THE FINITE REPRESENTATION PROPERTY FOR COMPOSITION, INTERSECTION, ANTIDOMAIN AND RANGE

BRETT MCLEAN

ABSTRACT. We prove that the finite representation property holds for representation by partial functions for the signature consisting of composition, intersection, antidomain and range. It follows that the finite representation property holds for any expansion of this signature with the fixset, preferential union, maximum iterate and opposite operators. The proof shows that for all these signatures the size of base required is bounded by a double-exponential function of the size of the algebra. We also give an example of a signature for which the finite representation property fails to hold for representation by partial functions.

Keywords: Finite representation property, partial functions, intersection, antidomain, range.

1. INTRODUCTION

The investigation of the abstract algebraic properties of partial functions involves studying the isomorphism class of algebras whose elements are partial functions and whose operations are some specified set of operations on partial functions—operations such as composition or intersection for example. We refer to an algebra isomorphic to an algebra of partial functions as representable.

The primary aim is always to determine how simply the class of representable algebras can be axiomatised and to find such an axiomatisation. Often, the representation classes have turned out to be axiomatisable by finitely many equations or quasi-equations [6, 1, 3, 4, 2].

Another question to ask is whether every finite representable algebra can be represented by functions on some finite set. Interest in this finite representation property originates from its potential to help prove decidability of representability, which in turn can help give decidability of the equational or universal theories of the representation class.

Recently, Hirsch, Jackson and Mikulás established the finite representation property for many signatures, but were unable to find a proof for signatures containing the intersection, domain and range operators together [2].

In this paper we prove the finite representation property for half the outstanding signatures: those containing not just domain, but antidomain. These include a signature containing all the most commonly considered operations on partial functions. Our proof also allows us to obtain a double-exponential bound on the size of base set required for a representation. As an additional observation, we give an example showing that there are signatures for which the finite representation property does not hold for representation by partial functions.

2. Algebras of Partial Functions

In this section we give preliminary definitions and then make some elementary observations, which we rely on later, about the algebras we consider.

Given an algebra $\mathfrak{A}$, when we write $a \in \mathfrak{A}$ or say that a is an element of $\mathfrak{A}$, we mean that a is an element of the domain of $\mathfrak{A}$. We follow the convention that algebras are always nonempty.

Definition 2.1. Let σ be an algebraic signature whose symbols are a subset of $\{;, \wedge, A, R, 0, 1', D, F, \sqcup, {}^{\uparrow}, {}^{-1}\}$. An **algebra of partial functions** of the signature σ is an algebra of the signature σ whose elements are partial functions and with operations given by the set-theoretic operations on those partial functions described in the following.

Let X be the union of the domains and ranges of all the partial functions. We call X the **base**. In an algebra of partial functions

- the binary operation ; is composition of partial functions:
$$f \,;\, g = \{(x, z) \in X^2 \mid \exists y \in X : (x, y) \in f \text{ and } (y, z) \in g\},$$

- the binary operation $\wedge$ is intersection:
$$f \wedge g = \{(x, y) \in X^2 \mid (x, y) \in f \text{ and } (x, y) \in g\},$$

- the unary operation A is the operation of taking the diagonal of the **antidomain** of a function—those points of X where the function is not defined:
$$A(f) = \{(x, x) \in X^2 \mid \not\exists y \in X : (x, y) \in f\},$$

- the unary operation R is the operation of taking the diagonal of the range of a function:
$$R(f) = \{(y, y) \in X^2 \mid \exists x \in X : (x, y) \in f\},$$

- the constant 0 is the nowhere-defined function:
$$0 = \emptyset,$$

- the constant 1' is the identity function on X:
$$1' = \{(x, x) \in X^2\},$$

- the unary operation D is the operation of taking the diagonal of the domain of a function:
$$D(f) = \{(x, x) \in X^2 \mid \exists y \in X : (x, y) \in f\},$$

- the unary operation F is **fixset**, the operation of taking the diagonal of the fixed points of a function:
$$F(f) = \{(x, x) \in X^2 \mid (x, x) \in f\},$$

- the binary operation $\sqcup$ is **preferential union**:
$$(f \sqcup g)(x) = \begin{cases} f(x) & \text{if } f(x) \text{ defined} \\ g(x) & \text{if } f(x) \text{ undefined, but } g(x) \text{ defined} \\ \text{undefined} & \text{otherwise} \end{cases}$$

- the unary operation ${}^{\uparrow}$ is the **maximum iterate**:
$$f^{\uparrow} = \bigcup_{n \leq \omega} (f^n \,;\, A(f)),$$

- the unary operation $^{-1}$ is an operation we call **opposite**:

$$f^{-1} = \{(x,y) \in X^2 \mid (y,x) \in f \text{ and } ((z,x) \in f \implies z = y)\}.$$

The list of operations in Definition 2.1 does not exhaust those that have been considered for partial functions, but does include the most commonly appearing operations.

Definition 2.2. Let $\mathfrak{A}$ be an algebra of one of the signatures specified by Definition 2.1. A **representation of $\mathfrak{A}$ by partial functions** is an isomorphism from $\mathfrak{A}$ to an algebra of partial functions of the same signature. If $\mathfrak{A}$ has a representation then we say it is **representable**.

In [2], Hirsch, Jackson and Mikulás give a finite equational axiomatisation for the class representation class for the signature $(;, \wedge, A, R)$ and similarly for any expansion of this signature by operations in $\{0, 1', D, F, \sqcup\}$. For expanded signatures containing the maximum iterate operator they give axioms that, if we restrict attention to finite algebras, axiomatise the representable ones.

The operation that we call opposite is described in [5], where Menger calls the concrete operation 'bilateral inverse' and uses 'opposite' to refer to an abstract operation intended to model this bilateral inverse. The opposite operation appears again in Schweizer and Sklar's [7] and [8], but thereafter does not appear to have received any further attention. In particular, for signatures containing opposite, axiomatisations of the representation classes remain to be found.

Definition 2.3. Let σ be a signature. We say that σ has the **finite representation property** (for representation by partial functions) if whenever a finite algebra of the signature σ is representable by partial functions, it is representable on a finite base.[a]

In [2], Hirsch, Jackson and Mikulás establish the finite representation property for many signatures whose symbols are subsets of $\{;, \wedge, 0, 1', D, R, A, F, \sqcup, ^\uparrow\}$. Assuming composition is in the signature, they prove the finite representation property holds for any such signature not containing domain, any not containing range and almost all that do not contain intersection. This leaves one significant group of cases, which they leave as an open problem: signatures containing $\{;, \wedge, D, R\}$.

In this paper we prove that the signature $(;, \wedge, A, R)$ has the finite representation property and as a corollary that any expansion of $(;, \wedge, A, R)$ by operations that we have mentioned (including opposite) also has the finite representation property.

An arbitrary representable $(;, \wedge, A, R)$-algebra, $\mathfrak{A}$, is a meet-semilattice, with meet given by $\wedge$. Whenever we treat such an algebra as a poset, we are using the order induced by this semilattice. The algebra has a least element, 0, given by $A(a) ; a$ for any $a \in \mathfrak{A}$ and any representation of $\mathfrak{A}$ must represent 0 by the empty set. We can define $1' := A(0)$ and $D := A^2$ and in any representation these must be represented by the identity function and the domain operation respectively.

The down-set of any element $a \in \mathfrak{A}$ forms a Boolean algebra using the meet operation of $\mathfrak{A}$ and with complementation given by $\bar{b} := A(b) ; a$. In particular, the set $D(\mathfrak{A}) := \{D(a) \mid a \in \mathfrak{A}\}$ of **domain elements** forms such a Boolean algebra, being the down-set of 1'. Clearly $A(a)$ and $R(a)$ are domain elements for any $a \in \mathfrak{A}$.

---

[a]This property has also been called the *finite algebra on finite base property*.

Any representation θ of $\mathfrak{A}$ by partial functions restricts to representations of each $\downarrow a$ as a field of sets over $\theta(a)$. From this we see that θ turns any finite joins in $\mathfrak{A}$ into unions.

Definition 2.4. Let $\mathfrak{P}$ be a poset with a least element, 0. An **atom** of $\mathfrak{P}$ is a minimal nonzero element of $\mathfrak{P}$. We say that $\mathfrak{P}$ is **atomic** if every nonzero element is greater than or equal to an atom.

A *finite* representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra, $\mathfrak{A}$, is necessarily atomic and its domain elements form an atomic Boolean algebra. Any $a \in \mathfrak{A}$ can be expressed as a finite join of atoms of $\mathfrak{A}$ since, given $0 < a < b$, we can split b as $b = a \vee (\mathrm{A}(a) \,;\, b)$.

The purpose of the next example is to illustrate that, unlike Boolean algebras for example, the set of atoms in a finite representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra can be almost as large as the algebra itself.

Example 2.5. Let G be any finite group. We can make $G \cup \{0\}$ into an algebra of the signature $(;, \wedge, \mathrm{A}, \mathrm{R})$ by using the group operation for composition (and $g\,;\,0 = 0\,;\,g = 0$) and defining $g \wedge h = 0$ unless $g = h$, every antidomain of a nonzero element to be 0 (and $\mathrm{A}(0) = e$, the group identity) and every range of a nonzero element to be e (and $\mathrm{R}(0) = 0$). Then every nonzero element of $G \cup \{0\}$ is an atom. Augmenting the Cayley representation of G (the representation $\theta(g)(h) = hg$) by setting $\theta(0) = \emptyset$ demonstrates that $G \cup \{0\}$ is representable.

3. The Finite Representation Property

We now proceed towards our main result via a series of lemmas.

Lemma 3.1. *Let $\mathfrak{A}$ be a finite representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra and a and b be atoms of $\mathfrak{A}$. Then*

 (i) $\mathrm{D}(a)$ *is an atom,*
 (ii) $\mathrm{R}(a)$ *is an atom,*
(iii) $a\,;\,b$ *is either 0 or an atom.*

Proof.   (i) Suppose $\alpha \le \mathrm{D}(a)$ (so α is a domain element). Then $\alpha;a \le \mathrm{D}(a);a = a$. Since a is an atom, either $\alpha\,;\,a = 0$ or $\alpha\,;\,a = a$. By thinking of the elements as partial functions, it is clear that $\alpha = 0$ or $\alpha = \mathrm{D}(a)$ respectively. Hence $\mathrm{D}(a)$ is an atom.
  (ii) Completely analogous to (i).
 (iii) Let a be an atom and suppose that $a\,;\,b$ is nonzero. Let $c \in \mathfrak{A}$ and suppose $c \le a\,;\,b$. Then $\mathrm{D}(c) \le \mathrm{D}(a\,;\,b) \le \mathrm{D}(a)$. Hence if $\mathrm{D}(c)\,;\,a = 0$ then $c = 0$. If $\mathrm{D}(c);a > 0$, then we must have $\mathrm{D}(c);a = a$ and hence $\mathrm{D}(c) = \mathrm{D}(a;b) = \mathrm{D}(a)$. Therefore $c = a\,;\,b$. So $a\,;\,b$ is an atom. $\qquad\square$

The next lemma says that we only need worry about the atoms.

Lemma 3.2. *Every representation, θ, of a finite representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra is 'the same as' a labelled directed graph (with loops allowed) having the following properties.*

 (i) *Every edge is labelled by a single atom.*
 (ii) *Every atom appears as a label.*
(iii) *There is at most one edge with any given label starting from any given vertex.*

(iv) *For all atoms a and b and vertices x and z, there is an edge from x to z labelled with $a \,;\, b$ if and only if there exists a vertex y and edges from x to y labelled by a and from y to z labelled by b.*
(v) *Every vertex has a loop.*
(vi) *For every atom a and vertex x, there exists a vertex y and an edge from x to y if and only if the loop at x is labelled $\mathrm{D}(x)$.*
(vii) *For every atom a and vertex y, there exists a vertex x and an edge from x to y if and only if the loop at y is labelled $\mathrm{R}(y)$.*

Formally, we mean the following. Every graph with these properties determines a representation, given by $(x, y) \in \theta(a)$ if and only if there is an edge from x to y labelled with one of the atoms below a. Further, this operation of forming a representation from a graph is inverse to the operation of (in the obvious way) forming a graph from a representation.

Proof. Routine. □

Henceforth, we will identify representations with their atom-labelled graphs, so we will refer to vertices, edges and edge-labels of representations.

Definition 3.3. Let θ be a representation of a finite $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra. For a domain element α we will call a vertex x of θ an α-**vertex** if the loop at x is labelled α. We extend this definition to sets of domain elements in the obvious way.

Given a representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra, for α and β both domain elements and atoms, we will write $\alpha \lesssim \beta$ if there is an atom a with $\mathrm{D}(a) = \alpha$ and $\mathrm{R}(a) = \beta$. Then $\lesssim$ is easily seen to be a preorder.

For the next lemma we need the concept of a groupoid. We use the category-theoretic rather than the algebraic definition of a groupoid.

Lemma 3.4. *Let $\mathfrak{A}$ be a finite representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra. Let V be a $\lesssim$-equivalence class and E the set of all atoms, a, with both $\mathrm{D}(a)$ and $\mathrm{R}(a)$ in V. Then $(V, E, ;)$ is a groupoid.*

Proof. The statement should be read as saying that V is the set of objects of the groupoid and E the arrows, with $a : \mathrm{D}(a) \to \mathrm{R}(a)$ for each $a \in E$. Composition of arrows is given by $;$ and the identities given by $\mathrm{id}_\alpha := \alpha$.

We know that composition is associative. We have $\mathrm{D}(a) = \alpha \implies \alpha \,;\, a = a$ and $\mathrm{R}(a) = \alpha \implies a \,;\, \alpha = a$, so the id_α's really are identities. Hence we only need to show that every arrow has an inverse.

Let $a \in E$. Then $\mathrm{D}(a), \mathrm{R}(a) \in V$ and so $\mathrm{R}(a) \lesssim \mathrm{D}(a)$, meaning there is some b (necessarily in E) with $b : \mathrm{R}(a) \to \mathrm{D}(a)$.

Think of any representation of $\mathfrak{A}$. Some edge, starting at x say, will be labelled $a \,;\, b$ and all edges labelled $a \,;\, b$ will start and end at $\mathrm{D}(a)$-vertices. We can follow a path, starting at x, of $a \,;\, b$ edges. As $\mathfrak{A}$ is finite and all vertices on the path are reachable via an edge from x, the path will eventually revisit some vertex. Hence for some positive n there is an $(a \,;\, b)^n$ loop at a $\mathrm{D}(a)$-vertex. We conclude that $(a \,;\, b)^n = \mathrm{D}(a)$ and so a has inverse $b \,;\, (a \,;\, b)^{n-1}$. □

Given $\mathfrak{A}$, V and E as in Lemma 3.4, we know that in any representation of $\mathfrak{A}$, the subgraph induced by V-vertices will consist of disjoint copies of one particular labelled graph, G. Pick some $\alpha \in V$, then G is the complete graph with vertex set

those elements of E with domain α and for any pair (a, b) of vertices, the edge from a to b labelled by $a^{-1} \, ; b$. We call a copy of this graph a **representation of V**.

Definition 3.5. Let θ be a representation of an algebra of the signature $(;, \wedge, \mathrm{A}, \mathrm{R})$. Let W be a subset of the vertices of θ. We define the **forward closure** of W to be the subgraph induced by the vertices reachable via an edge *starting* in W.

Note that the forward closure really is a closure operator. Indeed, if there is an edge from $x \in W$ to y labelled a and from y to z labelled b then z is reachable via an edge starting in W, labelled $a \, ; b$. Each $x \in W$ is reachable from W via the loop at x.

Lemma 3.6. *Let $\mathfrak{A}$ be a finite $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra, θ and ϕ be representations of $\mathfrak{A}$ and V be a $\lesssim$-equivalence class of $\mathfrak{A}$. Let R and S be representations of V that are subgraphs of θ and ϕ respectively. Then the forward closures of R and S are isomorphic.*

Proof. Choose some $\alpha \in V$ and α-vertices x and y from R and S respectively. The labels of the edges starting from x are precisely the atoms whose domain is α (each atom appearing once). Hence there is a bijection between these atoms and the vertices reachable via an edge starting from x. Similarly for y. This gives a bijection between the two sets of vertices. The bijection is an isomorphism since there is an edge labelled b from the vertex reachable via a to the vertex reachable via c exactly when $a \, ; b = c$. □

For a $\lesssim$-equivalence class V, we call a graph isomorphic to the forward closure of a representation of V a **forward representation of V**.

Theorem 3.7. *The finite representation property holds for representation by partial functions for the signature $(;, \wedge, \mathrm{A}, \mathrm{R})$.*

Proof. Let $\mathfrak{A}$ be a finite representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra. From $\lesssim$, form the partial order of $\lesssim$-equivalence classes. The **depth** of a $\lesssim$-equivalence class V will be the length of the *longest* increasing chain in this partial order starting at V. (We take the length of a chain to be one less than the number of elements it contains, so a maximal $\lesssim$-equivalence class has depth zero.) The depth of an element of $\mathfrak{A}$ will be the depth of its $\lesssim$-equivalence class.

We construct a finite graph, G, step by step, by adding representations of $\lesssim$-equivalence classes of increasing depths and then argue that the graph is a representation of $\mathfrak{A}$. The idea of the proof is to 'add everything we can, $|\mathfrak{A}|$ times'.

Assume inductively that we have carried out steps $0, \ldots, n-1$ of our construction, giving us the graph G_{n-1}. We form G_n as follows. (For the base case of this induction, we let G_{-1} be the empty graph.) Let V be a $\lesssim$-equivalence class of depth n, and R a representation of V. A choice of edges from R to G_{n-1} labelled by atoms is **allowable** if adding R and these edges to G_{n-1} would make the forward closure of R isomorphic to the forward representation of V. For every allowable choice, we add to G_{n-1}: $|\mathfrak{A}|$ copies of R together with $|\mathfrak{A}|$ copies of the specified edges. The graph G_n is the graph we have once we have done this for every $\lesssim$-equivalence class.

Note that the order that $\lesssim$-equivalence classes of a given depth, n, are processed is immaterial since no allowable choice could have an edge ending at a vertex that had not been in G_{n-1}. By induction, each G_n is finite: assume that G_{n-1} is finite;

then as each 'R' is also finite and the set of atoms is finite we see that the number of allowable choices is finite, so G_n is finite. We take G to be G_N, where N is the maximum depth of any $\lesssim$-equivalence class of $\mathfrak{A}$.

We want to show that G satisfies the conditions of Lemma 3.2 and is therefore a representation of $\mathfrak{A}$. Condition (i) holds because each edge receives a label when and only when it is added to the graph. Conditions (iii)-(vi) hold because when any vertex is added, its forward closure is the forward representation of a $\lesssim$-equivalence class.

Condition (ii) will hold if we can show that for every $\lesssim$-equivalence class V, a nonzero number of representations of V are added at the appropriate stage of the construction. This is true for any V of depth 0 since G_0 consists of $|\mathfrak{A}|$ representations of each $\lesssim$-equivalence class of depth 0. For V of positive depth, n, this will follow, by induction on depth, if condition (vii) holds, because for some $\lesssim$-equivalence class V' of depth $n-1$ there will be V'-vertices in G that we know are hit by edges starting at V-vertices.

One direction of condition (vii) is clear: if G has an edge labelled a from x to y then it will have a loop labelled $R(a)$ at y. For the other direction, let y be an $R(a)$-vertex of G and F be a copy of the forward representation of the $\lesssim$-equivalence class of $D(a)$. Pick any $D(a)$-vertex, x', of F. Then F has an edge labelled a starting at x' and ending at y' say. We will show that it is possible to embed F into G in such a way that y' is mapped to y, which shows that there exists an a-labelled edge in G ending at y.

Since y' and y are both $R(a)$-vertices, there is an isomorphism, g, from the forward closure of y' to the forward closure of y, that maps y' to y. Let n be the depth of $D(a)$. We will argue by induction that for any $m \leq n$ there exists a map f_m such that

(i) f_m embeds the subgraph of F induced by vertices with loop-labels of depth at most m into G_m,

(ii) f_m and g agree where they are both defined.

As before the base case for the induction is depth -1, so we define f_{-1} to be the empty map. Of course the map f_n is the required embedding of F into G.

Suppose we have f_{m-1} as above. We can form f_m, an extension of f_{m-1}, as follows. First use g to extend f_{m-1} to those vertices in the forward closure of y' with loop-labels of depth m (and extend to the edges starting at these vertices). The remaining vertices we need to extend to are partitioned into representations of various $\lesssim$-equivalence classes of depth m. Each of these representations, R, together with f_{m-1} specifies an allowable choice on G_{m-1}, providing $|\mathfrak{A}|$ possible ways to extend f_{m-1} to R and to the edges starting from R. The number of vertices in F is the number of atoms with domain $D(a)$. So F certainly contains no more than $|\mathfrak{A}|$ representations of $\lesssim$-equivalence classes of depth m (including those in the forward closure of y'). Hence there exists a way to extend f_{m-1} to all these representations simultaneously. □

Corollary 3.8. *The finite representation property holds for representation by partial functions for any signature whose symbols contain $\{;, \wedge, A, R\}$ and are a subset of $\{;, \wedge, A, R, 0, 1', D, F, \sqcup, {}^{\uparrow}\}$.*

Proof. We know that 0, 1' and D are definable over $(;, \wedge, A, R)$. Hence, given an algebra $\mathfrak{A}$ over an expanded signature, a finite representation of $\mathfrak{A}$ is provided by

a finite representation of the $(;, \wedge, \mathrm{A}, \mathrm{R})$-reduct of $\mathfrak{A}$. Similarly for F, which is definable as $\mathrm{F}(a) := a \wedge \mathrm{D}(a)$.

We also know that representations of $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebras turn any finite joins into unions. When preferential union is present, the equation $a \sqcup b = a \vee (\mathrm{A}(a) \,;\, b)$ is valid in representable algebras, so this suffices to give the result for $\sqcup$. To see the result for maximum iterate, note that in a *finite* representable algebra, each $a^\uparrow$ must be the finite join $\bigvee_{n \leq |\mathfrak{A}|}(a^n \,;\, \mathrm{A}(a))$. $\qquad\qquad\square$

With a little more work we can expand the list of operators to include opposite.

Theorem 3.9. *The finite representation property holds for representation by partial functions for any signature whose symbols contain $\{;, \wedge, \mathrm{A}, \mathrm{R}\}$ and are a subset of $\{;, \wedge, \mathrm{A}, \mathrm{R}, 0, 1', \mathrm{D}, \mathrm{F}, \sqcup, {}^\uparrow, {}^{-1}\}$.*

Proof. Let $\mathfrak{A}$ be a finite representable algebra of one of the specified signatures. We will argue that the representation described in Theorem 3.7, of the $(;, \wedge, \mathrm{A}, \mathrm{R})$-reduct of $\mathfrak{A}$, represents opposite correctly.

First note that if $a \in \mathfrak{A}$ is an atom with domain and range in the same $\lesssim$-equivalence class, V, then the opposite of a in $\mathfrak{A}$ and the inverse of a in the groupoid V coincide, so there is no risk of confusion when we write a^{-1}. Now for an arbitrary $a \in \mathfrak{A}$ let $\varphi_a(b)$ be the assertion that

(i) b is an atom, with $b \leq a$,
(ii) $\mathrm{D}(b)$ and $\mathrm{R}(b)$ are in the same $\lesssim$-equivalence class,
(iii) for any atom b' with $b' \leq a$, if $\mathrm{R}(b) = \mathrm{R}(b')$ then $b = b'$.

We claim that

(1) $$a^{-1} = \bigvee\{b^{-1} \mid \varphi_a(b)\}.$$

Suppose that c is an atom, with $c \leq a^{-1}$. Then the domain and range of c must be in the same $\lesssim$-equivalence class, so that c^{-1} is an inverse in a groupoid and hence $(c^{-1})^{-1} = c$. In any representation, if $(x, y) \in c^{-1}$ then $(y, x) \in c$ so $(y, x) \in a^{-1}$, which implies $(x, y) \in a$. So $c^{-1} \leq a$. If b' is an atom, with $b' \leq a$ and $\mathrm{R}(c^{-1}) = \mathrm{R}(b')$ then y is an $\mathrm{R}(b')$-vertex, so there is a b'-labelled edge ending at y. Since $(y, x) \in a^{-1}$, this b'-labelled edge must be the edge starting at x, which we know is labelled with c^{-1}. As c^{-1} and b' are atoms they must therefore be equal. So $\varphi_a(c^{-1})$ holds and therefore $c = (c^{-1})^{-1}$ is one of the atoms in the join $\bigvee\{b^{-1} \mid \varphi_a(b)\}$. Hence we have $\leq$ in Equation (1).

For the reverse inequality, consider any representation and let b be such that $\varphi_a(b)$ holds. If $(x, y) \in b^{-1}$ then $(y, x) \in b$ and so $(y, x) \in a$. There is no other edge ending at x and labelled with b since b is an inverse in a groupoid and there is no edge labelled by any other atom below a since b does not share its range with any other atom below a. Hence (y, x) is the unique edge ending at x labelled by a and so $(x, y) \in a^{-1}$.

It easy to verify that for atoms with domain and range in the same $\lesssim$-equivalence class, the opposite is represented correctly in any representation of the $(;, \wedge, \mathrm{A}, \mathrm{R})$-reduct. Let θ be the representation of the $(;, \wedge, \mathrm{A}, \mathrm{R})$-reduct as constructed in

Theorem 3.7 and a be an arbitrary element of $\mathfrak{A}$. Using Equation (1) we see that

$$
\begin{aligned}
& (x, y) \in \theta(a^{-1}) \\
\implies & (x, y) \in \theta(b^{-1}) && \text{for some } b \text{ such that } \varphi_a(b) \\
\implies & (x, y) \in \theta(b)^{-1} && \text{for some } b \text{ such that } \varphi_a(b) \\
\implies & (y, x) \in \theta(b) && \text{for some } b \text{ such that } \varphi_a(b) \\
\implies & (y, x) \in \theta(a) && \text{and } y \text{ is unique in this regard} \\
\implies & (x, y) \in \theta(a)^{-1}
\end{aligned}
$$

The first implication holds because the join $\bigvee \{ b^{-1} \mid \varphi_a(b) \}$ is a finite join and we know finite joins are turned into unions. In the fourth implication, we conclude the uniqueness of y because of the condition $\varphi_a(b)$ on b.

Conversely, if $(x, y) \in \theta(a)^{-1}$ then $(y, x) \in \theta(a)$, so $(y, x) \in \theta(b)$ for some atom $b \leq a$. Further, y is the unique vertex having an edge to x labelled by an atom below a. This implies that b satisfies condition (iii) of $\varphi_a(b)$. If $\mathrm{D}(b)$ and $\mathrm{R}(b)$ were in different $\lesssim$-equivalence classes then, in the terminology of Theorem 3.7, the depth of $\mathrm{D}(b)$ would be greater than the depth of $\mathrm{R}(b)$. Then because of the way θ is constructed, if there is one b-labelled edge to x, there are at least $|\mathfrak{A}|$ b-labelled edges to x. Knowing that the trivial algebra has a finite representation, we may assume $|\mathfrak{A}| > 1$ and hence b must satisfy condition (ii) of $\varphi_a(b)$. So we have

$$
\begin{aligned}
& (y, x) \in \theta(b) && \text{for some } b \text{ such that } \varphi_a(b) \\
\implies & (x, y) \in \theta(b)^{-1} && \text{for some } b \text{ such that } \varphi_a(b) \\
\implies & (x, y) \in \theta(b^{-1}) && \text{for some } b \text{ such that } \varphi_a(b) \\
\implies & (x, y) \in \theta(a^{-1})
\end{aligned}
$$

completing the proof that $^{-1}$ is represented correctly by θ. $\qquad\square$

It is clear that from the proof of Theorem 3.7 we can extract a bound on the size required for the base.

Proposition 3.10. *Every finite representable $(;, \wedge, \mathrm{A}, \mathrm{R})$-algebra $\mathfrak{A}$ is representable over a base of size*

$$
|\mathfrak{A}|^{|\mathfrak{A}|^{|\mathfrak{A}|}}.
$$

Proof. If $|\mathfrak{A}| = 1$ then $\mathfrak{A}$ is representable over an empty base. So we may assume $|\mathfrak{A}| \geq 2$. Let G and $(G_n)_{n \geq -1}$ be as in the proof of Theorem 3.7. Let V be a $\lesssim$-equivalence class of depth n and let R be a representation of V. An allowable choice from R to G_{n-1} is determined by the labelled edges from a single vertex of R. There are at most $|\mathfrak{A}|$ labels, so at most $|G_{n-1}|^{|\mathfrak{A}|}$ allowable choices (unless V is of depth 0, in which case there is a single allowable choice). When G_n is constructed from G_{n-1}, for each allowable choice, $|\mathfrak{A}|$ copies of R are added, so $|\mathfrak{A}||R|$ vertices are added. The sum, over all $\lesssim$-equivalence classes of depth n, of the number of vertices in a representation of each class, is at most $|\mathfrak{A}|$. Hence at most $|\mathfrak{A}|^2 |G_{n-1}|^{|\mathfrak{A}|}$ vertices are added when G_n is constructed from G_{n-1}. We obtain

$$
|G_0| \leq |\mathfrak{A}|^2,
$$
$$
|G_n| \leq |G_{n-1}| + |\mathfrak{A}|^2 |G_{n-1}|^{|\mathfrak{A}|} \text{ for } n \geq 1,
$$

from which we may prove by induction that

$$|G_n| \leq |\mathfrak{A}|^{2(n+1)|\mathfrak{A}|^n}.$$

There are no more than $|\mathfrak{A}| - 1$ atoms in $\mathfrak{A}$, so the construction of G is completed by depth $|\mathfrak{A}| - 2$. Hence

$$|G| \leq |\mathfrak{A}|^{2(|\mathfrak{A}|-1)|\mathfrak{A}|^{|\mathfrak{A}|-2}} \leq |\mathfrak{A}|^{|\mathfrak{A}|^{|\mathfrak{A}|}}. \qquad \square$$

For comparison, note that in [2], whenever a signature is shown to have the finite representation property, a bound on the size required for the base is derived that has polynomial or exponential asymptotic growth.

As we alluded to in the introduction, finding a bound on the size required for a representation is sometimes used to show that representability of finite algebras is decidable. However, for the signatures that we have considered, decidability has already been established because a finite equational axiomatisation for the representable $(;, \wedge, A, R)$-algebras is known.

Finally, one might reasonably wonder if it is possible for the finite representation property *not* to hold for algebras of partial functions. After all, for every signature for which it has been settled, the finite representation property has been shown to hold. We finish with a simple example showing that we can indeed force a finite representable algebra of partial factions to fail to have representations over finite bases.

Example 3.11. Let U be the unary operation on partial functions given by

$$U(f) = \{(y, y) \in X^2 \mid \exists! x \in X : (x, y) \in f\}.$$

Let $\mathfrak{F}$ be the algebra of partial functions, of the signature $(;, \wedge, A, R, U)$ and with base $\omega \times 2$, containing the following elements.

- 0, the empty function and 1', the identity function on $\omega \times 2$,
- d, the identity function on $\omega \times \{0\}$ and r, the identity function on $\omega \times \{1\}$,
- f, the function with domain d and range r sending each $(n, 0)$ to $(n, 1)$,
- g, a function with domain d and range r. Each $(n, 1) \in \omega \times \{1\}$ has precisely two g-preimages: the least two elements of $\omega \times \{0\}$ that are neither the f-preimage $(n, 0)$ nor g-preimages of $(m, 1)$ for $m < n$. See Figure 1.
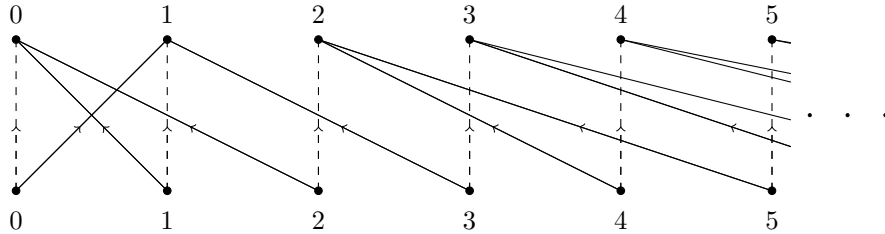

FIGURE 1. The algebra $\mathfrak{F}$. Dashed lines for f, solid lines for g.

Since $\mathfrak{F}$ *is* an algebra of partial functions, it is certainly representable by partial functions. It is easy to see that $\mathfrak{F}$ cannot be represented over a finite base. Indeed, $R(f) = U(f)$ so in any representation f is a bijection from its domain, the d-vertices, to its range, the r-vertices. On the other hand, $R(g) \neq U(g)$ so g maps

the d-vertices onto the r-vertices, but not injectively. Hence these sets of vertices cannot have finite cardinality.

By including the operator U in less expressive signatures, it is possible to give slightly simpler examples than Example 3.11. However we chose an expansion of the signature $(;, \wedge, A, R)$ in order to contrast with the other expansions that are the subject of this paper, for which we have seen that the finite representation property does hold.

References

1. Wiesław A. Dudek and Valentin S. Trokhimenko, *Functional Menger $\mathcal{P}$-algebras*, Comm. Algebra **30** (2002), no. 12, 5921–5931.
2. Robin Hirsch, Marcel Jackson, and Szabolcs Mikulás, *The algebra of functions with antidomain and range*, preprint: arXiv:1410.3981 (2014).
3. Marcel Jackson and Tim Stokes, *Partial maps with domain and range: extending Schein's representation*, Comm. Algebra **37** (2009), no. 8, 2845–2870.
4. Marcel Jackson and Tim Stokes, *Modal restriction semigroups: towards an algebra of functions*, Internat. J. Algebra Comput. **21** (2011), no. 7, 1053–1095.
5. Karl Menger, *An axiomatic theory of functions and fluents*, The Axiomatic Method (A. Heyting L. E. J. Brouwer, E. W. Beth, ed.), North Holland, 1959, pp. 454–472.
6. Boris M. Schein, *Restrictively multiplicative algebras of transformations*, Izv. Vyssh. Uchebn. Zaved. Mat. **95** (1970), no. 4, 91–102 (Russian).
7. B. Schweizer and A. Sklar, *The algebra of functions*, Math. Ann. **139** (1960), no. 5, 366–382.
8. ———, *The algebra of functions. III*, Math. Ann. **161** (1965), no. 3, 171–196.

Department of Computer Science, University College London, Gower Street, London WC1E 6BT

E-mail address: b.mclean@cs.ucl.ac.uk