The Bulletin of Symbolic Logic
Volume 17, Number 3, Sept. 2011

REVIEWS

The Association for Symbolic Logic publishes analytical reviews of selected books and articles in the field of symbolic logic. The reviews were published in ***The Journal of Symbolic Logic*** from the founding of the Journal in 1936 until the end of 1999. The Association moved the reviews to this Bulletin, beginning in 2000.

The Reviews Section is edited by Steve Awodey (Managing Editor), John Burgess, Mark Colyvan, Anuj Dawar, Marcelo Fiore, Noam Greenberg, Rahim Moosa, Ernest Schimmerling, Carsten Schürmann, Kai Wehmeier, and Matthias Wille. Authors and publishers are requested to send, for review, copies of books to *ASL, Box 742, Vassar College, 124 Raymond Avenue, Poughkeepsie, NY 12604, USA.*

In a review, a reference "JSL XLIII 148," for example, refers either to the publication reviewed on page 148 of volume 43 of the Journal, or to the review itself (which contains full bibliographical information for the reviewed publication). Analogously, a reference "BSL VII 376" refers to the review beginning on page 376 in volume 7 of this Bulletin, or to the publication there reviewed. "JSL LV 347" refers to one of the reviews or one of the publications reviewed or listed on page 347 of volume 55 of the Journal, with reliance on the context to show which one is meant. The reference "JSL LIII 318(3)" is to the third item on page 318 of volume 53 of the Journal, that is, to van Heijenoort's *Frege and vagueness*, and "JSL LX 684(8)" refers to the eighth item on page 684 of volume 60 of the Journal, that is, to Tarski's *Truth and proof*.

References such as 49***5*** or 280***1*** are to entries so numbered in *A bibliography of symbolic logic* (the Journal, vol. 1, pp. 121–218).

Stephen Cook and Phuong Nguyen. ***Logical foundations of proof complexity.*** Perspectives in Logic. Cambridge University Press, New York, 2010, 15 + 479 pp.

It is no secret that computational complexity theory has its origins deeply rooted in mathematical logic. By this we mean not only that the original definitions and early results were inspired by concepts from classical computability theory, but also that some of the most influential results in the area have a deep logical meaning. For example, one of the important early results in the area, the Cook–Levin Theorem, states that the satisfiability problem for propositional logic is **NP**-complete or, dually, that the problem of detecting propositional tautologies is **co-NP**-complete. One immediate consequence of this is that, unless **NP** = **co-NP**, the standard textbook proof systems for propositional logic are not *polynomially bounded*: although they are complete in the sense that all tautologies have proofs, they are not efficiently so in the sense that not all tautologies have short (i.e., polynomial-size) proofs. Perhaps surprisingly, this straight consequence of the hypothesis that **NP** ≠ **co-NP** is not known to hold unconditionally except for some weak proof systems. This takes us to an active field of research called propositional proof complexity, one of whose aims is to classify the relative strength of the existing propositional proof systems, with the goal of understanding and exploiting this sort of incompleteness phenomenon at the level of propositional logic.

The subject of ***Logical foundations of proof complexity*** revolves precisely around the theme of propositional proof complexity. It does so by developing a general framework by which

1079-8986/11/1703-0006/$1.90

a given complexity class $\mathcal{C}$ gets associated both a propositional proof system $\mathcal{P}$ and a corresponding weak theory of arithmetic $\mathcal{T}$. The definable functions in $\mathcal{T}$ are precisely those computable in $\mathcal{C}$. The bounded theorems of $\mathcal{T}$ translate into tautologies with polynomial-size proofs in $\mathcal{P}$. And a sort of converse to this last statement holds too since the theory $\mathcal{T}$ proves the soundness of $\mathcal{P}$. Thus, in the language of the previous paragraph, the proof system $\mathcal{P}$ is not only complete, but efficiently so, with respect to the propositional translations of bounded theorems in $\mathcal{T}$.

At this point an example will help and will also illustrate one of the constructions in the book. Let **PH** be the class of languages in the polynomial-time hierarchy. In logic terms, these are the sets of binary strings that are described by a bounded first-order formula in a two-sorted language, with one sort for natural numbers and one sort for binary strings indexed by these numbers, equipped with the basic arithmetic language 0, 1, $+$, $\times$ and $\leq$ on numbers, the length function $|\cdot|$ on strings, and the bit indicator relationship $\in$ between indices and strings.

The propositional proof system associated to **PH** is called **G**. This is an extension of the standard propositional sequent calculus to deal with quantified Boolean formulas. To be precise, **G** adds the following two rules to the usual propositional rules for the sequent calculus:

$$\frac{A(p), \Gamma \longrightarrow \Delta}{(\exists x)(A(x)), \Gamma \longrightarrow \Delta} \qquad \frac{\Gamma \longrightarrow \Delta, A(B)}{\Gamma \longrightarrow \Delta, (\exists x)(A(x))},$$

where in the first rule p is a free variable that does not appear in the lower sequent, and in the second $A(B)$ is the result of replacing each free occurrence of x in A by B.

The weak theory of arithmetic associated to **PH** is called $\mathbf{V}^{\infty}$. This is a two-sorted theory in the first-order language described above, with basic axioms for 0, 1, $+$, $\times$, $\leq$, $|\cdot|$, and $\in$, and the comprehension axiom for each bounded first-order formula. More precisely, if $\phi(x)$ is a bounded first-order formula, perhaps with additional free variables, the comprehension axiom for ϕ states that, for every length y, the binary string of length y that has ones precisely at the positions $z < y$ for which $\phi(z)$ holds, exists. In symbols:

$$(\exists X, |X| \leq y)(\forall z < y)(z \in X \leftrightarrow \phi(z)).$$

Alternatively, in this case one may replace the comprehension axiom scheme by the axiom scheme of induction for bounded formulas and obtain the same theory.

One of the main points of the book is that this intimate *ménage a trois* between a complexity class, its associated propositional proof system, and its associated weak theory of arithmetic can be arranged naturally and in a uniform way across all the spectrum of classical complexity classes in the neighborhood of **P**. The example we gave corresponds precisely to the upper end of this spectrum **PH**, but the book actually starts the development of the theory at the bottom end of the spectrum $\mathbf{AC}^0$. The studied complexity classes include:

$$\mathbf{AC}^0 \subseteq \mathbf{AC}^0(2) \subseteq \mathbf{ACC} \subseteq \mathbf{TC}^0 \subseteq \mathbf{NC}^1 \subseteq \mathbf{L} \subseteq \mathbf{NL} \subseteq \mathbf{NC}^2 \subseteq \mathbf{P} \subseteq \mathbf{PH}.$$

The declared motivation for developing this theory is to set up the background to start a program that the authors call "bounded reverse mathematics", in analogy with the ongoing program of reverse mathematics put forward by Harvey Friedman and others. In bounded reverse mathematics one would look for the weakest theory of bounded arithmetic in which the most important theorems that are relevant to computer science can be proved. The book includes several such examples from the recent literature in this area, such as feasible proofs of the Cayley–Hamilton Theorem of linear algebra, or a graphical version of the Jordan Curve Theorem.

The treatment of bounded arithmetic in the book is somewhat unusual in that it takes the two-sorted language with one sort for numbers and one sort for strings as the preferred language for the theory. This setup has its origins in Buss' celebrated thesis ***Bounded arithmetic***, Bibliopolis, 1986, for complexity classes beyond **PH**, and following Zambella *Notes on polynomially bounded arithmetic*, ***The Journal of Symbolic Logic***, vol. 61 (1996), no. 3, pp. 942–966, the authors adapt the method to **PH** and below. Of course one could say that the difference is merely cosmetic and that no real gain is obtained by considering two-sorted theories instead of one-sorted theories as in Peano arithmetic, but the point that the authors make is that their treatment makes the subject both closer and more accessible to mainstream computational complexity theory, where the string plays an important role as the basic representation device for inputs to Turing machines and, more importantly, to Boolean circuits. Indeed the outcome of the theory is an elegant setup with a clean language that, for example, avoids the introduction of *smash functions* and *rounded halving* from Buss' theories S_2^i and T_2^i. One nice consequence of this is that the formulas that describe the languages in the standard complexity classes become syntactically quite close to those used in the capturing results of descriptive complexity theory *à la* Immerman in ***Descriptive complexity***, Springer, 1999. As a matter of fact, the subject of the book can almost be thought as developing the proof theory that is missing from the descriptive complexity approach to understanding complexity classes through logic.

Before we close this review, a comparison to Krajíček's book ***Bounded arithmetic, propositional logic and computational complexity***, Cambridge University Press, 1995, is in order. Both books cover the topics of bounded arithmetic and propositional proof complexity through the lens of computational complexity. However, one of the declared goals of Cook and Nguyen's monograph, which is to prepare the ground for developing the program of bounded reverse mathematics, is barely touched in Krajíček's book, and in the reverse direction, one of the main focuses in Krajíček's book, which is to prove independence results for bounded arithmetic theories and super-polynomial lower bounds for concrete propositional proof systems, is barely covered in Cook and Nguyen's monograph. In this respect, the two books complement each other quite well. A piece of useful information for the student of the subject is that one consequence of this difference in focus is that the book under review makes a more accessible text for learning the basic theory from the start, yet at a good pace.

ALBERT ATSERIAS

Universitat Politècnica de Catalunya, Barcelona, Spain. atserias@lsi.upc.edu.

MANUEL LERMAN. ***A framework for priority arguments.*** Lecture Notes in Logic, vol. 34. Cambridge University Press, New York, 2010, xvi + 176 pp.

This is a book about one of the main methods in the area of computability (or recursion) theory. The *priority method* was invented by Friedberg and independently Muchnik for the solution of one of the main problems in the area at the time, namely Post's problem (whether there are computably enumerable degrees other than zero and the degree of the halting problem). Since then, the priority method has been developed considerably and is the main type of argument employed for the investigation of the structure of computably enumerable degrees. Moreover, it has found applications in other areas like computable mathematics, but it is ultimately connected to theorems that involve computably enumerable objects.

Since the Friedberg/Muchnik papers, priority arguments have become very complex and sophisticated. There have been various attempts to make them more systematic or even create a sort of framework for these arguments. For example Harrington gave his own account of a notoriously hard priority argument of Lachlan (the Lachlan non-splitting theorem) which was highly influential in understanding and presenting such proofs. This and other more systematic approaches (like one by Ash and one by Knight for constructions in computable model theory as well as unpublished work by Groszek and Slaman) have influenced the framework presented in this book.

Frameworks for priority arguments are usually needed when a theorem involves many levels (typically, infinitely many) of an underlying inductive structure. For example, the framework presented in this book is based on a framework developed by the author and Steffen Lempp, which was ultimately used in to prove the decidability of the existential theory of the poset of computably enumerable degrees in a language containing predicates representing the relations $\mathbf{a}^{(n)} \leq \mathbf{b}^{(n)}$ for all integers n. The same holds for Ash, who developed a framework in order to prove a theorem in computable model theory that referred to all levels of the hyperarithmetical hierarchy.

Apart from such cases where a framework was developed out of the need to prove a theorem, it is arguably desirable to have a uniform approach that encompasses all priority arguments. Indeed, on the one hand priority proofs tend to require lengthy presentations and on the other hand a priority proof seems to have a lot in common with other priority proofs. The point is that if we could isolate the general common principles that are used in these arguments (and are repeatedly demonstrated in each of these proofs) we could adopt a more generic approach in which proofs become shorter by making references to general 'framework theorems'. Such an approach is adopted in set theory with forcing arguments. It is fair to say, however, that in computability theory only a very small proportion of proofs are done in a formal framework. Even forcing arguments in computability theory are most often presented in a self-contained manner, without appealing to a forcing framework. It is arguable that presenting a priority proof with reference to a general framework tends to hide the intuition and the main ideas that underly the argument. It definitely separates the proof from the theorem it proves.

The book starts with a brief chapter giving the definition of the computably enumerable degrees along with notation and related notions. It also briefly discusses the notion of *sequences of trees* which plays a central role in the framework for priority arguments. Chapter 2 lays down the basics of the framework: *systems of trees of strategies*. As usual in priority arguments, we start by writing out a set of requirements which guarantee the theorem we wish to prove. Then for each requirement we introduce a basic module, which describes the way the requirement is going to be satisfied. According to the framework, basic modules are finite binary trees whose edges and non-terminal nodes are labeled with sentences. The *directing sentence* is the one that directs the action taken. The framework involves a decomposition of requirements that starts at some high level tree T^n and continues successively to lower levels until we reach the computable level T_0 at which the construction takes place. Further details are illustrated by means of examples, like the Friedberg–Muchnik theorem.

It has to be said that for someone who is familiar with the standard proofs of such basic theorems, the framework looks quite involved. It basically amounts to learning how to do something we already know, but in a different way. In any case, it seems necessary to spend considerable time learning how to prove well known theorems in this manner, before one can benefit from the framework. Moreover this presentation of priority arguments may be quite intimidating for someone starting out learning computability theory.

Various devices associated with the framework, like *links* and *blocks* are discussed before the *framework theorem* is presented at the end of Chapter 2. Various additional framework lemmas are needed for the proof of various theorems. These are stated throughout the book (in the proofs of the various theorems) but their proof is deferred until Chapters 8 and 9.

One of the criticisms about this framework (or the early versions of it) is that it needs a lot of patches and extensions in order to accommodate various theorems whose proof does not exactly fit the basic setting. This however seems unavoidable since, despite the apparent overlap in many priority proofs, new theorems often require new ideas and non-trivial modifications and devices in their proof. *No other attempt to organize priority arguments in a generic way has been carried out in such a detailed manner, covering most of the fundamental priority constructions in recursion theory.*

Chapters 3–7 present various representative priority arguments within the framework. These are organized in a familiar way, roughly speaking according to the number of quantifiers needed in order to fully describe how the underlying requirements are satisfied. Chapter 3 deals with Σ_1 constructions like the low computably enumerable non-computable degree and the properly d.c.e. degree. Such constructions are sometimes called 'finite injury of bounded type' since the number of injuries of a requirement can be bounded by a computable function on the indices of the requirements. These arguments are perhaps too simple to demonstrate advantages of the framework approach.

Chapter 4 deals with Δ_2 constructions. These are sometimes called 'finite injury of unbounded type' arguments. The Sacks splitting theorem is a typical example. Another example is the construction of a computably enumerable degree which is incomparable to a given incomplete non-computable computably enumerable degree (also due to Sacks). In these arguments there is no recursive bound on the (finite) number of injuries that a strategy may endure.

Chapter 5 deals with Π_2 constructions. These are often called *infinite injury arguments* since strategies may endure an infinite number of injuries (but in a controlled manner which ensures their satisfaction). Various typical examples of such theorems are presented within the framework, like the construction of a high incomplete computably enumerable degree and the jump inversion theorem. A special representative of this class of arguments is the minimal pair construction, which is also presented in this chapter along with the embedding of the pentagon.

Chapter 6 deals with Δ_3 constructions, which are a type of infinite injury argument. The Sacks density theorem (asserting that the structure of the computably enumerable degrees is dense) is presented in the framework. Chapter 7 deals with Σ_3 constructions. The first example of such a construction was Lachlan's non-splitting theorem which was dubbed the *monster theorem* due to its complexity. The example chosen to demonstrate this type of construction in the framework is a curious one. It is the construction of a 'strong minimal pair' or a 'super minimal pair' as it has often been called in the literature. This is a pair of degrees $\mathbf{a}, \mathbf{b}$ such that $\mathbf{a} \not\leq \mathbf{b}$ and the least upper bound of every $\mathbf{z}$ with $\mathbf{0} < \mathbf{z} < \mathbf{a}$ with $\mathbf{b}$ is $\geq \mathbf{a}$. This theorem is attributed to Slaman, but no written proof has appeared until now. It has been discussed by various researchers, including myself, and it plays a role in one of the main open problems in the computably enumerable degrees: whether the 2-quantifier theory is decidable. I have attempted to prove it myself in the past, but found it quite tricky. The account given in Chapter 7 (within the framework) is impressively short. However (having not studied the framework sufficiently) I can not get an impression as to how the author deals with various obstacles one encounters in a standard attempt to prove it.

Chapter 7 concludes with various reflective remarks on the framework. The author conjectures that any priority argument proof of properties of the computably enumerable degrees can be presented in the framework. On the other hand he admits that in certain cases (like the embedding of the 1-3-1 lattice) the proof using the framework may be more cumbersome.

Chapters 8 and 9 are devoted to the proof of the various framework lemmas that are used in the proofs that we just discussed. They are quite technical and require considerable effort to read. Finally in Chapter 10 the author discusses higher level priority arguments, which are really the type of arguments that seem to require a formal framework.

In conclusion, this is an impressive technical piece of work. It is hard to say how influential it is going to be in the future, but as of today most theorems in the computably enumerable degrees are proved without reference to a formal framework. Having said that, higher order priority arguments that refer to some underlying inductive structure (like the jump hierarchy) seem to require a more organized and formal approach, making them perfect candidates for this framework. The value of this approach will ultimately depend on the new theorems that are proved using it.

GEORGE BARMPALIAS

State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing 100190, People's Republic of China. barmpalias@gmail.com.

ILIJAS FARAH. *All automorphisms of the Calkin algebra are inner*. ***Annals of Mathematics***, vol. 173 (2010), no. 2, pp. 619–661.

Farah has produced this brilliant work between set theory and operator algebra. We have a lot to describe in little space, so let us jump right into it.

Let H be a separable infinite-dimensional complex Hilbert space, and let $\mathcal{B}(H)$ be the collection of bounded linear operators on H. A natural way to think of a linear operator A on H is as infinite dimensional matrix with entries $A_{i,j} = \langle Ae_j, e_i\rangle$ where $\{e_i\}$ is a given orthonormal basis for H. Such an A is bounded iff it is continuous, which is just to say that multiplication by its matrix is continuous. Generalizing concepts from finite linear algebra, for $A, B \in \mathcal{B}(H)$, we have the operator norm, $\|A\| = \sup\{\|Ax\| \mid x \in H$ and $\|x\| = 1\}$, scalar multiplication, $(\lambda A)x = \lambda(Ax)$, and addition and multiplication of operators, $(A+B)x = Ax+Bx$ and $(AB)x = A(Bx)$. The adjoint operator A^* is implicitly defined by the equation $\langle Ax, y\rangle = \langle x, A^*y\rangle$. In terms of matrices, $A^*_{i,j} = \overline{A_{j,i}}$. With the operator norm and these four operations, but without the Hilbert space structure, $\mathcal{B}(H)$ is the motivating example of a C^*-algebra. The *Gelfand–Naimark theorem* says that every C^*-algebra is isomorphic to one that is *concrete*, that is, a norm-closed $*$-subalgebra of $\mathcal{B}(H')$ for some Hilbert space H', which need not be separable.

A bounded linear operator A is *compact* iff A is a norm limit of finite rank operators. Define $\mathcal{K}(H) = \{A \in \mathcal{B}(H) \mid A \text{ is compact}\}$. Here is a practical description. Write P_n for the projection map from H onto the finite subspace spanned by $\{e_i\}_{i<n}$. Then, A is compact iff at least one of the following equivalent conditions holds: $\|A - P_nAP_n\| \to 0$, $\|A - AP_n\| \to 0$, or $\|A - P_nA\| \to 0$. The identity operator $I : H \to H$ is an example of an operator that is not compact.

As in finite linear algebra, A is *normal* iff $AA^* = A^*A$. Recall that every finite normal matrix has an orthonormal basis of eigenvectors. The same statement also holds for compact normal operators; it is an instance of the *spectral theorem*. Let us describe it in terms of matrices. An operator U is *unitary* iff $UU^* = U^*U = I$. Conjugation by a unitary corresponds to a change of basis. This version of the spectral theorem says that if A is normal and compact, then there is a unitary U such that UAU^* is diagonal with eigenvalues tending to zero along the diagonal. There is a more general spectral theorem that applies to normal operators that are not necessarily compact. The conclusion tells us that up to unitary equivalence, normal operators look like multiplication operators $M_f : g \mapsto fg$ on $L^2(X)$ for a σ-finite measure space X and $f \in L^\infty(X)$.

With the same norm, arithmetic and adjoint, $\mathcal{K}(H)$ also forms a C^*-algebra. One important difference between $\mathcal{B}(H)$ and $\mathcal{K}(H)$ is that the latter is not unital because I is not compact. We write $[A]$ for the equivalence class of A modulo compact operators. In other words, $[A] = \{A + C \mid C \in \mathcal{K}(H)\}$. Another way to view $\mathcal{K}(H)$ is as a nontrivial $*$-closed, norm-closed, two-sided ideal in $\mathcal{B}(H)$. In fact, there are no others given that H is separable. The *Calkin algebra* is the quotient unital C^*-algebra $\mathcal{C}(H) = \mathcal{B}(H)/\mathcal{K}(H)$.

A standard fact is that an element $[U]$ of $\mathcal{C}(H)$ is unitary iff there exists $V \in [U]$ such that V is a partial isometry between cofinite dimensional subspaces of H. This is a consequence of Atkinson's theorem; a good reference is Arveson (***A short course in spectral theory***, Graduate Texts in Mathematics 209, Springer, 2000). A question that has drawn much attention, asked by Brown, Douglas and Fillmore (*Extensions of C^*-algebras and K-homology*, ***Annals of Mathematics***, vol. 105 (1977), pp. 265–324), is whether every C^*-algebra automorphism Φ of $\mathcal{C}(H)$ is *inner* in the sense that it amounts to conjugation by a unitary of $\mathcal{C}(H)$. In other words, whether there exists a partial isometry $U \in \mathcal{B}(H)$ such that $\Phi([A]) = [UAU^*]$ for every $A \in \mathcal{B}(H)$. The opposite of inner is *outer*. Phillips and Weaver (*The Calkin algebra has outer automorphisms*, ***Duke Mathematics Journal***, vol. 139 (2007), pp. 185–202) proved the negative answer to this question assuming the Continuum Hypothesis (CH). Farah proves the positive answer assuming Todorcevic's Axiom (TA), which we define below. Combined, these theorems show that the existence question for outer automorphisms of the Calkin algebra is independent of ZFC. The first independence result in C^*-algebra is a groundbreaking accomplishment to be celebrated!

Let X be a separable metric space and consider the off-diagonal subset of the product: $[X]^2 = \{\{x, y\} \mid x, y \in X \text{ and } x \neq y\}$. Suppose K_0 is an open subset of $[X]^2$ and $K_1 = [X]^2 - K_0$. We say that $Y \subseteq X$ is *K_i-homogeneous* iff $[Y]^2 \subseteq K_i$. By definition, TA says that in every such situation, one of the following holds:

- There is an uncountable K_0-homogeneous set.
- There is a countable family of K_1-homogeneous sets whose union is X.

TA is a consistent relative to ZFC. For, if the second option fails for a given open coloring and CH holds, then there is a ccc forcing that adds an uncountable K_0-homogeneous set. Assuming that $\diamondsuit_{\omega_2}$ holds too, a finite support iteration of length ω_2 makes TA hold in a generic extension. TA had been called OCA but the name was changed to distinguish it from closely related statements.

Farah's theorem has been called a non-commutative analog of an earlier theorem by Velickovic regarding the Boolean algebra quotient of $\mathcal{P}(\omega)$ modulo the ideal of finite sets. Namely, $\mathrm{TA} + \mathrm{MA}_{\omega_1}$ implies that, for every Boolean algebra automorphism Φ of $\mathcal{P}(\omega)/\mathrm{Fin}$, there is a function e on ω such that $\Phi([A]) = [e''A]$ for every $A \subseteq \omega$. In turn, Velickovic built on a consistency proof by Shelah. Here is one way to explain how the two rigidity theorems are related. First, by Stone duality, we have a Boolean algebra isomorphism $\mathcal{P}(\omega)/\mathrm{Fin} \simeq S(\beta\omega - \omega)$ where $S(X)$ is the collection of clopen subsets of X. Second, as C^*-algebras, $\ell^\infty/c_0 \simeq C(\beta\omega - \omega)$, a fact explained in the fine survey article by Weaver (*Set theory and C^*-algebras*, ***The Bulletin of Symbolic Logic***, vol. 13 (2007), pp. 1–20). Next, observe that the characteristic function of a clopen set is continuous, so $S(\beta\omega - \omega)$ sits inside of $C(\beta\omega - \omega)$. Finally, thinking about the sequence of diagonal matrix entries of a diagonal operator, we see how ℓ^∞ sits inside $\mathcal{B}(H)$ and c_0 sits inside $\mathcal{K}(H)$. Not all operators are diagonal, or even normal, which is where the non-commutativity enters. Elements of Velickovic's proof appear in new forms in the paper under review. Other central ideas grew out of Farah (*Analytic Quotients*, ***Memoirs of the American Mathematical Society***, vol. 148 (2000), no. 702).

There is an easy decomposition lemma from §1 that gets the project started. Let $f\colon \omega \to \omega$ be a strictly increasing function with $f(0) = 0$. Then f gives rise to a partition $E = \langle E_n \mid n < \omega\rangle$ of ω into finite intervals, where $E_n = [f(n), f(n+1))$. In general, define $\mathcal{D}[E]$ to be the collection of $A \in \mathcal{B}(H)$ whose matrix is block diagonal with the blocks determined in the obvious way by E. Throwing in $E_{-1} = \emptyset$, we define overlapping coarser partitions: $E_n^{\mathrm{even}} = E_{2n} \cup E_{2n+1}$ and $E_n^{\mathrm{odd}} = E_{2n-1} \cup E_{2n}$. Now let $A \in \mathcal{B}(H)$ be arbitrary. The lemma concludes that we have a partition E, operators $A^{\mathrm{even}} \in \mathcal{D}[E^{\mathrm{even}}]$ and $A^{\mathrm{odd}} \in \mathcal{D}[E^{\mathrm{odd}}]$, and a compact operator C such that $A = A^{\mathrm{even}} + A^{\mathrm{odd}} + C$. In other words, in the Calkin algebra, $[A] = [A^{\text{even}}] + [A^{\text{odd}}]$. For this reason, the paper focuses on C^*-algebras of the form $\mathcal{D}[E]$. In particular, Farah defines a family $\mathcal{F}$ of pairs to be a *coherent family of unitaries* iff (1) in the first coordinate, $\mathcal{F}$ consists of all partitions of ω into finite intervals, and (2) for all pairs $(E, U), (F, V) \in \mathcal{F}$, U and V are partial isometries between cofinite dimensional subspaces of H, and $[UAU^*] = [VAV^*]$ for every $A \in D[E] \cap D[F]$. If $\mathcal{F}$ is a coherent family of unitaries, then there is a unique automorphism Φ of $\mathcal{C}(H)$ such that $\Phi([A]) = [UAU^*]$ for all $(E, U) \in \mathcal{F}$ and $A \in D[E]$. In §1 of the paper, Farah gives a new proof of the Phillips–Weaver theorem by constructing a coherent family of unitaries whose corresponding automorphism of $\mathcal{C}(H)$ is outer. In this proof, CH is only used to list in order type ω_1 all partial isometries U whose matrix is diagonal and, for every $n < \omega$, there is an angle $\theta_n \in [0, 2\pi)$ such that $U_{n,n} = e^{i\theta_n}$. By a C^*-algebraic lemma, for the most part, it suffices to consider coherent families of this sort of unitary operators. In §3, Farah shows that TA has the extreme opposite effect as CH. Namely, TA implies that if an automorphism of $\mathcal{C}(H)$ is determined by coherent family of unitaries, then it is inner. It is hard to come up with a meaningful slogan that summarizes how TA is applied in this part of the proof, so the reviewer will not try. This reduction is the first of three main steps towards the proof of Farah's theorem.

We mentioned that Velickovic used MA_{ω_1} in his theorem on $\mathcal{P}(\omega)/\text{Fin}$. Groundwork in §4 includes a lemma with a short and simple proof that Farah explains is the reason why he does not use MA_{ω_1} at the analogous points in §§6–7.

Recapitulating the result from §3, under TA, in order to show an automorphism Φ of $\mathcal{C}(H)$ is inner, it suffices to show it is inner on the Calkin algebra of $\mathcal{D}[E]$ for every partition E of ω into finite intervals. The second main step, handled in §6, does this under the additional assumption that Φ has a representation on $\mathcal{D}[E]$ that is C-measurable on the unit ball $\mathcal{D}[E]_{\leq 1} = \{A \in \mathcal{D}[E] \mid \|A\| \leq 1\}$ with the strong topology. The additional assumption means that there is a function φ on $\mathcal{D}[E]$ such that $\Phi([A]) = [\varphi(A)]$ for every $A \in \mathcal{D}[E]$, and φ restricted to $\mathcal{D}[E]_{\leq 1}$ is measurable with respect to the σ-algebra generated by the analytic subsets of $\mathcal{D}[E]_{\leq 1}$. In each of several intermediate steps, the author argues that there is a certain set of operators on which a higher quality representation of Φ exists. The quality improves from C-measurable, to strongly continuous, to Borel-measurable approximately $*$-homomorphic, to $*$-homomorphic, and finally to conjugation by a partial isometry. Each intermediate step involves lots of hands-on analysis and combinatorics, and one of the steps uses a deeper fact about C^*-algebra proved earlier in §5. (Set theorists acquainted with the material in the book by Arveson mentioned above have the background to read the paper if they are willing to take on faith a few intelligible results quoted from operator algebra journal articles.)

Putting together the two main steps described so far gives that, under TA, if Φ has a C-measurable representation, then Φ is inner. But it is not difficult to show that if Φ has a C^*-measurable representation, then it is possible to assert that Φ is inner in a Σ^1_2 way. Recalling that TA holds in a forcing extension and applying Shoenfield's Σ^1_2-absoluteness theorem, we see that, so far, we have not really used TA, only ZFC!

What remains is to manufacture a C-measurable representation φ of Φ for a given $\mathcal{D}[E]$. This part of Farah's proof, §7, resembles the part of Velickovic's proof that does not use MA_{ω_1}. Namely, TA implies that if Φ is an automorphism of $\mathcal{P}(\omega)/\text{Fin}$, then there is an infinite set $S \subseteq \omega$ such that Φ has a representation on $\mathcal{P}(S)/\text{Fin}$ that is covered by countably many Borel-measurable functions on $\mathcal{P}(S)$. For the end game of Farah's proof, the reviewer found it helpful to learn the corresponding part of the Boolean algebra proof and stare at the beginning of §7 until he saw how they are related. (Lectures by Farah and Velickovic, and reading the paper together with a sharp graduate student, Paul McKenney, made a big difference overall.)

The main theorem and the virtuosity the proof displays are astounding, as are the avenues for interdisciplinary research the paper expands or creates. The reviewer cannot recommend it highly enough.

ERNEST SCHIMMERLING

Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, USA. eschimme@andrew.cmu.edu.