Synthesising axioms by games

Robin Hirsch and Ian Hodkinson*

Abstract

It is an important problem in algebraic logic to find explicit axiomatisations of various classes of algebras, atom structures, etc. We will discuss some of the history of the problem, and in what way games can contribute to solving it.

Contents

*Research partially supported by UK EPSRC grants GR/K54946, GR/L85978, GR/L85961. Thanks to Ágnes Kurucz and Szabolcs Mikulás for comments.

1 Introduction

We would like to begin by hoping that you had a very happy birthday (we assume you found more enjoyable ways to spend it than reading this), and we wish you many more years of fruitful research. In this short article we would like to discuss from our current perspective the problem of providing axioms for classes of algebras, and the way in which games can contribute to solving it. Towards the end, we will describe general settings in which this can be done. We will be rather discursive and opinionated — you may well disagree with what we say (and we'd like to hear from you about it), but we are certainly not trying to provoke you or implying that you do disagree with us. We will use natural language, and mostly avoid mathematical symbols. Some (much?) of what we say will be familiar, but we include it to help tell the story. We will concentrate on the algebraic logic side, as we believe most Dutch people nowadays can easily mentally translate this into modal logic or whatever as they go along.

2 Story

We begin in the nineteenth century, when the pre-modern tradition of algebraic logic began. We find Augustus De Morgan in a restive frame of mind, wishing to extend the Aristotelian syllogistic tradition to encompass more complicated situations. His eye alighted on the slightly earlier work of George Boole, who had put forward a highly successful algebra of propositions [Boo48]. De Morgan wrote:

> When the ideas thrown out by Mr Boole shall have borne their full fruit, algebra, though only founded on ideas of number in the first instance, will appear like a sectional model of the whole form of thought. Its forms, considered apart from their matter, will be seen to contain all the forms of thought in general. The anti-mathematical logician says that it makes thought a branch of algebra, instead of algebra a branch of thought. It *makes* nothing; it *finds*: and it finds the laws of thought symbolized in the forms of algebra. [Mor60]

De Morgan's project, then, was to use an algebraic formalism, like boolean algebra, to reason about the laws of thought in general; criticising the limitations of Aristotle's syllogisms, he turned to binary (and higher-order) relations in particular. This line of research led to what we now call relation algebra: see [Mad91] for a thorough historical survey. We note particularly the insistence on algebra as the intended paradigm; later, we will discuss how far modern developments have kept to it.

By analogy with boolean algebra, what is wanted is a set of algebraic axioms, preferably equations as in the boolean case, that identify precisely the true laws of binary relations. Of course, one has to specify the desired algebraic operations on these relations — composition and converse lead to the relation algebra approach, while quantification (cylindrification) leads to first-order logic or cylindric algebras. The difference in spirit between these two approaches, though very interesting, is not really germane here; we will focus on relation algebras, but much of what we say applies more widely.

Around the beginning of the twentieth century, Peirce and Schröder worked towards identifying these true laws. Literally hundreds of true equations about binary relations were discovered, but there seemed no end in sight, and in fact Peirce commented that it appeared that one must create an ad-hoc argument to show validity of the equation in each case: 'this algebra cannot be subjected to hard and fast rules like those of the Boolian calculus; and all that can be done in this place is to give a general idea of the way of working with it' [Pei33]. This turned out to be very near the truth.

The modern stage of the subject began with Tarski and his school at Berkeley in the 1940s. Jónsson and Tarski laid down the modern axioms for relation algebra in 1948; they consist of the boolean axioms plus half a dozen additional equations governing the specifically relational operations — composition, converse, identity. Jónsson and Tarski hoped that these equations would be sufficient to characterise all true properties of binary relations. Unfortunately for them, in around 1960 Lyndon produced examples of relation algebras arising from projective planes, some of which were 'representable' as 'true' algebras of binary relations and some of which were not; Monk soon used these [Mon64] to prove that no finite set of equations or indeed first-order axioms would capture all of the true properties of binary relations. So Peirce had been far-sighted in his view.

At this point, we perceive something of a crisis in the development of the subject. The hopes of workers over a hundred years to produce a (simple, elegant) set of algebraic properties that captured exactly the true properties of binary relations (the representable relation algebras) were not to be realized. This impasse is still creating employment today.

Several different stratagems to get round the problem were evolved. One, promulgated by Tarski especially, was to find elegant sufficient conditions for representability, such as when the atoms of the given (atomic) relation algebra are functions. This approach has continued to the present, being refined by workers such as Andréka et. al [AGM+98] and Venema [Ven96, using games], looking at dense subsets of the algebra rather than atoms, and Maddux with the notion of tabulation; this was recently used in process algebras. Fork algebras are a related approach. Step-by-step constructions and occasionally games have been used by adherents to this stratagem. This is also true for another stratagem, where classes such as Crs with relativised notions of representation admit more positive results; this of course is related to *dynamic logic.*

A third stratagem, persevering in the attempt to obtain necessary and sufficient conditions for representability, is more relevant here. It brings us back to the work of Lyndon. In [Lyn50] we see an early example of a step-by-step construction, to build a representation of an (atomic) relation algebra. The step-by-step view is quite explicit in Lyndon's paper, though he does not mention games per se. Further, he described first-order axioms, which we like to call the *Lyndon conditions,* which he claimed would fully axiomatise the real binary relations (at least for atomic relation algebras). Though this claim was later seen to be too strong, as his conditions were ironically sufficient but not necessary, it is correct for finite algebras of binary relations; this is an instance of a more general phenomenon which we will discuss later. The Lyndon conditions can be written as equations, but their meaning is really in terms of the step-by-step construction. This is the aspect of Lyndon's stratagem that distinguishes it from Tarski's; the game-theoretic tools used later can be seen as mere variations on it. *The Lyndon conditions are essentially the literal translation into first-order logic of the statement that the proposed step-by-step construction of a representation can be carried through successfully.*

The question arises as to whether such a statement can be regarded as truly algebraic in nature, in the way that De Morgan had wanted. What 'form of thought' do the Lyndon conditions codify? Although this question is imprecise and cannot be given a definitive answer, by De Morgan or anyone else, it has been of some importance historically. There was a problem posed in the monograph of Henkin, Monk, and Tarski [HMT71, p461]: to find a 'simple intrinsic characterization' of the true n-ary relations. Though the question was really about cylindric algebras, we will take representable relation algebras (binary relations) to be included too. This problem seems to be regarded as still open in spite of the existence of several axiomatisations, including a corrected one by Lyndon in a later paper [Lyn56], game-theoretic ones by us [HH97c], and others. It seems that these axiomatisations are regarded as unsatisfactory in some way: they are perhaps too complicated — or too trivial, just paraphrasing the original problem. There seems to be a clash between the step-by-step approach to building representations, and something more 'algebraic'. A similar situation, perhaps, occurs in modal logic, when people say that techniques such as Gabbay's irreflexivity rule involve introducing variables 'by the back door' and are inimical to the true spirit of modal logic.

We will have more to say about this later, but we do believe that such a dispute can be unproductive. The acid test of success of a research viewpoint should be what further work it stimulates and what problems can be solved by adopting it. We should learn from nature, rather than impose our views on it. The step-by-step technique of building representations, especially when viewed as a game, is extremely potent. Not only does it allow the construction of axiomatisations for relation algebras and other kinds of algebra, but close examination of the way that games can be played on given algebras will elicit very fine and detailed information about their structure. The use of games may seem just a presentational matter but is important, because games make life so much easier and so permit us to go much further in this analysis. This is no longer controversial: the step-by-step approach, whether by games or otherwise, is now widely accepted and has been used by — hardly exaggerating — many workers, from Andréka to Zakharyaschev.

3 To the games

Let us now jump forward to a modern perspective on the use of games to provide axiomatisations. The methodology has four ingredients:

1. Building a 'representation' of an algebra or structure by a 'step by step'

construction.

2. Doing this by a game of usually infinite length played on the algebra. Following Keisler and Hodges [Kei65, Hod85, Hod97], we call the players ∀ and ∃, and we view ∃ as a Ph.D. student in the Faculty of Representability of Algebras and ∀ as examiner of her dissertation on the algebra in question. During play, ∃ tries to build a representation by a sequence of approximations; in each step or round of the game, ∀ challenges her to refine her current approximation in one way or another. ∃ wins the game if she survives all his challenges.

   We are actually interested in whether or not ∃ has a winning strategy in the game, not in whether she wins a particular play. This is because if the algebra is countable, ∃ has a winning strategy in the game if and only if the algebra is representable. (Mostly the games are *determined,* in that one of the players has a winning strategy; so assuming that ∀ is a thorough examiner, this presents no problem for Hodges' metaphor.)

3. Approximating the infinite game by finite ones. In these, ∃'s examination is curtailed after some arbitrarily large finite number of rounds. The hope is that this does not lead to a drop in standards: that if she can pass any such examination, she could pass the full exam of infinite length. For finite algebras and for those games when ∃ is never asked to choose an element of the algebra, this hope is realized, by a König tree lemma argument. But in general it is too naïve, because in a finite game, however long, ∃ may be able to prevaricate and not be caught out in the available time.[1] Still, ∃'s surviving all finite examinations does entitle her to an M.Phil. or similar qualification, with the meaning that some suitably saturated algebra elementarily equivalent to the original will be representable. Whether employers would understand such a qualification is doubtful, but it is often sufficient for the original algebra to be representable, by dint of the class of representable algebras being known by other means to be elementary.

4. Writing out first-order axioms expressing that the finite games can be won. We write a single axiom for the game of each finite length. This is where it helps to have finite-length games — an axiom for an infinite-length game would likely be infinitary. The process is not difficult but just technical (our students laugh when a lecturer says that), and we only note that occurrences of $\forall, \wedge$ in the resulting sentences come from moves that ∀ can make in the game, and $\exists, \vee$ from ∃'s responses — whence the names of the players.

[1] This is essentially the mistake in [Lyn50], and is why the Lyndon conditions are successful in capturing representability of finite relation algebras but not of arbitrary ones. The remedy in [Lyn56], seen anachronistically from the games viewpoint, involves changing the game so that ∃ never does have to choose algebra elements; this is also the case for the game doing the same job in [HH97c]. However, there is a subtlety, to do with the difference between *complete* representations that respect all meets and joins existing in the algebra, and ordinary representations, which need not. For atomic relation algebras, the Lyndon conditions are sufficient but not necessary for representability, and necessary but not sufficient for complete representability; ∃'s 'examination' is in the latter. See [HH97a].

As the history of these ideas is quite interesting, we will now detail some of the many precursors for each of these four ingredients in turn. Useful general references are [AHT65, Hod93].

1. As we said, there is now a lot of work in algebraic logic and modal and temporal logic on building a representation or model in a 'step by step' fashion. Samples include [Gab81, Bur82, Mad82, AT88, And95, VM98]. Though such arguments can usually be expressed in terms of a game, this was not done explicitly in the original. The step-by-step construction is carried out in a setting with pre-determined properties sufficient to ensure it can be concluded successfully.

2. Games to build structures are familiar in model-theoretic forcing [Hod85], and even as well-known a device as Henkin's completeness theorem for first-order logic can be seen this way. As in (1), these constructions generally take place in a setting satisfying certain properties carefully chosen in advance, and the game is used at the meta-level to build the required structure. The game is not usually brought into the object level to characterise what properties are needed, which is what we essentially do when writing axioms expressing a winning strategy. For example, [Ven96] used games to prove representability of any diagonal-free cylindric algebra that has the well-chosen property of *rectangular density;* and Henkin's construction to prove that a *consistent* first-order theory has a model is not usually thought of as a test for consistency of the given theory, though it could be used that way.

3. We cannot think of any precedents for approximating infinite games by finite ones in isolation, but [Kei65] does it in combination with (4) below.

4. Expressing a winning strategy in a game by logical sentences has a long history, notably exemplified in the famous Ehrenfeucht–Fraïssé characterisation of equivalence in first-order logic of a given quantifier depth [Ehr61] and the related theorem of Karp [Kar65] for fragments of $L_{\infty\omega}$. These games (perhaps disguised as bisimulations or back-and-forth systems) do not usually involve constructing anything, but rather, comparing two existing structures. For corresponding games on a single structure, see [Hin73], building on earlier work of Henkin (1971), and in an infinitary context, [Kei65, Sco65]. The finite model theorists have pursued these lines recently, especially using 'k-pebble games' [Bar77, IK87, DLW95].

As we said, combinations and blends of the ingredients (3) and (4) also occur in the literature. Your own theorem showing that two modal models satisfying the same modal formulas have bisimilar elementary extensions [vB76, vB85] seems to fit here (we wonder if you see it this way), as does work of de Rijke [dR95a, dR95b]. The very striking [Kei65] should certainly be mentioned too; there was a great deal of interest in games and infinitary logic in the 1960s and 70s, leading (e.g.) to work in admissible set theory. The construction aspect of (1,2) is generally missing in these examples; it is present in Lyndon's papers [Lyn50, Lyn56], which, although they don't explicitly use games, are notable (among other things) for combining (1,3,4).

4 Generalisations

Thus far, we have tried to get a feel for how games can be *used* to determine representability of algebras. We have not described how to *devise* a game suitable for axiomatising a given class of algebras. In fact, defining new games *ad hoc* for each kind of algebra and representation encountered can be misleading and tiring to do. It still treats games at the meta-level, whereas the spirit of Lyndon's stratagem seems to be to move them into the object level. Can we do this and take further advantage of it by developing a more general meta-style reasoning about games?

Some more general results in this vein exist. In [Kei65], a whole class of games is defined, one for each sentence of a certain infinitary logic. But these games do not construct objects. Closer to algebraic logic is [HMV99], where the class of complex algebras over an arbitrary variety is axiomatised. Here, the 'representation' is a certain expansion of a structure in the given variety, and the game is defined by the variety. The method covers many classes of algebras from algebraic logic. We also note that [PKS99], while not concerned with axiomatising it, is attempting to capture the notion of representability (of PDL) by a meta-language — in this case, situation calculus.

Here, we will discuss an approach which can be seen as a blend of [HMV99, PKS99], though it did not come about that way. If we want to use games to build representations, we should consider what a representation is. It turns out that being a representation of an algebra is usually easily definable by a suitable first-order theory in a 2-sorted language. The first sort of a model of this *defining theory* will be the algebra itself, and the second sort will be a representation of it. The defining theory specifies the relationship between the two, and its axioms depend on what kind of representation we are considering. Thus, the representable algebras are simply those that are the first sort of some model of the defining theory.

Not all kinds of representation are so easily written, but — as a practical observation — many of them are. Let us see some examples. The class of representable relation algebras is one. To give the idea, the parts of the defining theory expressing that representations respect boolean join and relative product of binary relations will read as follows:

$$\begin{aligned}\forall abxy(\texttt{holds}(a+b,x,y) &\leftrightarrow \texttt{holds}(a,x,y)\vee\texttt{holds}(b,x,y))\\ \forall abxy(\texttt{holds}(a\,;b,x,y) &\leftrightarrow \exists z(\texttt{holds}(a,x,z)\wedge\texttt{holds}(b,z,y))),\end{aligned}$$

where `holds` is a ternary predicate with first argument of the 'algebra' sort and the others of the 'representation' sort, and $\texttt{holds}(a,x,y)$ is intended to mean that the element a of the algebra is a binary relation that holds on the arrow (x,y), for elements x,y of the representation. The same formulas occur in [PKS99].

Other examples of classes expressible this way include:

1. The class of completely representable relation algebras — those relation algebras with a representation respecting arbitrary meets and joins.
2. The class of relation algebras with infinite representations.

3. The class of group relation algebras, and its closure under products.
4. The class of relation algebras with a permutational representation (with automorphism group acting transitively on isomorphism types of points) and those with a homogeneous representation (the automorphism group is transitive on n-tuples of a given isomorphism type, for all finite n).
5. The class $\mathbf{S}\mathsf{RaCA}_\alpha$ of subalgebras of relation algebra reducts of α-dimensional cylindric algebras (for any ordinal $\alpha \geq 3$).
6. The class RA_n of [Mad89], for finite n.
7. The class of atom structures of representable relation algebras.
8. The class of atom structures of atomic algebras in V, where V is any variety of boolean algebras with operators.
9. RCA_n (for any fixed finite n), Crs_n, D_n, G_n, and $\mathbf{S}\mathsf{Nr}_m\mathsf{CA}_n$ for finite $m \leq n$. Polyadic algebras, Jónsson Q-algebras, sequential algebras, and others can also be handled.

Difficult cases include the Kleene star operator from PDL, whose definition is not self-evidently first-order. One way to cope with this is to add a third sort consisting of a model of a suitable fragment of Peano arithmetic, and try to use it to define the star. This approach is taken for a different purpose in [PKS99]; how successful it might be in the current context is a matter for further research. Another is the class of atom structures of representable relation algebras that are full complex algebras. One might make some progress on this one by adding a third sort consisting of a model of ZFC. The method of [Hod85, theorem 7.3.1] may be relevant.

The class of all structures that arise as the first sort of a model of a fixed two-sorted first-order theory is of course a venerable old notion in model theory, introduced by Mal'tsev in the 1940s and since studied by Makkai, Svenonius, Tarski, and others. It is often used in algebraic logic to show classes to be varieties (see the recent survey by Andréka–Németi–Sain). It is known as a *pseudo-elementary class*.[2] We have seen that many classes of representable algebras from algebraic logic can be expressed as pseudo-elementary classes. The defining theory is usually finite and very natural and simple, and essentially always recursively enumerable, because we certainly expect that a Turing machine should be able to write down what we mean by a representation. In the light of the examples above, we would go so far as to say that a (fairly but not completely general) *definition* of the notion of representation of an algebra is just the second sort of a model of some two-sorted (perhaps r.e.) first-order theory, the first sort of the model being the algebra.

Now model-theoretic forcing, as seen for example in [Hod85] and indeed in the classical construction of Henkin to show completeness of first-order logic, typically involves constructing a model of some first-order theory by a game.

[2] We mean a PC'_Δ class, but expressed by two-sorted logic. The term 'pseudo-elementary class' strictly means PC_Δ, where the second sort is empty, but the two notions were proved equivalent by Makkai. Any elementary class is pseudo-elementary, but not vice versa. See [Hod93, chapters 5, 6] for more information about these classes.

The game builds the model in just the same way as other games build a representation of an algebra, step-by-step, elements of the model being introduced by the second player $\exists$ in response to criticisms by $\forall$. (In forcing, during the game itself the elements of the model are treated syntactically, e.g. as Henkin witnesses; only at the end of play do they become semantic as elements of the actual model. Here, this is only of notational significance and we don't worry about it further.)

This suggests that we combine forcing-games with the pseudo-elementary approach to representations, by *using a game to build the second sort of a model of the defining theory whose first sort is already fixed to be the algebra whose representability is at issue.* We take the defining theory of the pseudo-elementary class to be given — this defines the notion of representation to be axiomatised. The game is just as in many forcing arguments. We apply the process described earlier, approximating the infinite-length game by finite ones and writing down axioms expressing that $\exists$ has a winning strategy in them. These axioms can be obtained recursively from the defining theory of the class.

In this way, we obtain:

THEOREM 1 *Let $\mathcal{K}$ be a pseudo-elementary class of algebras in some recursive signature L, $\mathcal{K}$ being defined by a two-sorted first-order theory S in a recursive language extending L.*

1. *There is a set T of first-order L-sentences that axiomatises the elementary closure of $\mathcal{K}$.*
2. *Assume that $\mathcal{K}$ is closed under subalgebras. Then there is a set U of universal first-order L-sentences that axiomatises $\mathcal{K}$.*

T and U can be obtained effectively from S in a uniform way. If S is recursively enumerable then so are T and U. In the context of discriminator varieties, U can be effectively transformed into equations.

This theorem generalises all our own results using games to obtain axiomatisations, and those axiomatisations are similar to ones obtained by the theorem. It is a pity we did not find it earlier. We thank Szabolcs Mikulás and Yde Venema for stimulating discussions that helped to get there.

Before we leave this theorem, we stress that we can hardly claim much originality for it. As we have seen, the ideas behind it are old — Keisler in 1965 would not have been surprised by it. Indeed, the theorem itself is hardly new: here is a result which is implicit in [Mal71, from 1941] and explicit in [Tar55].

THEOREM 2 *Let $\mathcal{K}, L, S$ be as in theorem 1. Assume that $\mathcal{K}$ is closed under subalgebras. Then $\mathcal{K}$ is axiomatised by a universal L-theory U.*

For a proof (a gentle exercise in compactness) see [Hod93, theorem 6.6.7]. The proof can be tweaked to show that U can be taken to be recursively enumerable, if S is. By Craig's trick, for any (universal) r.e. theory there is an equivalent (universal) recursive theory. So the above result also proves theorem 1(2), and an adaptation of it will prove part 1 too. But the axiomatisation U that it gives is rather obscure. A similar comment can be made about the transformation in [Sve65] of the defining theory of a pseudo-elementary class into what we might anachronistically call game-normal form, equivalent to the original for finite and countable structures. And as many algebraic logic classes are known anyway to be varieties, a theorem giving an obscure r.e. axiomatisation is not telling us all that much new. As Dov Gabbay once said in a similar situation [GHR94, p. 201]: 'We want to see some axioms and rules'. So we prefer the game-theoretic proof, which does allow us to write down the axioms.

5 Discussion

Finally, we return to De Morgan's original desire to find algebraic forms of the laws of thought, and to the (we think) related problem of [HMT71] to find 'simple, intrinsic' characterizations for various classes of algebras in algebraic logic. Theorem 1 can be used to provide explicit axioms for a wide range of these classes, 'automatically' and directly from their definition. On the basis that a general procedure is worth two *ad hoc* ones, there is at least something to be said for this. The axioms do presumably constitute an intrinsic characterization in the sense that they are evaluated solely in the algebra. But are they simple, and consonant with the laws of thought? These are in the end questions for you and the community, not us, but nonetheless we will end with some remarks on this matter.

FIrst, simplicity. Well, one person's complications are another's non-trivial theorems. Substantial effort has gone into proving negative results about what kinds of axiomatisation are possible for certain classes. For example, no finite axiomatisation exists for the varieties RRA, RCA_n of representable relation algebras and cylindric algebras, for $3 \leq n < \omega$ [Mon64, Mon69]; they cannot be axiomatised by equations using only k variables, for any finite k [Jón91, And97], nor by Sahlqvist equations [Ven97]. The lesson is presumably that *any* axiomatisation is going to be complicated. The axioms obtained via theorem 1 can be written fairly simply, by an inductive definition. If we eliminate this, to really 'see' the axioms, they do become much more complicated in appearance. But their *meaning* is always clear — it comes from the games.

Let us now consider a point nearer De Morgan's hopes. A sceptic might argue that the axioms produced by theorem 1 do not give much 'algebraic insight', adding up to little more than the original statement that the algebra is representable, and that inasmuch as they are 'really' speaking about representations, they are not really intrinsic at all. In contrast, elegant finite axiomatisations of the classes $\mathsf{D}_n, \mathsf{G}_n$ are known [AT88, And95], and a similar situation applies to the class WA of weakly associative algebras [Mad82]. The axiomatisations obtained for these classes by theorem 1 get nowhere near that ideal.

Three responses can be made to this. First, it is true that axioms expressing winning strategies in games arise very directly from the pseudo-elementary definition of the class being axiomatised. But this is not the end of the matter. RRA, for example, can be defined as a pseudo-elementary class in more than one way: by letting the second sort be a representation, as already explained, or alternatively, by dint of work of Monk, by letting it be an ω-dimensional cylindric algebra of which the given relation algebra is a 'relation algebra reduct'. Theorem 1 will synthesise axioms for RRA by either of these characterisations, and others. So we have a question for our sceptic: which of these axiomatisations has the *least* algebraic insight? The difference between them is simply because of the differing characterisations of the class. Plausibly, an axiomatisation of a pseudo-elementary class obtained via theorem 1 carries about the same insight as the defining theory of the class, and that is where attention should be directed.

Second, we admit that theorem 1 is not of itself much help in telling whether a given class is finitely axiomatisable, or providing a finite set of axioms if it is. Of course, many (most?) algebraic logic classes are not finitely axiomatisable, and here game-theoretic axiomatisations come into their own. But we can surely agree that for many classes, winning strategies in games will serve to build 'representations' of algebras. In the cases of $\mathsf{D}_n, \mathsf{G}_n$, and WA, algebraic skill shows that the conditions for such strategies to exist are finitely axiomatisable. For WA, for example, the characterising game is defined in such a way that a winning strategy for $\exists$ in the game curtailed to two rounds — a finitely axiomatisable property — is already enough to guarantee a winning strategy for her in the full, infinite game. But from then on, the proof of completeness of the axioms is similar to that of theorem 1 — in each case it was presented as a 'step by step' argument and this can be recast as a game.

So do we then admit that our axioms impart no 'algebraic insight'? This brings us to our third point. The existence of a winning strategy in a game can often be expressed in a different way, by inventing some 'algebraic device' to represent it. The simplest examples, hardly a shift at all, are bisimulations and back-and-forth systems; but there are more profound transformations than this. We have in mind Fraïssé's characterisation of homogeneity by amalgamation, the related cylindric and relational bases of Maddux, the IRR theories of Gabbay [GHR94], and the 'mosaic' method of Németi. If such a device can be found, it can itself be used as the second sort in an alternative pseudo-elementary definition of the class in question, to produce new axioms for it. But further, if the device is finite, it can be used in other ways. Examples include the decidability and complexity results of [VM98]. In the appendix [HHM+98] to that paper, the finite model property was obtained by amalgamating copies of mosaics indexed by a finite group; this approach will prove the finite model property for relativised arrow logic and (a result of Grädel) the guarded fragment. There is some 'algebraic insight' in this, though it is at some remove from the games we started with. But games themselves yield their own kind of insight. With colleagues, we have found that close attention to the way that games can be played on algebras can shed quite considerable light on the nature of the algebras [HH97a, HH98, HH97b, HH99, HM98].

Would De Morgan would have been satisfied with games and the axioms they produce, as a contribution to understanding the laws of thought? Well, he did not know of Monk's negative result. We are wiser because of it, but not sadder — after all, thought has an object, and it can only be right that real-world considerations creep into our attempts to ascertain the true laws of thought. It is surely a positive move to face up to the ability of algebras to encapsulate information about their representations, and to seek algebraic content on the representation side as well as in the algebras themselves. Games are one of the more powerful and pleasant ways of doing this. As you have hinted in your recent lecture notes with Marc Pauly, perhaps games, being so much part of our lives, do themselves encapsulate part of the laws of thought. Maybe our futile speculations as to how happy De Morgan would be should await the discovery of how true this is.

References

[AGM+98] H Andréka, S Givant, Sz Mikulás, I Németi, and A Simon. Notions of density that imply representability in algebraic logic. *Ann. Pure Appl. Logic*, 91:93–190, 1998.

[AHT65] J W Addison, L Henkin, and A Tarski, editors. *The Theory of Models*. Proc. 1963 Internat. Symposium at Berkeley. North-Holland, Amsterdam, 1965.

[AMN91] H Andréka, J D Monk, and I Németi, editors. *Algebraic Logic*. Number 54 in Colloq. Math. Soc. J. Bolyai. North-Holland, Amsterdam, 1991.

[And95] H Andréka. A finite axiomatization of locally square cylindric-relativized set algebras. Preprint, Math. Inst. Hung. Acad. Sci., 1995.

[And97] H Andréka. Complexity of equations valid in algebras of relations, Part I: Strong non-finitizability. *Annals of Pure and Applied Logic*, 89:149–209, 1997.

[AT88] H Andréka and R J Thompson. A Stone type representation theorem for algebras of relations of higher rank. *Trans. Amer. Math. Soc.*, 309(2):671–682, 1988.

[Bar77] J Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296, 1977.

[Boo48] G Boole. *The mathematical analysis of logic, being an essay towards a calculus of deductive reasoning*. Basil Blackwell, 1948.

[Bur82] J P Burgess. Axioms for tense logic I: “since” and “until”. *Notre Dame J Formal Logic*, 23(2):367–374, 1982.

[DLW95] A Dawar, S Lindell, and S Weinstein. Infinitary logic and inductive definability over finite structures. *Information and Computation*, 119:160–175, 1995.

[dR95a] M de Rijke. A Lindström theorem for modal logic. In A. Ponse, M. de Rijke, and Y. Venema, editors, *Modal logic and process algebra*, volume 53 of *Lecture notes*, pages 217–230. CSLI, Stanford, CA, 1995. ISBN 1-881526-95-X.

[dR95b] M de Rijke. Modal model theory. Technical report, CWI, Amsterdam, 1995.

[Ehr61] A. Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fund. Math.*, 49:128–141, 1961.

[Gab81] D M Gabbay. An irreflexivity lemma with applications to axiomatizations of conditions on tense frames. In U Monnich, editor, *Aspects of Philosophical Logic*, pages 67–89. Reidel, Dordrecht, 1981.

[GHR94] D Gabbay, I Hodkinson, and M Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects*, volume 1. Clarendon Press, Oxford, 1994.

[HH97a] R Hirsch and I Hodkinson. Complete representations in algebraic logic. *Journal of Symbolic Logic*, 62(3):816–847, 1997.

[HH97b] R Hirsch and I Hodkinson. Relation algebras with n-dimensional relational bases. 1997. Annals Pure Appl. Logic, to appear.

[HH97c] R Hirsch and I Hodkinson. Step by step — building representations in algebraic logic. *Journal of Symbolic Logic*, 62(1):225–279, 1997.

[HH98] R Hirsch and I Hodkinson. Representability is not decidable for finite relation algebras. *Trans. Amer. Math. Soc.*, 1998. To appear.

[HH99] R Hirsch and I Hodkinson. Relation algebras from cylindric algebras, II. Submitted, 1999.

[HHM+98] R Hirsch, I Hodkinson, M Marx, Sz Mikulás, and M Reynolds. Mosaics and step-by-step. In Orłowska [Orł98], pages 158–167. Appendix to 'A Modal Logic of Relations' by Y Venema and M Marx.

[Hin73] J Hintikka. Quantifiers vs. quantification theory. *Dialectica*, 27:329–358, 1973.

[HM98] I Hodkinson and Sz Mikulás. Nonfinite axiomatisability of reducts of relation and cylindric algebras. *Algebra Universalis*, 1998. Submitted. ILLC prepublication series, ML-1997-04.

[HMT71] L Henkin, J D Monk, and A Tarski. *Cylindric Algebras Part I.* North-Holland, 1971.

[HMV99] I Hodkinson, Sz Mikulás, and Y Venema. Axiomatizing complex algebras by games. In preparation, 1999.

[Hod85] W Hodges. *Building Models by Games*. Number 2 in London Mathematical Society Student Texts. Cambridge University Press, 1985.

[Hod93] W Hodges. *Model theory*, volume 42 of *Encylopedia of mathematics and its applications*. Cambridge University Press, 1993.

[Hod97] W Hodges. Games in logic. In P Dekker, M Stokhof, and Y Venema, editors, *Proc. 11th Amsterdam Colloquium*, pages 13–18, Amsterdam, December 1997. ILLC/Department of Philosophy, Universiteit van Amsterdam.

[IK87] N Immerman and D Kozen. Definability with bounded number of bound variables. In *LICS87, Proceedings of the Symposium on Logic in Computer Science, Ithaca, New York*, pages 236–244, Washington, 1987. Computer Society Press.

[Jón91] B Jónsson. The theory of binary relations. In Andréka et al. [AMN91], pages 245–292.

[Kar65] Carol Karp. Finite-quantifier equivalence. In Addison et al. [AHT65], pages 407–412.

[Kei65] H J Keisler. Finite approximations of infintely long formulas. In Addison et al. [AHT65], pages 158–169.

[Lyn50] R Lyndon. The representation of relational algebras. *Annals of Mathematics*, 51(3):707–729, 1950.

[Lyn56] R Lyndon. The representation of relation algebras, II. *Annals of Mathematics*, 63(2):294–307, 1956.

[Mad82] R Maddux. Some varieties containing relation algebras. *Transactions of the AMS*, 272(2):501–526, 1982.

[Mad89] R D Maddux. Non-finite axiomatizability results for cylindric and relation algebras. *Journal of Symbolic Logic*, 54(3):951–974, 1989.

[Mad91] R Maddux. The origin of relation algebras in the development and axiomatization of the calculus of relations. *Studia Logica*, 3/4:421–456, 1991.

[Mal71] A I Mal'tsev. A general method for obtaining local theorems in group theory. In *The metamathematics of algebraic systems: Collected papers 1936–1967*, pages 15–21. North-Holland, 1971. B F Wells III, translator.

[Mon64] J D Monk. On representable relation algebras. *Michigan Mathematics Journal*, 11:207–210, 1964.

[Mon69] J D Monk. Nonfinitizability of classes of representable cylindric algebras. *Journal of Symbolic Logic*, 34:331–343, 1969.

[Mor60] A De Morgan. On the syllogism, no. iv, and on the logic of relations. *Trans. Cambr. Philos. Soc.*, 10:331–358, 1860. Republished in Rare masterpieces of philosophy and science, Routledge & Kegan Paul, 1966.

[Orł98] E. Orłowska, editor. *Logic at Work*, volume 24 of *Studies in Fuzziness and Soft Computing*. Physica-Verlag, Heidelberg/New York, 1998. Essays dedicated to the memory of Elena Rasiowa. ISBN 3-7908-1164-5.

[Pei33] C S Peirce. *Collected papers*. Harvard University Press, Cambridge, Mass., 1933. C Hartshorne and P Weiss, eds.

[PKS99] Y Permpoontanalarp, R Kowalski, and F Sadri. Relating situation calculus and dynamic logic. Preprint, 1999.

[Sco65] D Scott. Logic with denumerably long formulas and finite strings of quantifiers. In Addison et al. [AHT65], pages 329–341.

[Sve65] L Svenonius. On the denumerable models of theories with extra predicates. In Addison et al. [AHT65], pages 376–389.

[Tar55] A Tarski. Contributions to the theory of models. In *Proc. Koninkl. Nederl. Akad. van Wetensch Ser. A*, volume 57 and 58, pages 572–581 and 582–588 resp., 1954 and 1955.

[vB76] J van Benthem. *Modal correspondence theory*. PhD thesis, Mathematical Institute, University of Amsterdam, 1976.

[vB85] J F A K van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Naples, 1985.

[Ven96] Y Venema. Rectangular games. Technical Report 96-176, Mathematics Department, Victoria University, Wellington, New Zealand, 1996.

[Ven97] Y Venema. Atom structures and Sahlqvist equations. *Algebra Universalis*, 38:185–199, 1997.

[VM98] Y Venema and M Marx. A modal logic of relations. In Orłowska [Orł98], pages 124–167.