k=1$, then C is homomorphically equivalent to C_1. If $k>1$, then $C\le C_k<P_2$. Since finite structures can only pp construct finite structures, we have that P_2 cannot pp construct the core digraphs P_∞, P^∞, $P_\infty+P^\infty$, and P_∞^∞. Conversely, these graphs can pp construct P_2 with the same formula $\phi(x_1,x_2,y_1,y_2):=E(y_1,x_1)\wedge E(x_2,y_2)\wedge x_1=x_2$. Clearly, P_∞ and P^∞ pp construct each other. We do not know whether these graphs are maximal lower bounds of P_2 in the class of all digraphs.

6. Concluding remarks

Primitive positive constructibility orders finite digraphs H by their 'strength' with respect to the H-coloring problem. Many deep combinatorial statements about graphs and digraphs can be phrased in terms of this order. We showed that at least the top region of the resulting poset can be described completely. A full description of the entire poset $\mathfrak{P}_{\text{Digraphs}}$ would be highly desirable.

We already mentioned that the pp constructibility poset on disjoint unions of cycles has been described in [11]; in particular, it contains no infinite ascending chains and is a lattice. Note that this result combined with the result of the present paper shows that for exploring $\mathfrak{P}_{\text{Digraphs}}$ it remains to explore the interval between K_3 and T_3: if a finite digraph H does not have a Maltsev polymorphism, then we proved that it is below T_3 (and above K_3); otherwise, it is homomorphically equivalent to a directed path or a disjoint union of cycles and hence falls into the region that has already been completely described.

We state three concrete open problems.

1. Is $\mathfrak{P}_{\text{Digraphs}}$ a lattice? (Primitive positive constructibility is known to form a meet semilattice on the class of all finite relational structures factored by pp interconstructibility, but it is not clear to the authors whether the clone product construction for the meet used there can be carried out in the category of digraphs.)
2. Does $\mathfrak{P}_{\text{Digraphs}}$ contain infinite ascending chains? (We have seen an infinite antichain in this article; an infinite descending chain of digraphs with a Maltsev polymorphism can be found in [11] and the existence of infinite descending chains of digraphs without a Maltsev polymorphism follows from results of [17], and also from results in [18].)
3. What are the maximal lower bounds of T_3 in $\mathfrak{P}_{\text{Digraphs}}$?

Acknowledgement. The authors would like to thank the anonymous referees for thoroughly reading our article as well as giving helpful comments that improved the final article.

References

[1] J. Nešetřil and C. Tardif: Duality theorems for finite structures, *Journal of Combinatorial Theory, Series B* **80** (2000), 80–97.

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

[3] T. Feder and M. Y. Vardi: The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory, *SIAM Journal on Computing* **28** (1999), 57–104.

[4] J. Bulin, D. Delic, M. Jackson and T. Niven: A finer reduction of constraint problems to digraphs, *Log. Methods Comput. Sci.* **11** (2015).

[5] A. A. Bulatov: A dichotomy theorem for nonuniform CSPs, in: *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17*, 319–330, 2017.

[6] D. Zhuk: A proof of the CSP dichotomy conjecture, *J. ACM* **67** (2020), 1–78.

[7] V. Dalmau: Linear datalog and bounded path duality of relational structures, *Logical Methods in Computer Science* **1** (2005).

[8] L. Egri, B. Larose and P. Tesson: Symmetric datalog and constraint satisfaction problems in logspace, in: *Proceedings of the Symposium on Logic in Computer Science (LICS)*, 193–202, 2007.

[9] A. Kazda: n-permutability and linear Datalog implies symmetric Datalog, *Log. Methods Comput. Sci.* **14** (2018).

[10] L. Barto, J. Opršal and M. Pinsker: The wonderland of reflections, *Israel Journal of Mathematics* **223** (2018), 363–398.

[11] M. Bodirsky, F. Starke and A. Vucaj: Smooth digraphs modulo primitive positive constructability and cyclic loop conditionsm *International Journal on Algebra and Computation* **31** (2021), 939–967.

[12] M. Bodirsky and A. Vucaj: Two-element structures modulo primitive positive constructability, *Algebra Universalis*, **81**, 2020.

[13] P. Hell and J. Nešetřil: *Graphs and Homomorphisms*, Oxford University Press, Oxford, 2004.

[14] M. Bodirsky: Cores of countably categorical structures, *Logical Methods in Computer Science (LMCS)* **3** (2007), 1–16.

[15] C. Carvalho, L. Egri, M. Jackson and T. Niven: On Maltsev digraphs, *Electr. J. Comb.* **22** (2015), P1.47.

[16] B. L. Bauslaugh: Cores and compactness of infinite directed graphs, *Journal of Combinatorial Theory, Series B* **68** (1996), 255–276.

[17] M. Bodirsky, A. Mottet, M. Olšák, J. Opršal, M. Pinsker and R. Willard: ω-categorical structures avoiding height 1 identities, *Transactions of the American Mathematical Society*, accepted.

[18] M. Jackson, T. Kowalski and T. Niven: Complexity and polymorphisms for digraph constraint problems under some basic constructions, *Int. J. Algebra Comput.* **26** (2016), 1395–1433.

Manuel Bodirsky, Florian Starke

Institute of Algebra

Technische Universität Dresden

01062 Dresden, Germany

{manuel.bodirsky,florian.starke}@tu-dresden.de