Axiomatizing complex algebras by games

Ian Hodkinson[*] Szabolcs Mikulás[†] Yde Venema[‡]

January 10, 2000

Abstract

Given a variety V, we provide an axiomatization $\Phi(\mathsf{V})$ of the class SCmV of *complex algebras* of algebras in V. $\Phi(\mathsf{V})$ can be obtained effectively from the axiomatization of V; in fact, if this axiomatization is recursively enumerable, then $\Phi(\mathsf{V})$ is recursive.

1991 *Mathematics Subject Classification.* 06E25, 03G15, 03C25.
Key words and phrases. Boolean algebra with operators, complex algebra, axiomatization, complex group, variety, games.

[*]Department of Computing, Imperial College, 180 Queen's Gate, London SW7 2BZ. E-mail: `imh@doc.ic.ac.uk`. Partially supported by UK EPSRC No. GR/L85978.

[†]Department of Computer Science, King's College London, Strand, London WC2R 2LS. E-mail: `szabolcs@dcs.kcl.ac.uk`. Supported by EPSRC No. GR/L82441.

[‡]Institute of Logic, Language and Computation, University of Amsterdam, Plantage Muidergracht 24, 1081 TV Amsterdam. E-mail: `yde@wins.uva.nl`. The research of the third author has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

1 Introduction

The construction of complexes of structures is a standard procedure in mathematics. Probably the oldest and best known example is found in group theory: given a group, consider the algebra whose carrier is the power set of the group elements and whose operations are the power *lifts* of the group operations, for instance,

$$X \circ Y \;=\; \{x \circ y \mid x \in X, y \in Y\}.$$

In lattice theory it is well known that the set of ideals of a distributive lattice $\mathcal{L}$ again forms a lattice, of which the meet and join coincide with the lifted meet and join operations of $\mathcal{L}$, respectively:

$$\begin{aligned} I_1 \vee I_2 &= \{a_1 \vee a_2 \mid a_1 \in I_1, a_2 \in I_2\}, \\ I_1 \wedge I_2 &= \{a_1 \wedge a_2 \mid a_1 \in I_1, a_2 \in I_2\}. \end{aligned}$$

And as a last example we mention formal language theory, where we may see the product of two languages as the lift of word concatenation:

$$L_1; L_2 \;=\; \{w_1 w_2 \mid w_1 \in L_1, w_2 \in L_2\}.$$

Obviously, this construction can be carried out for an arbitrary operation, giving rise to the *power algebra* of an algebra (formal definitions are found in the next section). Since the universe of such an algebra is a power *set* algebra, it is natural to include the Boolean operations into the similarity type; thus we obtain the *full complex algebra* $\mathcal{L}^+$ of an algebra $\mathcal{L}$. If instead of *all* subsets of $\mathcal{L}$ we take as carrier of the algebra some non-empty collection of subsets of $\mathcal{L}$ that is closed under the Boolean operations and under the lifted operations, we get an arbitrary *complex algebra* over $\mathcal{L}$; formulated more concisely, a complex algebra over $\mathcal{L}$ is any subalgebra of $\mathcal{L}^+$.

For some notation, given a class K of algebras, we denote the class of full complex algebras over (algebras in) K by CmK; SCmK denotes the class of isomorphism types of complex algebras over K, and $\mathsf{Var}_{\mathsf{K}}$, the variety generated by CmK.

The construction gives rise to various questions of a universal algebraic nature, for instance concerning the relation between a class K of algebras and the class SCmK of associated complex algebras. For a survey of known results and references to the literature we refer the reader to Brink [1] and Goldblatt [4] (the second paper takes a more general perspective, considering complex algebras of arbitrary *relational* structures).

In this paper, we are interested in finding an axiomatization of the class of complex algebras of a given variety V. It seems that in the general case, not much is known. There are some known results relating the validity of an equation in an algebra to its validity in the power algebra. For instance, a result by Gautam [3] states that the validity of an equation is preserved under moving to the power algebra if and only if every variable in the equation occurs exactly once on each side of the equation. This makes it improbable that an equational axiomatization of a variety V will be of *direct* use in finding an axiomatization of SCmV.

Recently, Goranko and Vakarelov [5] have given complete axiomatizations of the modal logic of various classes of relational structures, including varieties of algebras. Translated into algebraic terms, their result yields a derivation system for the set $Equ(\mathsf{CmV})$ of equations valid in the class CmV for an arbitrary variety V. Their result crucially involves the extension of the lifted algebraic language with a so-called difference operator, and an extension of the derivation system with an non-structural derivation rule. However, for some varieties V, including groups and (thus) Boolean algebras, this difference operator is term-definable over the class CmV. Hence, for such a variety V, the result of Goranko and Vakarelov provides a derivation system for the equational theory $Equ(\mathsf{CmV})$ *within* the language of the complex algebras — but since this system has a non-structural rule, it is not an equational axiomatization in the traditional sense, or an equational characterization of the variety $\mathsf{Var}_{\mathsf{V}}$. Independently, Venema [17] obtained the same result for the case of groups.

In the case of groups, some other results are known. Complex algebras of groups appear in the literature on algebraic logic as *group relation algebras*, GRAs. Tarski [15] showed that GRA is axiomatizable by a set of equations over the class of integral relation algebras, while McKenzie [13] proved that no finite axiomatization of GRA can be found. McKenzie [13, p.282] writes:

> "It would certainly be of interest to have a reasonably elegant system of first-order axioms characterizing [GRA]."

The aim of this paper is to give such a characterization, not just for group relation algebras but in general for the class of complex algebras of any (recursively axiomatizable) variety of algebras. We will use two-player games in the characterization, and translate the existence of a winning strategy for one of the players into a set of first-order axioms; thus, we find, for an arbitrary class of the form SCmV, an axiomatization with strong intuitive content. Similar techniques were used to construct axiomatizations in [16, 6, 8, 7, 14]. The method is implicitly used in the much earlier [12], although games are not mentioned per se.

Formulated precisely, in this paper we will prove the following Theorem.

Theorem 1.1 *Let V be a variety of Σ-algebras, where Σ is a finite functional similarity type. There is a set $\Phi(\mathsf{V})$ of universal first-order sentences in the language of complex algebras over V such that whenever $\mathcal{A}$ is a Boolean algebra with Σ-operators, $\mathcal{A} \models \Phi(\mathsf{V})$ if and only if $\mathcal{A}$ is representable as a complex algebra over V. The set $\Phi(\mathsf{V})$ is effectively constructible from an equational axiomatization of V.*

There is no special reason to restrict ourselves to either a finite similarity type or to complex algebras over a variety. Similar techniques serve to axiomatize the class of complex algebras over any universally axiomatized class of relational structures and indeed, over any elementary class (for instance, by using Skolem functions to reduce to universal case). This covers representable relation algebras, and representable cylindric algebras of finite or countable dimension. In a more model-theoretic vein, any pseudo-elementary class of structures (see, e.g., [2]) that is closed under substructures can also be universally axiomatized by games [9].

We have mentioned the universal form of the axiomatization explicitly because of the following. Suppose that we are (also) interested in an equational axiomatization of the *variety* $\mathsf{Var}_{\mathsf{K}}$. Now if we have a *discriminator term* at our disposal for the class CmK (which is the case for e.g., group relation algebras), then the universal axiomatization $\Phi(\mathsf{K})$ can be effectively converted into an equational axiomatization for the variety $\mathsf{Var}_{\mathsf{K}}$. This can be seen as follows. Let $c(x)$ be a unary discriminator term over the class CmK. It is well-known (cf. Jipsen [11]) that there is a set of equations D_c such that (i) the variety V_c of Boolean algebras with operators defined by D_c is generated by the algebras for which c is a discriminator term, and (ii) c is a unary discriminator term in all subdirectly irreducible members of V_c. It is equally well-known that given a unary discriminator term c, there is an effective translation

$(\cdot)^c$ mapping universal formulas to equations such that φ^c is equivalent to φ in every algebra for which c is a discriminator term. From this it is straightforward to show that $\mathrm{Var}(\mathsf{K})$ is axiomatized by the set of equations $D_c \cup \{\varphi^c \mid \varphi \in \Phi(\mathsf{K})\}$, together with the set of equations axiomatizing Boolean algebras with operators of this similarity type. (Also, in such a case it follows that Var_K is identical to the class SPCmK.)

The paper is organized as follows. In the next section, we recall the basic definitions. In section 3, we introduce a two-player game, and in section 4, a game characterization is given for representability as a complex algebra. In the last section, we turn this into a first-order axiomatization.

Acknowledgements We would like to thank Peter Jipsen for stimulating discussions.

2 Preliminaries

A *similarity type* is a set of function symbols, each of which comes with a non-negative *arity*; we denote the arity of a symbol ∇ as $ar(\nabla)$. Throughout the paper, we will abbreviate $q = ar(\nabla)$ in order to ease some notational burden. Given a similarity type Σ, a Σ-*algebra* is a pair $\mathcal{L} = (L, I)$, where L is some non-empty set and I is a function interpreting each function symbol ∇ in Σ as an operation $I_\nabla : {}^qL \to L$.

As the similarity type of Boolean algebras we take the set $BA = \{+, -, 0\}$ where '+' denotes the join operation (union $\cup$ in fields of sets), '$-$' denotes complementation, and '0' represents the least element (the empty set $\varnothing$ in set algebras). The other function symbols such as $\cdot$ and 1 are taken as abbreviations. Operations interpreting the Boolean function symbols are denoted by the function symbols themselves (for instance, we do not write I_+, but rather $+$ or $\cup$). We assume that the reader is familiar with Boolean algebras and various notions pertaining to them, such as ultrafilters; see, e.g., [2].

Given a similarity type Σ, let Σ_{BA} be the similarity type consisting of the disjoint union of Σ and the Boolean function symbols. (So if Σ should already contain the Boolean symbols, we add new copies of them; these are not identified with the ones already in Σ; this avoids confusion in the case we are dealing with complex algebras of Boolean algebras.) In order to distinguish Σ_{BA}-algebras from Σ-algebras, we will usually denote the interpretation function of a Σ_{BA}-algebra by a diamond symbol; for example, in an abstractly given Σ_{BA}-algebra we denote by $\Diamond_\nabla$ the operation interpreting the symbol ∇.

We define the notion of a complex algebra over a Σ-algebra $\mathcal{L} = (L, I)$ as follows. For any function symbol ∇, the operation $I_\nabla^\Diamond$ is defined as the *lift* of the operation I_∇. That is, for subsets $X_1, \ldots, X_q$ of L we define

$$I_\nabla^\Diamond(X_1, \ldots, X_q) = \{I_\nabla(x_1, \ldots, x_q) \mid x_i \in X_i \text{ for all } i\}.$$

Now the *power algebra* of $\mathcal{L}$ is defined as the algebra

$$(\mathcal{P}(L), I_\nabla^\Diamond)_{\nabla \in \Sigma},$$

where $\mathcal{P}(-)$ denotes power set, whereas the *full complex algebra* $\mathcal{L}^+$ of $\mathcal{L}$ is given as

$$\mathcal{L}^+ = (\mathcal{P}(L), \cup, -, \varnothing, I_\nabla^\Diamond)_{\nabla \in \Sigma}.$$

Any subalgebra of $\mathcal{L}^+$ is called a *complex algebra* over $\mathcal{L}$. For a class K of Σ-algebras, we let CmK denote the class of full complex algebras of algebras in K. Since we use S as the class operation giving isomorphic copies of subalgebras, this means that SCmK denotes the class of isomorphism types of complex algebras over K. We say that a Σ_{BA}-algebra $\mathcal{A}$ is *representable* over a class of Σ-algebras K if it belongs to this class SCmK. If the class K is in fact a set consisting of one algebra $\mathcal{L}$, we also say that $\mathcal{A}$ is representable over $\mathcal{L}$; observe that this is equivalent to saying that there is a *representation of* $\mathcal{A}$ *over* $\mathcal{L}$, that is, an embedding $rep : \mathcal{A} \rightarrowtail \mathcal{L}^+$.

Complex algebras are the prime examples of Boolean algebras with operators. An *operator* on a Boolean algebra $(A, +, -, 0)$ is an operation on A that is normal (meaning that its value equals 0 whenever one of its arguments equals 0) and additive (that is, it distributes over $+$ in each of its arguments). Given a similarity type Σ, a *Boolean algebra with* Σ*-operators* is a Σ_{BA}-algebra $(A, +, -, 0, \diamond)$ such that each operation $\diamond_\nabla$ ($\nabla \in \Sigma$) is an operator on the Boolean algebra $(A, +, -, 0)$.

We will need the following fact.

Theorem 2.1 *For any variety* V *of* Σ*-algebras,* SCmV *is an elementary class of Boolean algebras with* Σ*-operators.*

PROOF. It is easy to see that complex algebras are Boolean algebras with operators. In order to prove that SCmV is an elementary class, by the Keisler–Shelah theorem it suffices to show that it is closed under ultraproducts and ultraroots.

The latter is straightforward: take any Σ_{BA}-type algebra $\mathcal{A}$ and an ultrapower $\mathcal{A}^J/U$ of it such that $\mathcal{A}^J/U$ belongs to SCmV. Since $\mathcal{A}$ can be embedded in $\mathcal{A}^J/U$ via the diagonal embedding, it is immediate that $\mathcal{A}$ belongs to $\mathsf{SSCmV} = \mathsf{SCmV}$.

Now suppose that $(\mathcal{A}_j)_{j\in J}$ is a family of Σ_{BA}-algebras in SCmV. That is, for each $j \in J$ there is an $\mathcal{L}_j$ in V such that $\mathcal{A}_j \rightarrowtail {\mathcal{L}_j}^+$. Consider an ultraproduct $\mathcal{A} = (\prod_{j\in J} \mathcal{A}_j)/U$. It is a well-known fact (cf. [4], Lemma 3.6.5) that

$$(\prod_{j\in J} {\mathcal{L}_j}^+)/U \;\rightarrowtail\; \big((\prod_{j\in J} \mathcal{L}_j)/U\,\big)^+.$$

But $(\prod_{j\in J} \mathcal{L}_j)/U$ belongs to V, since V is a variety and hence closed under ultraproducts. Thus the structure $((\prod_{j\in J} \mathcal{L}_j)/U)^+$ is in CmV; it therefore follows from $\mathcal{A} \rightarrowtail (\prod_{j\in J} {\mathcal{L}_j}^+)/U$ that $\mathcal{A}$ belongs to SCmV, as required. QED

This result actually holds for any elementary class V.

3 Games

Let us fix, for the rest of the paper, a finite[1] similarity type Σ, and a variety V of Σ-algebras. We will also fix an enumeration $\{\varepsilon_i : i < \omega\}$ of a set of equations defining V.

It is our aim in this section to define the game that we will use to characterize complex algebras. The key concept employed in our game — the playing board as it were — is that of a *network*. In order to define this, we use the notion of a *partial algebra*.

Definition 3.1 A *partial* Σ*-algebra* is a structure $\mathcal{N} = (N, I)$ such that I is a function interpreting each function symbol ∇ in Σ as a partial operation I_∇ on N of arity q.

[1]This is for simplicity; our results hold for any recursive similarity type Σ.

Analogously to the case of total Σ-algebras, a Σ-term $\tau(x_1,\ldots,x_n)$ can be evaluated in a partial Σ-algebra $\mathcal{N}=(N,I)$ under any assignment θ of its free variables to values in N; we denote the resulting value by τ^θ. The evaluation is partial in that τ^θ need not always exist. Now we say that an equation $\sigma\approx\tau$ *holds* in a partial Σ-algebra $\mathcal{N}$, or that $\mathcal{N}$ *satisfies* the equation, if for every assignment θ of the free variables of σ and τ, if both σ^θ and τ^θ exist then they are equal. A partial Σ-algebra $\mathcal{N}$ is called a *partial* V*-algebra* if it satisfies all the equations of V, and a *partial* V*-algebra of grade* r (where $r\leq\omega$) if it satisfies the equations $\{\varepsilon_i : i<r\}$.

Let (N,I) and (N',I') be two partial Σ-algebras. Then we say that (N,I) is a *partial subalgebra* of (N',I') if $N\subseteq N'$, and for any $\nabla\in\Sigma$ and $k_1,\ldots,k_q\in N$, if $I_\nabla(k_1,\ldots,k_q)$ is defined then $I'_\nabla(k_1,\ldots,k_q)$ is also defined and $I'_\nabla(k_1,\ldots,k_q)=I_\nabla(k_1,\ldots,k_q)$. $\lhd$

Observe that in particular, the constant (zero-ary) function symbols need not obtain an interpretation in a partial algebra.

Definition 3.2 Given a Σ_{BA}-algebra $\mathcal{A}=(A,+,-,0,\Diamond)$, a *network over* $\mathcal{A}$ is a structure $\mathcal{N}=(N,I,\lambda)$ such that (N,I) is a finite partial Σ-algebra and λ is a map: $N\to A$. Elements of N are called *nodes*, and λ is called the *labelling* of the network. The *empty network* $(\varnothing,\varnothing,\varnothing)$ is denoted as $\mathcal{N}_\varnothing$.

A network (N,I,λ) is called a V*-network (of grade* r*)* if (N,I) is a partial V-algebra (of grade r). (N,I,λ) is said to be *coherent* if $\lambda(k)\neq 0$ for each node $k\in N$, and in addition, λ satisfies the following condition, for each function symbol ∇ and all nodes $k_1,\ldots,k_q\in N$ such that $I_\nabla(k_1,\ldots,k_q)$ is defined:

$$\lambda(I_\nabla(k_1,\ldots,k_q))\cdot\Diamond_\nabla(\lambda(k_1),\ldots,\lambda(k_q))\;\neq\;0.$$

$\lhd$

Where $\mathcal{A}$ can be recovered from the context, we will simply say 'network' instead of 'network over $\mathcal{A}$'.

The intuition about this notion is that a network over a Σ_{BA}-algebra $\mathcal{A}$ provides structured information about the representability potential of $\mathcal{A}$. In particular, suppose that $\mathcal{A}$ is representable, say, via the representation *rep* over the algebra $\mathcal{L}$ in V. Further, suppose that the network $\mathcal{N}=(N,I,\lambda)$ *matches* with *rep* in the sense that (i) (N,I) is a partial subalgebra of $\mathcal{L}$ and (ii) λ satisfies $k\in rep(\lambda(k))$ for all nodes $k\in N$. Then $\mathcal{N}$ is a coherent V-network — as an easy calculation shows. Such a network can be seen as a finite approximation of the representation *rep*: the network only provides partial information concerning the representation.

We are interested in certain relations between networks, like one coherent V-network approximating a representation better than another. In general, we need to define when one network *extends* or provides more information than another.

Definition 3.3 A network $\mathcal{N}'=(N',I',\lambda')$ over an algebra $\mathcal{A}$ is said to *extend* or to be an *extension of* a network $\mathcal{N}=(N,I,\lambda)$, notation: $\mathcal{N}\lhd\mathcal{N}'$, if (N,I) is a partial subalgebra of (N',I') and λ' is a *tightening* of λ: that is, $\mathcal{A}\models\lambda'(k)\leq\lambda(k)$ for all $k\in N$. $\lhd$

Note that if $\mathcal{N}\lhd\mathcal{N}'$ and $\mathcal{N}'$ is a coherent V-network of grade r then so is $\mathcal{N}$.

In the sequel, we will be interested in a number of ways to extend a network, in particular, the following three:

1. adding new points, that is, enlarging the network,
2. tightening the labelling,
3. providing more values for the partial operations interpreting the function symbols.

Definition 3.4 Let $\mathcal{N} = (N, I, \lambda)$ be a network over the Σ_{BA}-algebra $\mathcal{A}$.

1. For an object n (either being a node of the network or not), $\mathcal{N}(n)$ is defined to be the network $(N \cup \{n\}, I, \lambda')$, where the labelling λ' is given by
$$\lambda'(x) = \begin{cases} 1 & \text{if } x = n \notin N, \\ \lambda(x) & \text{otherwise.} \end{cases}$$
2. For a node $k \in N$ and an element $a \in A$, $\mathcal{N}(k : a)$ denotes the network (N, λ', I), where the labelling λ' is given by
$$\lambda'(x) = \begin{cases} \lambda(x) \cdot a & \text{if } x = k, \\ \lambda(x) & \text{otherwise.} \end{cases}$$
3. For an operator symbol $\nabla \in \Sigma$ and nodes $k_0, k_1, \ldots, k_q$ of $\mathcal{N}$, the network $\mathcal{N}(\nabla, \overline{k} \mapsto k_0)$ (where $\overline{k}$ is the tuple $(k_1, \ldots, k_q)$) is defined to be (N, I', λ), where $I'_\heartsuit = I_\heartsuit$ for all function symbols $\heartsuit$ different from ∇, while the interpretation of ∇ is given by
$$I'_\nabla(\overline{x}) = \begin{cases} k_0 & \text{if } \overline{x} = \overline{k} \text{ and } I_\nabla(\overline{x}) = \uparrow, \\ I_\nabla(\overline{x}) & \text{otherwise.} \end{cases}$$
Here '$I_\nabla(\overline{x}) = \uparrow$' denotes that $I_\nabla(\overline{x})$ is undefined.

$\triangleleft$

We are now ready to define the games.

Definition 3.5 Let $\mathcal{A}$ be some Σ_{BA}-algebra, let $\mathcal{N}$ be some network over $\mathcal{A}$ and let $\alpha \leq \omega$ be an ordinal. We define a *game* $G_\alpha(\mathcal{N}, \mathcal{A}, \bigvee)$ between two players: $\forall$ (male) and $\exists$ (female).

A *match* of the game consists of α *rounds*, numbered $0, 1, \ldots, i, \ldots$ for $i < \alpha$. The match starts with the network $\mathcal{N}_0 = \mathcal{N}$; during the match, the players build a sequence of networks $(\mathcal{N}_{i+1})_{i<\alpha}$. All networks are over $\mathcal{A}$. Each round consists of a *move* made by $\forall$ and a *response* move made by $\exists$. In round i, for each $i < \alpha$, the playing board consists of the network $\mathcal{N}_i$. The actions of the players during the round define a new network $\mathcal{N}_{i+1}$ which forms the playing board for the next round, $i + 1$; and so on.

The moves of the players are subject to the following constraints. In each round of the game, $\forall$ has a choice between four kinds of move, listed below. Suppose that he is about to make a move in round i of the game and that $\mathcal{N}_i = (N, I, \lambda)$ is the network forming the playing board. His move can be seen as a proposal to extend $\mathcal{N}_i$ in some way; in her response, $\exists$ can choose either to *accept* or *reject* his proposal. There is some fixed infinite set Q at $\exists$'s disposal from which to draw new nodes, if her response entails enlarging $\mathcal{N}_i$.

(α) In the first type of move, $\forall$ chooses a node k of the network $\mathcal{N}_i$ and an element a of the algebra $\mathcal{A}$. If $\exists$ accepts this move, $\mathcal{N}_{i+1}$ is defined as the network $\mathcal{N}_i(k : a)$; otherwise, it is $\mathcal{N}_i(k : -a)$.

(β) The second type of move consists of $\forall$ choosing a node k of $\mathcal{N}_i$, a function symbol $\nabla \in \Sigma$ of arity q, say, and elements $a_1, \ldots, a_q$ of $\mathcal{A}$. In this case, rejection by $\exists$ gives the network
$$\mathcal{N}_{i+1} = \mathcal{N}_i(k : -\Diamond_\nabla(a_1, \ldots, a_q)).$$
If, on the other hand, she accepts $\forall$'s proposal, she must choose objects $m_1, \ldots, m_q$ (which may or may not be nodes of $\mathcal{N}$) such that if $I_\nabla(m_1, \ldots, m_q)$ exists then it is k, and define the new network as
$$\mathcal{N}_{i+1} = \mathcal{N}_i(\nabla, k, \overline{m}, \overline{a}),$$
where $\overline{a} = (a_1, \ldots, a_q)$, $\overline{m} = (m_1, \ldots, m_q)$ and
$$\begin{aligned}&\mathcal{N}_i(\nabla, k, \overline{m}, \overline{a}) = \\ &\quad \mathcal{N}_i(m_1) \cdots (m_q)(\nabla, \overline{m} \mapsto k)(m_1 : a_1) \cdots (m_q : a_q).\end{aligned}$$

(γ) In the third kind of move, $\forall$ points out a function symbol $\nabla \in \Sigma$ and nodes $k_1, \ldots, k_q$ of $\mathcal{N}_i$. In this case, $\exists$ has no other choice but to accept, and she does so by choosing a point m (which may or may not be a node of the old network). The new network is defined as
$$\mathcal{N}_{i+1} = \mathcal{N}_i(m)(\nabla, (k_1, \ldots, k_q) \mapsto m).$$

(δ) Finally, in the fourth kind of move, $\forall$ simply picks a non-zero element a of the algebra $\mathcal{A}$. $\exists$ has to accept, by providing an object k; the new network is defined as
$$\mathcal{N}_{i+1} = \mathcal{N}_i(k)(k : a).$$

$\exists$ is said to *win* the match if $\mathcal{N}_0$ and every $\mathcal{N}_{i+1}$ $(i < \alpha)$ is a coherent V-network of grade α; if she does not win, then $\forall$ does. $\lhd$

It is in a sense the aim of the first player, $\forall$, in the game $G_\alpha(\mathcal{N}, \mathcal{A}, \mathsf{V})$ to show that the starting network $\mathcal{N}$ is *not* an approximation of a representation of the algebra $\mathcal{A}$ over some V-algebra, while the second player $\exists$ wants to show the contrary. Less confrontationally, we can view $\exists$ as a Ph.D. student in the Faculty of Representability of Algebras, and $\forall$ as examiner of her dissertation on $\mathcal{A}$ [10]. The idea of a move in the game is that $\forall$ challenges $\exists$ to provide more information about a possible representation over nodes of the networks played during the match. $\forall$ makes a type (α) move if he wants to find out whether $\exists$ thinks a or $-a$ 'holds' at some point k of the representation. In a real representation, if $\Diamond_\nabla(a_1, \ldots, a_q)$ holds at a point k then there must be points m_i at which a_i holds $(1 \leq i \leq q)$ and with $I_\nabla(m_1, \ldots, m_q) = k$. A move of type ($\beta$) tests this. Any function I_∇ must be defined on all tuples of points of appropriate length, and this is tested in type (γ) moves. Moves of type (δ) test injectivity of the representation, as we will see later on.

Definition 3.6 $\exists$ is said to have a *winning strategy* for the game $G_\alpha(\mathcal{N}, \mathcal{A}, \mathsf{V})$ if there is a set of rules that tells her how to respond to $\forall$ in each round of a match, depending on play so far, such that she wins any match in which she follows these rules. $\lhd$

The notion of a winning strategy can be formalized by certain functions but it is not helpful here to do so.

In the sequel, we will make tacitly use of the observation that if $\exists$ uses a winning strategy in a match of the game $G_\omega(\mathcal{N}, \mathcal{A}, \mathsf{V})$, then at each round i she has a winning strategy for the game $G_\omega(\mathcal{N}_i, \mathcal{A}, \mathsf{V})$. A similar observation holds for the finite-length games, except that we must remember that games of different finite lengths require networks of different grades (the reason for this requirement will be seen in section 5). Observe furthermore that it follows from the definitions that the sequence of networks obtained in any match of the game is in fact a *chain*: $\mathcal{N}_0 \lhd \mathcal{N}_1 \lhd \cdots$.

Remark 3.7 From Definition 3.5, it may seem that the game $G_\alpha(\mathcal{N}, \mathcal{A}, \mathsf{V})$ is not well-defined, in that it depends on the set Q of 'new' nodes available to $\exists$. It may also seem that to some moves of $\forall$, $\exists$ has an *infinite* choice of networks to respond with; it is crucial for our later results that this infinite choice in fact boils down to a finite one. These two issues are closely related. We felt that dealing with them more formally in the definition of the game would have gone at the cost of transparency. Nevertheless, in section 5 we need more precision concerning this issue, so let us discuss it now in some detail.

As to the first point, it is easily seen that if $\exists$ has a winning strategy in the game $G_\alpha(\mathcal{N}, \mathcal{A}, \mathsf{V})$ using *some* set Q of spare nodes, even a finite one, then she has a winning strategy with $Q = \omega$, because it does not matter what the elements of Q actually are, and during the α rounds of any match she will only need at most ω new nodes. So we will formally take $Q = \omega$ in the game definition, but sometimes allow other sets Q in practice.

In more detail, and addressing the second point, consider the situation arising after $\forall$ has made a type (γ) move, choosing the function symbol ∇ and the nodes $k_1, \ldots, k_q$ (recall that q denotes the arity of ∇), and suppose that $I_\nabla(k_1, \ldots, k_q)$ is undefined for the old network $\mathcal{N} = (N, I, \lambda)$. Then $\exists$ is forced to give a value for the new interpretation function I'_∇ on the tuple $(k_1, \ldots, k_q)$. It is obvious that there are only finitely many old candidates, but if she chooses to enlarge the network with a new object, isn't there the whole infinite set Q to choose from? The formal answer to this question is of course affirmative, but the point is that *if* $\exists$ chooses to enlarge the network, it does not matter at all *which* object from the set Q she chooses. We might as well have required that whenever $\exists$ needs to extend a network $\mathcal{N}$ with a new node, she takes some *canonically chosen* object $\#_N$ from $Q \setminus N$. In fact, we may (and later on, will) assume that $\#$ is a recursive function on sets of nodes of networks. For instance, we may require that the nodes of a network are always taken from the set of natural numbers, and indeed that the set of nodes of any network form a set of the form $\{0, \ldots, i\}$. This would mean that we could take $\#_N$ to be the *size* of the network — this example in fact inspired our notation. In this alternative but equivalent set-up, it is clear that to any type (γ) move of $\forall$, the only choice that $\exists$ has is which object to pick from the finite set $N^* = N \cup \{\#_N\}$.

Obviously, the same applies to the other kinds of move for the first player. Concerning $\exists$'s response to a type (β) move of $\forall$, it is convenient to introduce some notation. Recall that in case $\exists$ accepts $\forall$'s move, she has to choose q witnesses (where q denotes the arity of the function symbol involved) which may but need not be nodes of the old network. Assume that she chooses the nodes $m_1, \ldots, m_q$ in this order, and that when she chooses a *new* node m_i then this will canonically be the object $\#_{N \cup \{m_1, \ldots, m_{i-1}\}}$. In other words, let $K^q(N)$ be the set of those q-tuples $(m_1, \ldots, m_q)$ such that for all i with $1 \leq i \leq q$, $m_i \in (N \cup \{m_1, \ldots, m_{i-1}\})^*$. Then an affirmative answer of $\exists$ to a type (β) move of $\forall$ on the network $\mathcal{N}$ consists of choosing a tuple from the finite set $K^q(N)$.

4 A game characterization

It is our aim in this section to prove the following result:

Theorem 4.1 *Let $\mathcal{A}$ be a countable Boolean algebra with Σ-operators. Then $\exists$ has a winning strategy for the game $G_\omega(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$ if and only if $\mathcal{A}$ belongs to SCmV.*

The proof of this theorem naturally falls out in two parts. First we show the soundness part of the game characterization (that is, the right to left direction of the Theorem). Here we do not need to restrict to countable algebras.

Proposition 4.2 *Let $\mathcal{A}$ be an algebra in SCmV. Then $\exists$ has a winning strategy for the game $G_\omega(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$.*

Proof. If $\mathcal{A} = (A, \diamond)$ belongs to SCmV, then there is some algebra $\mathcal{L} = (L, f)$ and a representation map $rep : A \to \mathcal{P}(L)$ which embeds $\mathcal{A}$ into $\mathcal{L}^+$.

Now let $\mathcal{N} = (N, I, \lambda)$ be a network over $\mathcal{A}$. We say that $\mathcal{N}$ *matches with rep* if (N, I) is a partial subalgebra of $\mathcal{L}$, while λ satisfies

$$k \in rep(\lambda(k)) \text{ for all nodes } k \in N. \tag{1}$$

In words, $\mathcal{N}$ matches with rep if it can be seen as an approximation for the representation.

By Remark 3.7, it suffices to prove the proposition under the assumption that the set Q of spare nodes available to $\exists$ during play is L. The idea of the winning strategy for $\exists$ is that during a match of the game, she will maintain the condition that the current network matches with rep. Evidently, any such network is a coherent V-network. So in order to show that this strategy works, it is sufficient to prove that it sees $\exists$ through one single round of the game. To be more precise, we have to prove that if $\mathcal{N} = (N, I, \lambda)$ is a network matching with rep, then $\exists$ can counter any move of $\forall$ on this network with a new network $\mathcal{N}' = (N', I', \lambda')$ that also matches with rep.

We only treat the case in which $\forall$ makes a type (β) move — say, he picks the function symbol ∇, the node m, and the elements $a_1, \ldots, a_q$ of A (here, q denotes the arity of ∇). The first thing that $\exists$ does is to check whether m belongs to $rep(\diamond_\nabla(\overline{a}))$; if this is *not* the case, then (naturally) she rejects the proposal, whence the new network is defined as $\mathcal{N}' = \mathcal{N}(m : -\diamond_\nabla(\overline{a}))$. Since the only difference between $\mathcal{N}$ and $\mathcal{N}'$ concerns the new label of m, in order to show that $\mathcal{N}'$ matches with rep it suffices to check (1) for m and λ'. But $\lambda'(m) = \lambda(m) \cdot -\diamond_\nabla(\overline{a})$, whence $rep(\lambda'(m)) = rep(\lambda(m)) \setminus rep(\diamond_\nabla(\overline{a}))$. By our assumption on $\mathcal{N}$ we have that $m \in rep(\lambda(m))$ and by our case assumption we have that $m \notin rep(\diamond_\nabla(\overline{a}))$. Thus we find that indeed, $m \in rep(\lambda'(m))$.

Now suppose that on the other hand, m does belong to $rep(\diamond_\nabla(\overline{a}))$. Note that since rep is a homomorphism, we have that

$$rep(\diamond_\nabla(\overline{a})) = f_\nabla^\diamond(rep(a_1), \ldots, rep(a_q)).$$

Thus by definition of $f_\nabla^\diamond$, there must be elements $k_1, \ldots, k_q \in L$ such that $m = f_\nabla(\overline{k})$ and $k_i \in rep(a_i)$ for each i. Naturally, $k_1, \ldots, k_q$ are the objects that $\exists$ picks as her response to $\forall$'s move. The new network $\mathcal{N}'$ is defined as $\mathcal{N}(\nabla, m, \overline{k}, \overline{a})$. It is obvious from the definitions that $\mathcal{N}'$ satisfies (1), so let us check now that the underlying partial algebra (N', I') of $\mathcal{N}'$ is a partial subalgebra of $\mathcal{L}$. But since (N, I) is a partial subalgebra of $\mathcal{L}$ and I' is like I save perhaps for its value on $\overline{k}$, this follows from our assumption that $f_\nabla(\overline{k}) = m$.

It turns out that in this case, $\exists$ manages to reach the end of the round with a network that matches the representation. The proof for the other cases is very similar. QED

The completeness part of Theorem 4.1 is the hard direction.

Proposition 4.3 *Let $\mathcal{A}$ be a countable Boolean algebra with Σ-operators, and suppose that $\exists$ has a winning strategy for the game $G_\omega(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$. Then $\mathcal{A}$ is representable as a complex algebra over V.*

PROOF. We will consider a match of the game in which both players play according to a special strategy. Basically, the strategy of $\forall$ will consist of listing all moves that become possible during the match and actually making each one of these moves at some stage of the play; obviously, $\exists$ will use her winning strategy. We will prove that we can 'read off' a representation of the algebra $\mathcal{A}$ from the chain of networks arising during this particular match of the game.

For notational simplicity, we assume that the similarity type Σ has only one function symbol ∇, of arity q. First of all, we define the countable set

$$P := (Q \times A) \cup (Q \times {}^qA) \cup ({}^qQ) \cup A$$

coding the *potential moves* for $\forall$, where Q is the (countable) set of objects from which the nodes of the networks are taken. For any given network $\mathcal{N}$ with nodes in Q, such a move may be *possible* or not, according to whether the elements of Q occurring in the potential move are nodes of $\mathcal{N}$ or not.

We may assume that the strategy of $\forall$ is thus that whenever a potential move p is possible in some network during the match, then $\forall$ will make this move at some stage during the game. (This could be done by enumerating the set P with order type ω, and asking $\forall$ at each round to make the move p that has the least index among all possible moves that he has not made yet.)

Consider the chain of networks arising in such a match of the game: $\mathcal{N}_0 \lhd \mathcal{N}_1 \lhd \cdots$. We will use this chain to define a partial Σ-algebra $\mathcal{L} = (N, I)$. Its carrier N is given as the union of all the N_i:

$$N = \bigcup_{i \in \omega} N_i.$$

Recall, for the definition of the function $I_\nabla : {}^qN \to N$, that the sequence of networks associated with any match of the game form a chain. In particular, the underlying algebra of a network will be a partial subalgebra of the underlying algebra of any network arising later; hence, if $I^i_\nabla(\overline{k})$ is defined at some stage i, we have that $I^i_\nabla(\overline{k}) = I^j_\nabla(\overline{k})$ for all stages j at which $I^j_\nabla(\overline{k})$ is defined. This means that the following is an unambiguous definition of the interpretation I_∇ of the function symbol ∇ on N in $\mathcal{L}$:

$$I_\nabla(\overline{k}) = \begin{cases} I^i_\nabla(\overline{k}) & \text{for any } i \text{ such that } I^i_\nabla(\overline{k}) \text{ is defined,} \\ \uparrow & \text{if no } I^i_\nabla(\overline{k}) \text{ is defined.} \end{cases}$$

Finally, we define a labelling $\lambda : N \to \mathcal{P}(A)$ by

$$\lambda(k) = \{a \in A \mid a \geq \lambda_i(k) \text{ for some } i \in \omega\}.$$

Observe that unlike the labellings that we have seen so far, λ labels nodes with *subsets* of the algebra $\mathcal{A}$, not with elements of it.

It is our ultimate aim to show that the following map is a representation of $\mathcal{A}$ over (N, I):

$$rep : \quad a \mapsto \{k \in N \mid a \in \lambda(k)\}.$$

To this end, we will prove the following claims. We will use the fact that, since $\exists$ plays according to her winning strategy, she wins this match of the game; this means that each $\mathcal{N}_i$ is a coherent V-network.

Claim 1 (N, I) is a total algebra and belongs to V.

PROOF OF CLAIM We first show that I_∇ is a total operation. Take some elements $k_1, \ldots, k_q$ in N. By our definition of N, there must be some stage i of the match at which each object k_i is present as a node of the network $\mathcal{N}_i$. By our assumption on $\forall$'s strategy, this means that at some stage j of the game, $\forall$ makes a type (γ) move picking ∇ as the operator symbol together with the nodes $k_1, \ldots, k_q$. It then follows from the rules of the game that $I_{j+1}(\overline{k})$ is defined; but then by our definition of I_∇, we have that I_∇ is also defined on $\overline{k}$.

Then, once we know that (N, I) is a total algebra, it is trivial to show that it satisfies all equations of V. For, since each $\mathcal{N}_i$ is a coherent V-network, (N, I) is a partial V-algebra, so by its totality it belongs to V. ◀

Claim 2 For all $k \in N$, $\lambda(k)$ is an ultrafilter over $\mathcal{A}$.

PROOF OF CLAIM It follows that $\lambda(k)$ is a filter for each node k, by $\lambda_i(k) \leq \lambda_j(k)$ for $i \geq j$. Since each network $\mathcal{N}_i$ $(i < \omega)$ is coherent, these filters $\lambda(k)$ $(k \in N)$ must all be proper.

In order to show that λ in fact labels with ultrafilters, take some node $k \in N$ and an element a of the algebra. It follows from our assumption on the strategy of $\forall$ that at some stage i of the match, $\forall$ makes the type (α) move picking k and a. This means that at the next stage of the game we have either $\lambda_{i+1}(k) \leq a$ or $\lambda_{i+1}(k) \leq -a$. In the first case we obtain that $a \in \lambda(k)$, in the second case that $-a \in \lambda(k)$. This proves that indeed $\lambda(k)$ is an ultrafilter. ◀

Claim 3 *rep* is a homomorphism.

PROOF OF CLAIM Since $\lambda(k)$ is an ultrafilter of $\mathcal{A}$ for each node $k \in N$, it is straightforward to show that *rep* is a homomorphism with respect to the Boolean operations. Hence, we restrict ourselves to proving that *rep* is a homomorphism with respect to the operation interpreting ∇. In other words, we have to show that

$$rep(\Diamond_\nabla(a, b)) = I^\Diamond_\nabla(rep(a), rep(b)) \quad \text{for all } a, b \in A$$

(in order to keep notation simple here, we assume that ∇ is binary).

We first establish the left-to-right inclusion: assume that $k \in rep(\Diamond_\nabla(a, b))$. By definition, this means that $\Diamond_\nabla(a, b) \geq \lambda_i(k)$, for some stage $i \in \omega$. It follows from our assumption on $\forall$'s strategy that at some stage j of the match, he makes a type (β) move picking k, ∇, a, and b. It is clear that $\exists$ does not reject this proposal, for if she did, we would have $k \in rep(-\Diamond_\nabla(a, b))$ contradicting $k \in rep(\Diamond_\nabla(a, b))$. So she accepts: this means that at stage $j+1$, there are nodes k_a and k_b with $I^{j+1}_\nabla(k_a, k_b) = k$, $\lambda_{j+1}(k_a) \leq a$ and $\lambda_{j+1}(k_b) \leq b$. From this it easily follows that $I_\nabla(k_a, k_b) = k$, $k_a \in rep(a)$ and $k_b \in rep(b)$ and thus that $k \in I^\Diamond_\nabla(rep(a), rep(b))$.

For the other inclusion, assume that $k \in I^\Diamond_\nabla(rep(a), rep(b))$. It is easily seen that this implies that at some stage i of the match, there are nodes $k_a, k_b \in N_i$ such that $I^i_\nabla(k_a, k_b) = k$, $\lambda_i(k_a) \leq a$, $\lambda_i(k_b) \leq b$ and either $\lambda_i(k) \leq \Diamond_\nabla(a, b)$ or $\lambda_i(k) \leq -\Diamond_\nabla(a, b)$. But since $\lambda_i(k) \leq -\Diamond_\nabla(a, b)$ would contradict the coherency of $\mathcal{N}_i$, it must be the case that $\lambda_i(k) \leq \Diamond_\nabla(a, b)$. From this it follows that $k \in rep(\Diamond_\nabla(a, b))$. ◀

Claim 4 *rep* is injective.

Proof of Claim It is sufficient to prove that $rep(a) \neq 0$ for any non-zero element a of the algebra. But this is taken care of in a type (δ) move of $\forall$: at some stage i of the game, he will play the element a. Since $\mathcal{N}_{i+1}$ is defined as the network $\mathcal{N}_i(k)(k : a)$, with k being the object chosen by $\exists$, we have $\lambda_{i+1}(k) \leq a$ and hence, $a \in \lambda(k)$ and $k \in rep(a)$. ◀

It is immediate by these claims that indeed, *rep* is a representation embedding $\mathcal{A}$ in the full complex algebra $(N, I)^+$ of the V-algebra (N, I). QED

5 The axiomatization

Recall that we fixed a variety V of Σ-algebras, where Σ is a finite similarity type, and a sequence $\langle \varepsilon_i : i < \omega \rangle$ of equations defining V. It is the aim of this section to prove the main theorem of this paper. That is, we will provide a collection $\Phi(\mathsf{V})$ of universal first-order sentences (in the algebraic language of similarity type Σ_{BA}) that axiomatize the class SCmV. We will proceed in three steps. First, we will prove that for any $\mathcal{N}$ and $\mathcal{A}$, $\exists$ has a winning strategy in the game $G_\omega(\mathcal{N}, \mathcal{A}, \mathsf{V})$ if and only if she has winning strategies for all games of finite length — the $G_i(\mathcal{N}, \mathcal{A}, \mathsf{V})$ for $i \in \omega$. Second, we will recursively define a collection of sentences $(\varphi_i)_{i \in \omega}$ such that for each i, φ_i holds in a Σ_{BA}-algebra $\mathcal{A}$ if and only if $\exists$ has a winning strategy for the game $G_i(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$. In the third and last part of this section, we show that these two results provide sufficient material for proving the main theorem.

As we announced, we first show that for $\exists$, having a winning strategy in a game of infinite length is equivalent to having winning strategies in infinitely many games of finite length.

Theorem 5.1 *For any Σ_{BA}-algebra $\mathcal{A}$ and any network $\mathcal{N}$ over $\mathcal{A}$, $\exists$ has a winning strategy in the game $G_\omega(\mathcal{N}, \mathcal{A}, \mathsf{V})$ if and only if she has a winning strategy for every game $G_i(\mathcal{N}, \mathcal{A}, \mathsf{V})$ of finite length $i \in \omega$.*

Proof. The left-to-right direction of the Theorem is obvious, so we will only prove the other direction. Assume that $\exists$ has a winning strategy for each game $G_i(\mathcal{N}, \mathcal{A}, \mathsf{V})$ of finite length i. We have to supply her with a winning strategy for the game $G_\omega(\mathcal{N}, \mathcal{A}, \mathsf{V})$.

Call a network $\mathcal{N}$ *safe for* $\exists$ if for infinitely many j, she has a winning strategy in the game $G_j(\mathcal{N}, \mathcal{A}, \mathsf{V})$. Note that the initial network $\mathcal{N}$ is safe for $\exists$ by assumption, and that any safe network is a coherent V-network of grade α for every $\alpha < \omega$, and hence is a coherent V-network. Now the idea of $\exists$'s strategy in $G_\omega(\mathcal{N}, \mathcal{A}, \mathsf{V})$ is to maintain the condition that the current network is safe for her. Obviously, in order to show that this is a *winning* strategy, it suffices to show that she can survive one round of the game maintaining this condition.

Hence, suppose that we are in the i-th round of the game $G_\omega(\mathcal{N}, \mathcal{A}, \mathsf{V})$; let $\mathcal{N}_i$ be the network board of this round, and assume that $\mathcal{N}_i$ is safe for $\exists$. Now assume that $\forall$ makes his i-th move; as we saw in Remark 3.7, $\exists$ has a *finite* choice of networks to respond with. Since there are infinitely many j for which she has a winning strategy in the game $G_j(\mathcal{N}_i, \mathcal{A}, \mathsf{V})$, this means that there must be at least one of these responses, say $\mathcal{N}'$, on which she has a winning strategy in the game $G_{j-1}(\mathcal{N}', \mathcal{A}, \mathsf{V})$ for infinitely many j. Obviously, this means that this $\mathcal{N}'$ is safe for her; hence, if she chooses it to be her response in the i-th round of the game of infinite length, she has maintained her condition. QED

We have now arrived at the second and hard part of the section in which we have to provide the first-order formulas characterizing SCmV. The crucial concept that we employ here is that of a *term network*.

Definition 5.2 A *term network* is a structure $\mathcal{N} = (N, I, \tau)$ such that (N, I) is a finite partial Σ-algebra and τ is a *term labelling*, that is, a map assigning a Σ_{BA}-term τ_k to each node $k \in N$ of the network.

Given a term network $\mathcal{N}$, $\mathrm{Var}(\mathcal{N})$ denotes the set of (algebraic) variables occurring in the term labels of $\mathcal{N}$. ◁

We will use notation in line with that adopted for networks to denote extensions of term networks: for instance, given a term network $\mathcal{N} = (N, I, \tau)$, a node k of $\mathcal{N}$, and a Σ_{BA}-term σ, we let $\mathcal{N}(k:\sigma)$ denote the term network (N, I, τ') where τ' is defined by $\tau'_x = \tau_x$ for $x \in N$ such that $x \neq k$, while τ'_k is the Σ_{BA}-term $\tau_k \cdot \sigma$. We assume that there is a canonical way of adding a new node to a term network (cf. Remark 3.7); this new node is denoted $\#_N$ and we assume that $\#$ is in fact a recursive function.

The basic idea is that term networks provide *structure to the indices* of the variables occurring in the formulas characterizing the class SCmV. Later we will come to this point in more detail; let us first see how term networks relate to ordinary networks. The connection is that given a Σ_{BA}-algebra $\mathcal{A}$, a term network corresponds to a family of (ordinary) networks over $\mathcal{A}$, in a sense to be made precise in the definition below.

Definition 5.3 Let $\mathcal{N} = (N, I, \tau)$ be a term network and let $\mathcal{A}$ be a Σ_{BA}-algebra. Given an assignment $\theta : \mathrm{Var}(\mathcal{N}) \to A$, $\mathcal{N}^\theta$ denotes the network (N, I, λ^θ) over $\mathcal{A}$, where λ^θ is given by

$$\lambda^\theta(k) = \tau_k^{(\mathcal{A},\theta)}, \text{ for any node } k \in N.$$

◁

In words, the network $\mathcal{N}^\theta$ over $\mathcal{A}$ that we associate with an assignment θ and a term network $\mathcal{N}$ consists of the finite partial algebra underlying the term network, while the label $\tau_k^{(\mathcal{A},\theta)}$ of a node k is obtained by interpreting the label term τ_k of k in $\mathcal{A}$ according to the assignment θ.

The definition of the formula set $\Phi(\mathsf{V})$ is given in Figure 1. Recall that we use the abbreviations $q = ar(\nabla)$, $\overline{k} = (k_1, \ldots, k_q)$, etc., $K^q(N)$ and N^* are as defined in Remark 3.7, $\mathcal{N} = (N, I, \lambda)$ is an arbitrary term network, and $i, r < \omega$ are arbitrary.

Observe that the definition of the sentences φ_i uses some auxiliary formulas expressing properties of Σ_{BA}-algebras. It is easily checked that all formulas involved are universal. Each of these formulas is indexed by a term network; the free variables of the formula are the variables occurring in (the term labels of) this network. Roughly speaking, if we are given a Σ_{BA}-algebra $\mathcal{A}$, a term network $\mathcal{N}$, and an assignment θ of the variables occurring in $\mathcal{N}$ to elements of $\mathcal{A}$, then a formula indexed by the network $\mathcal{N}$, evaluated in $\mathcal{A}$ under θ, corresponds to some property that the network $\mathcal{N}^\theta$ induced by θ should satisfy. These correspondences are as follows:

$$\begin{array}{rcl} \psi_0^r(\mathcal{N}) & \sim & \mathcal{N}^\theta \text{ is a coherent } \mathsf{V}\text{-network of grade } r, \\ \psi_i^i(\mathcal{N}) & \sim & \exists \text{ has a winning strategy in } G_i(\mathcal{N}^\theta, \mathcal{A}, \mathsf{V}), \\ \varphi_i & \sim & \exists \text{ has a winning strategy in } G_i(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V}). \end{array}$$

In fact, the formulas are fairly literal transcriptions of the definitions of the game.

For instance, the formula $\psi_{i+1}^r(\mathcal{N})$ is a conjunction of four formulas, each corresponding to a kind of move that the first player may make. The first conjunct, $\alpha_{i+1}^r(\mathcal{N})$, of this

$$
\begin{aligned}
\pi^r(\mathcal{N}) &= \begin{cases} \top & \text{if } (N, I) \text{ is a partial } \mathsf{V}\text{-algebra of grade } r, \\ \bot & \text{otherwise} \end{cases} \\
\chi(\mathcal{N}) &= \bigwedge_{k \in N} \tau_k \not\approx 0 \;\wedge\; \bigwedge_{\nabla \in \Sigma} \bigwedge_{\substack{\overline{k}, k_0 \text{ in } N \\ k_0 = I_\nabla(\overline{k})}} \tau_{k_0} \cdot \nabla(\tau_{k_1}, \ldots, \tau_{k_q}) \not\approx 0. \\
\psi_0^r(\mathcal{N}) &= \pi^r(\mathcal{N}) \wedge \chi(\mathcal{N}) \\
\psi_{i+1}^r(\mathcal{N}) &= \alpha_{i+1}^r(\mathcal{N}) \wedge \beta_{i+1}^r(\mathcal{N}) \wedge \gamma_{i+1}^r(\mathcal{N}) \wedge \delta_{i+1}^r(\mathcal{N}) \\
\alpha_{i+1}^r(\mathcal{N}) &= \forall v \bigwedge_{k \in N} \big(\psi_i^r(\mathcal{N}(k : v)) \vee \psi_i^r(\mathcal{N}(k : -v))\big), \\
&\qquad\quad \text{where } v \text{ is a new variable.} \\
\beta_{i+1}^r(\mathcal{N}) &= \bigwedge_{k \in N} \bigwedge_{\nabla \in \Sigma} \forall \overline{v} \Big(\psi_i^r(\mathcal{N}(k : -\Diamond_\nabla(\overline{v}))) \;\vee \bigvee_{\overline{m} \in M} \psi_i^r(\mathcal{N}(\nabla, k, \overline{m}, \overline{v}))\Big), \\
&\qquad\quad \text{where } \overline{v} = v_1, \ldots, v_q \text{ is a tuple of new variables} \\
&\qquad\quad \text{and } M = \{\overline{m} \in K^q(N) : I_\nabla(\overline{m}) \in \{\uparrow, k\}\}. \\
\gamma_{i+1}^r(\mathcal{N}) &= \bigwedge_{\nabla \in \Sigma} \bigwedge_{\substack{\overline{k} \text{ in } N \\ I_\nabla(\overline{k}) = \uparrow}} \bigvee_{m \in N^*} \psi_i^r(\mathcal{N}(\nabla, \overline{k} \mapsto m)) \\
\delta_{i+1}^r(\mathcal{N}) &= \forall v \Big(v \not\approx 0 \to \bigvee_{m \in N^*} \psi_i^r(\mathcal{N}(m)(m : v))\Big) \\
\varphi_i &= \psi_i^i(\mathcal{N}_\varnothing) \\
\Phi(\mathsf{V}) &= \{\varphi_i \mid i \in \omega\}
\end{aligned}
$$

Figure 1: The Axioms

formula corresponds to a type (α) move of $\forall$; it states that for every choice of an element a of the algebra (represented by the universal quantification $\forall v$) and of a node of the network (represented by the conjunction $\bigwedge_{k\in N}$), the second player has a winning strategy for the game of length i over each network arising by tightening the labelling of $\lambda^\theta(k)$ either with a or with $-a$ (all of this is recursively represented by the formula $\psi_i^r(\mathcal{N}(k:v)) \vee \psi_i^r(\mathcal{N}(k:-v))$). The formulation of the other conjuncts of $\psi_{i+1}^r(\mathcal{N})$ is as direct, as will become evident further on.

The previous example shows how networks are used to bring structure in the information conveyed by the terms occurring in the χ- and ψ-formulas. This kind of structuring is ubiquitous in axiomatizations, but usually far less intricate. For instance, the reader will certainly be used to examples using *natural numbers* as indices to variables, in formulas such as $\bigwedge_{i\neq j} v_i + v_j > 0$. Slightly more complex examples are formulas in which terms are indexed by *pairs of natural numbers*, as in $\bigwedge_{i+j=k} v_{ik} \leq v_{ij}; v_{jk}$. Formulas using term networks are more complicated examples of the same idea.

It should be emphasized that term networks are finite structures that can be coded up and serve as input to Turing machines; all operations on networks that are used in the formulation of the axioms in Figure 1 can in fact be programmed as recursive functions on such codings of networks. For example, given a term network $\mathcal{N} = (N, I, \tau)$ and a number r, the question whether or not (N, I) is a partial V-algebra of grade r can be decided effectively by some Turing machine, and thus, such a machine may check whether a given formula is of the form $\pi^r(\mathcal{N})$. This observation may serve as one of the base cases of an inductive argument showing that for any $i < \omega$, given (i) an algorithm that recursively enumerates a set $\langle \varepsilon_i : i < \omega\rangle$ of equations defining V, (ii) a grade $r < \omega$ and (iii) a finite network $\mathcal{N}$, we can effectively decide whether a given formula is of the form $\psi_i^r(\mathcal{N})$. Since the formulas ψ_i^i grow in length when we increase i this shows that the collection $\Phi(\mathsf{V})$ is recursive, even if our axiomatization for V is only recursively enumerable.

The reader may wonder why the equations holding in the variety V do not show up explicitly in the axiomatization $\Phi(\mathsf{V})$. The point is that our axiomatization is about Σ_{BA}-algebras, not Σ-algebras; hence, we cannot use *variables* to refer to Σ-algebras, so we have to use a more roundabout method. The only way in which we *do* have access to information about V is through the underlying partial algebra of a network, and this is precisely what the first conjunct of $\psi_0^r(\mathcal{N})$ is about: the formula $\pi^r(\mathcal{N})$ is simply set to be *false* if the underlying partial algebra of $\mathcal{N}$ does not satisfy the first r equations defining V. This formula should be seen in the context of the ψ-formulas: for instance, if to some type (γ) move $(\nabla, k_1, \ldots, k_q)$ of $\forall$, none of the responses of the second player in a game of length 1 over a network $\mathcal{N}^\theta$ led to a network based on a partial V-algebra of grade 1, then the formula $\gamma_1^1(\mathcal{N})$ would evaluate to be *false* under the assignment θ, simply because each of the disjuncts (for $m \in N^*$)

$$\psi_0^1(\mathcal{N}(\nabla, \overline{k} : m)) \quad = \quad \pi^1(\mathcal{N}(\nabla, \overline{k} \mapsto m)) \wedge \chi(\mathcal{N}(\nabla, \overline{k} \mapsto m))$$

of $\gamma_1^1(\mathcal{N})$ would have $\perp$ as its first conjunct.

We can now prove the main technical result of this section.

Proposition 5.4 *Let $\mathcal{A}$ be a Boolean algebra with Σ-operators. Then $\mathcal{A} \models \Phi(\mathsf{V})$ if and only if $\exists$ has a winning strategy in the game $G_n(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$ for every $n < \omega$.*

Proof. This proposition is proved via a series of claims. Let $\mathcal{A}$ be a fixed Boolean algebra with Σ-operators.

Claim 1 For any term network $\mathcal{N}$, any assignment θ of the variables in the terms of $\mathcal{N}$ into $\mathcal{A}$, and any $r < \omega$, we have $(\mathcal{A}, \theta) \models \psi_0^r(\mathcal{N})$ if and only if $\mathcal{N}^\theta$ is a coherent V-network of grade r.

Proof of Claim Fix a term network $\mathcal{N} = (N, I, \tau)$ and an assignment θ. The truth of $\pi^r(\mathcal{N})$ in $(\mathcal{A}, \theta)$ corresponds to (N, I) being a partial V-algebra of grade r, while $\chi(\mathcal{N})$ takes care of coherence. ◀

Claim 2 For any term network $\mathcal{N}$, any assignment $\theta : \mathrm{Var}(\mathcal{N}) \to \mathcal{A}$, and any pair of natural numbers i and r, $(\mathcal{A}, \theta) \models \psi_i^r(\mathcal{N})$ if and only if $\exists$ has a winning strategy for the game $G_i(\mathcal{N}, \mathcal{A}, \mathsf{V})$ and moreover can arrange that the final network is comprised of a partial V-algebra of grade r.

Proof of Claim Simply by construction of the formulas $\psi_i^r(\mathcal{N})$, using induction on i. Cf. the foregoing remarks. ◀

Claim 3 For any natural number i, $\mathcal{A} \models \varphi_i$ if and only if $\exists$ has a winning strategy for the game $G_i(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$.

Proof of Claim Immediate by the previous claim and the fact that the empty term network does not contain any variables. ◀

Finally, the Proposition follows from Claim 3 and the definition of $\Phi(\mathsf{V})$. QED

As we mentioned in the introduction to this section, we have now gathered sufficient material to prove the main result of the paper.

Theorem 5.5 *Let $\mathcal{A}$ be a Boolean algebra with Σ-operators, and let V be a variety of Σ-algebras. Then the following are equivalent:*

1. *$\mathcal{A}$ belongs to the class SCmV,*
2. *$\exists$ has a winning strategy in the game $G_\omega(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$,*
3. *$\exists$ has a winning strategy in the game $G_n(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$, for any $n < \omega$,*
4. *$\mathcal{A} \models \Phi(\mathsf{V})$.*

Proof. The implications $1 \Rightarrow 2$, $2 \Rightarrow 3$ and $3 \Rightarrow 4$ follow from Proposition 4.2, Theorem 5.1 and Proposition 5.4, respectively.

For the other implication $4 \Rightarrow 1$, we first consider the countable case. Hence, let $\mathcal{A}$ be a countable Boolean algebra with Σ-operators and assume that $\mathcal{A} \models \Phi(\mathsf{V})$. It follows from Proposition 5.4 and Theorem 5.1 that $\exists$ can win the game $G_\omega(\mathcal{N}_\varnothing, \mathcal{A}, \mathsf{V})$. As $\mathcal{A}$ is countable, Proposition 4.3 yields that $\mathcal{A}$ must be representable over V.

Now suppose that $\mathcal{A}$ is an *arbitrary* (not necessarily countable) Boolean algebra with Σ-operators satisfying $\Phi(\mathsf{V})$. Since the φ_i are universal sentences, $\Phi(\mathsf{V})$ is valid in every subalgebra of $\mathcal{A}$. But then it follows by the previous argument that every countable subalgebra of $\mathcal{A}$ belongs to SCmV. Now SCmV is an elementary class by Theorem 2.1, and Σ_{BA} is countable, so a simple Löwenheim–Skolem argument shows that $\mathcal{A}$ itself must be representable over V as well. QED

Remark 5.6 A completely analogous result can be obtained for the class SPCmV by considering a version of the game where $\forall$ is only allowed to make a type δ move in the first round of the game.

References

[1] C Brink. Power structures. *Algebra Universalis*, **30** (1993) pp. 177–216.

[2] C C Chang and H J Keisler. *Model Theory.* North-Holland, Amsterdam, 3rd edition, 1990.

[3] N D Gautam. The validity of equations of complex algebras. *Archive der Mathematische Logik und Grundlagen der Mathematik*, **3** (1957) pp. 117-124.

[4] R Goldblatt. Varieties of complex algebras. *Annals of Pure and Applied Logic*, **44** (1989) pp. 173-242.

[5] V Goranko and D Vakarelov. Modal logic and universal algebra. Part 1: Modal axiomatizations of structures. In K Segerberg, M Zakharyaschev, M de Rijke and H Wansing (editors). *Advances in Modal Logic, Vol. 2*, Lecture Notes, CSLI Publications, Stanford, to appear.

[6] R Hirsch and I Hodkinson. Step by step — building representations in algebraic logic. *Journal of Symbolic Logic*, **62** (1997) pp. 225–279.

[7] R Hirsch and I Hodkinson. Relation algebras from cylindric algebras, I. Submitted, 1999.

[8] R Hirsch and I Hodkinson. Relation algebras with n-dimensional relational bases. *Annals of Pure and Applied Logic*, to appear.

[9] R Hirsch and I Hodkinson. *Relation algebras by games*, in preparation.

[10] W Hodges. Games in logic. In P Dekker, M Stokhof, and Y Venema, editors, *Proc. 11th Amsterdam Colloquium*, Amsterdam, December 1997. ILLC/Department of Philosophy, Universiteit van Amsterdam, pp. 13–18.

[11] P Jipsen. Discriminator varieties of Boolean algebras with residuated operators. In C Rauszer (editor). *Algebraic methods in Logic and Computer Science*, Banach Center Publications, Volume 28. Polish Academy of Sciences, Warszawa, 1993, pp. 239–252.

[12] R Lyndon. The representation of relational algebras. *Annals of Mathematics*, **51** (1950) pp. 707–729.

[13] R McKenzie. Representations of integral relation algebras. *Michigan Mathematical Journal*, **17** (1970) pp. 279–287.

[14] V Stebletsova and Y Venema. Q-algebras. *Algebra Universalis*, **40** (1998) pp. 19–49.

[15] A Tarski. Contributions to the theory of models. III. *Indagationes Mathematicae*, **17** (1955), pp. 56–64.

[16] Y Venema. Rectangular games. *Journal of Symbolic Logic*, **63** (1998) pp. 1549–1564.

[17] Y Venema. *Group Relation Algebras.* Talk given at the Algebraic Logic Extravaganza, Institute of Logic, Language and Computation, University of Amsterdam, april 1998.