The Cardinality of an Oracle in Blum-Shub-Smale Computation

Wesley Calvert
Murray State University
Murray, Kentucky 42071 USA
wesley.calvert@murraystate.edu

Ken Kramer   Russell Miller
Queens College of CUNY
65-30 Kissena Blvd., Flushing, NY 11367 USA
CUNY Graduate Center
365 Fifth Avenue, New York, NY 10016 USA
kkramer@qc.cuny.edu   Russell.Miller@qc.cuny.edu

We examine the relation of BSS-reducibility on subsets of $\mathbb{R}$. The question was asked recently (and anonymously) whether it is possible for the halting problem $\mathbb{H}$ in BSS-computation to be BSS-reducible to a countable set. Intuitively, it seems that a countable set ought not to contain enough information to decide membership in a reasonably complex (uncountable) set such as $\mathbb{H}$. We confirm this intuition, and prove a more general theorem linking the cardinality of the oracle set to the cardinality, in a local sense, of the set which it computes. We also mention other recent results on BSS-computation and algebraic real numbers.

1 Introduction

Blum, Shub, and Smale introduced in [2] a notion of computation with full-precision real arithmetic, in which the ordered field operations are axiomatically computable, and the computable functions are closed under the usual operations. A complete account of this model is given in [1]. A program for such a machine consists of a finite set of instructions as described there, and the instructions are allowed to contain finitely many real parameters, since a single real number is viewed as a finite object. The program can add, multiply, subtract, or divide real numbers in its cells, can copy or delete the content of a cell, and can use the relations $=$ and $<$ to compare the contents of two cells, forking according to whether the contents of those cells satisfy that relation. For our purposes, it will be convenient to assume that the forking instructions in the program compare the real number in a single given cell to 0, under either $=$ or $<$ or $>$. Such a machine has equivalent computing power to machines which can compare the contents of two different cells to each other.

Of course, the BSS model is not the only concept of computation on $\mathbb{R}$, nor should it be considered the dominant model. It corresponds to a view of the real numbers as a fixed structure, perhaps given axiomatically – defined, for instance, as the unique complete ordered field, with field operations vouchsafed unto us mathematicians; as opposed to a view of real numbers as objects defined by Cauchy sequences or by Dedekind cuts in the rational numbers $\mathbb{Q}$, with operations derived from the analogous operations on $\mathbb{Q}$. There is no obvious method of implementing BSS machines by means of digital computers. This failure invites a contrast with computable analysis, which treats real numbers as quantities approximated by rational numbers and is intended to reflect the capabilities of digital computers. However, the BSS model is of interest both for the analogy between it and the Turing model, which can be seen as BSS computation on the ring $\mathbb{Z}/(2\mathbb{Z})$, and because it reflects the intuitions of many mathematicians – dating back to the nineteenth century, and mostly outside of computer science – about the notion of algorithmic computation on $\mathbb{R}$

CCA 2010
EPTCS ??, 20??, pp. 1–11, doi:10.4204/EPTCS.??.??

This paper will consider sets of algebraic real numbers, and other sets of tuples from $\mathbb{R}$, as oracles for BSS machines, and will examine the relative difficulty of deciding membership in such sets under the BSS model of computation. We will focus in particular on questions about cardinality: to what extent the complexity of a subset of $\mathbb{R}$ allows us to draw conclusions about its cardinality. The previous paper [8] by Meer and Ziegler focused attention on these issues, and here we answer several of the questions raised there. Our method adapts a known technique from BSS computability, and should be comprehensible to casual readers as well as to logicians and computer scientists. It requires significant use of algebraic properties of the real numbers, in addition to computability, reinforcing the general perception of the BSS model as an essentially algebraic approach to computation on $\mathbb{R}$, treating real numbers as indivisible finite items. In contrast, the use of computable analysis normally results in a more analytic approach to computation on $\mathbb{R}$. As computable model theorists with experience in algorithms on (countable) Turing-computable fields, we the present authors are more familiar with the algebraic side.

Our notation generally follows that of [8]. The set of all finite tuples of real numbers is denoted $\mathbb{R}^\infty$; the inputs and outputs of BSS machines on $\mathbb{R}$ all lie in this set, and the collective content of the cells of a BSS machine at a given stage in a computation may also be regarded as an element of $\mathbb{R}^\infty$. We use $\mathbb{A}$ to denote the set of all real numbers which are algebraic over the subfield $\mathbb{Q}$ of rational numbers. $\mathbb{A}$ is partitioned into subsets $\mathbb{A}_{=d}$, for each $d \in \omega$: $\mathbb{A}_{=d}$ contains those algebraic real numbers of degree exactly d over $\mathbb{Q}$. (Recall that the *degree* of x over $\mathbb{Q}$ is the vector space dimension over $\mathbb{Q}$ of the field $\mathbb{Q}(x)$ generated by x; equivalently, it is the degree of the minimal polynomial of x in $\mathbb{Q}[X]$.) We also write $\mathbb{A}_d = \cup_{c \leq d} \mathbb{A}_{=c}$, the set of algebraic real numbers of degree $\leq d$. By the definition of degree, $\mathbb{A}_0$ is empty, and $\mathbb{A}_1$ contains exactly the rational numbers themselves. We mention [9] as an excellent source for these and other algebraic preliminaries, and [4] for more advanced questions about algorithms on fields.

The following lemma is well known, and clear by induction on stages. It reflects the fact that the four field operations are the only operations which a BSS machine is able to perform.

Lemma 1.1 *If M is a BSS machine using only the real parameters $\vec{z}$ in its program, then at every stage of the run of M on any input $\vec{x}$, the content of every cell lies in the field $\mathbb{Q}(\vec{z},\vec{x})$.* ∎

It is immediate from this lemma that the set $\mathbb{A}$ of algebraic real numbers cannot be the image of ω under any BSS-computable function, as it is not contained within any finitely generated field. (Here ω represents the set of nonnegative integers, viewed as a subset of $\mathbb{R}$.) We say that $\mathbb{A}$ is not *BSS-denumerable*. On the other hand, $\mathbb{A}$ does satisfy the definition of *BSS semidecidability*, which is the best analogue of Turing-computable enumerability and has been studied more closely in the literature.

Definition 1.2 A set $S \subseteq \mathbb{R}^\infty$ is *BSS-semidecidable* if there exists a (partial) BSS-computable function with domain S, and *BSS-denumerable* if there exists a partial BSS-computable function mapping ω onto S. S is *BSS-decidable* if its characteristic function χ_S is BSS-computable.

It is immediate that S is BSS-decidable if and only if both S and $(\mathbb{R}^\infty - S)$ are BSS-semidecidable. This justifies the analogy between BSS-semidecidability in $\mathbb{R}^\infty$ and computable enumerability in ω, and also dictates the use of the prefix "semi." The term *BSS-denumerable*, on the other hand, suggests that the set can be listed out, element by element, by a BSS machine, which is precisely the content of the definition above. (The adjective *denumerable* was once a synonym for *countable*, but has fallen out of use in recent years.) In the context of Turing computability, computable enumerability and semidecidability are equivalent, but in the BSS context, the set $\mathbb{A}$ distinguishes the two notions, being BSS-semidecidable but not BSS-denumerable. (On the other hand, every BSS-denumerable set is readily seen to be BSS-semidecidable.) The semidecision procedure for $\mathbb{A}$ is well-known: take any input x, and go through all nonzero polynomials $p(X) \in \mathbb{Q}[X]$, computing $p(x)$ for each. If ever $p(x) = 0$, the machine halts.

The ability to go through the polynomials in $\mathbb{Q}[X]$ follows from the BSS-denumerability of $\mathbb{Q}[X]$, which in turn follows from the BSS-denumerability of $\mathbb{Q}$. (A similar result applies to the set of algebraically dependent tuples in $\mathbb{R}^\infty$; see for instance [7].)

The question which gave rise to this paper was posed by Meer and Ziegler in [8]. (There they credit it to an anonymous referee of that paper.) It uses the notion of a *BSS reduction*, analogous to Turing reductions. A *oracle BSS machine* is essentially a BSS machine with the additional ability to take any finite tuple (which it has already assembled on the cells of its tape), ask an oracle set A whether that tuple lies in A, and fork according to whether the answer is positive or negative. The oracle A should be a subset of $\mathbb{R}^\infty$, of course, and we will write M^A to represent an oracle BSS program (or machine) equipped with an oracle set A. Oracle BSS programs can be enumerated (by tuples from $\mathbb{R}^\infty$) in much the same manner as regular BSS programs. If $B \subseteq \mathbb{R}^\infty$ and the characteristic function χ_B can be computed by an oracle BSS machine M^A with oracle A, then we write $A \leq_{BSS} B$, and say that A is *BSS-reducible* to B, calling M the *BSS reduction* of A to B. Should $A \leq_{BSS} B$ and also $B \leq_{BSS} A$, we write $A \equiv_{BSS} B$ and call the two sets *BSS-equivalent*. All this is exactly analogous to oracle Turing computation on subsets of ω.

Question 1.1 *Let $\mathbb{A}$ be the set of algebraic numbers in $\mathbb{R}$, i.e. those which are roots of a nonzero polynomial in $\mathbb{Q}[X]$. Also, let $\mathbb{H}$ be the Halting Problem for BSS computation on $\mathbb{R}$, as described in [1, §3.5]. Is it true that $\mathbb{H} \not\leq_{BSS} \mathbb{A}$? And more generally, could any countable subset of $\mathbb{R}^\infty$ contain enough information to decide $\mathbb{H}$?*

That $\mathbb{A} \leq_{BSS} \mathbb{H}$ is immediate. Let P be the BSS program which, on input $x \in \mathbb{R}$, plugs x successively into each nonzero polynomial $p(X)$ in (the BSS-denumerable set) $\mathbb{Q}[X]$ and halts if ever $p(x) = 0$. Then $x \in \mathbb{A}$ iff the program P halts on input x. (Similarly, every BSS-semidecidable set is BSS-decidable in $\mathbb{H}$, and indeed 1-reducible to $\mathbb{H}$ in the BSS model.) The focus of the question is on the lack of any reduction in the opposite direction. Section 2 gives the basic technical lemma used in this paper to address such questions, and Section 3 applies it to give a positive answer to Question 1.1. We also prove there a more general theorem relating BSS degrees to cardinality, showing that for infinite subsets $S \subseteq \mathbb{R}$ and $C \subseteq \mathbb{R}^\infty$, if $S \leq_{BSS} C$, then the local cardinality (in a technical sense defined in that section) of S cannot be greater than the (global, i.e. usual) cardinality of C.

2 BSS-Computable Functions At Transcendentals

Here we introduce our basic method for showing that various functions on the real numbers fail to be BSS-computable. In Section 3, this method will be extended to give answers about BSS-computability below certain oracles. However, even the non-relativized version yields straightforward proofs of several well-known results about BSS-decidable sets, as we will see shortly after describing the method.

In many respects, our method is equivalent to the method, used by many others, of considering BSS computations as paths through a finite-branching tree of height ω, branching whenever there is a forking instruction in the program. However, we think that the intuition for our method can be more readily explained to a mathematician unfamiliar with computability theory. Our straightforward main lemma says that near any transcendental input in its domain, a BSS-machine must be defined by rational functions. Where previous proofs usually made arguments about countable sets of terminal nodes in the tree of possible computations, we simply use the transcendence of this element.

Lemma 2.1 *Let M be a BSS-machine, and $\vec{z}$ the finite tuple of real parameters mentioned in the program for M. Suppose that $\vec{y} \in \mathbb{R}^{m+1}$ is a tuple of real numbers algebraically independent over the field $Q = \mathbb{Q}(\vec{z})$, such that M converges on input $\vec{y}$. Then there exists $\varepsilon > 0$ and rational functions $f_0, \ldots, f_n \in Q(\vec{Y})$,* *(that is, rational functions of the variables $\vec{Y}$ with coefficients from Q) such that for all $\vec{x} \in \mathbb{R}^{m+1}$ with $|\vec{x}-\vec{y}| < \varepsilon$, M also converges on input $\vec{x}$ with output $\langle f_0(\vec{x}), \dots, f_n(\vec{x})\rangle \in \mathbb{R}^{n+1}$.*

Proof. The intuition is that by choosing $\vec{x}$ sufficiently close to $\vec{y}$, we can ensure that the computation on $\vec{x}$ branches in exactly the same way as the computation on $\vec{y}$, at each of the (finitely many) branch points in the computation on $\vec{y}$. More formally, say that the run of M on input $\vec{y}$ halts at stage t, and that at each stage $s \leq t$, the non-blank cells contain the reals $\langle f_{0,s}(\vec{y}), \dots, f_{n_s,s}(\vec{y})\rangle$. Lemma 1.1 shows that all $f_{i,s}(\vec{y})$ lie in the field $Q(\vec{y})$, so each $f_{i,s}$ may be viewed as a rational function of $\vec{y}$ with coefficients in Q. Indeed, each rational function $f_{i,s}$ is uniquely determined in $Q(\vec{Y})$, since $\vec{y}$ was chosen algebraically independent over Q.

Let F be the finite set

$$F = \{f_{i,s}(\vec{Y}) : s \leq t \ \& \ i \leq n_s \ \& \ f_{i,s} \notin Q\},$$

the set of nonconstant rational functions used in the computation. Now for each $f_{i,s} \in F$, the preimage $f_{i,s}^{-1}(0)$ is closed in $\mathbb{R}^{m+1}$, and therefore so is the finite union

$$U = \bigcup_{f_{i,s} \in F} f_{i,s}^{-1}(0).$$

By algebraic independence, $\vec{y}$ does not lie in U, so there exists an $\varepsilon > 0$ such that the ε-ball $B_\varepsilon(\vec{y}) = \{\vec{x} \in \mathbb{R}^{m+1} : |\vec{x}-\vec{y}| < \varepsilon\}$, does not intersect the closed set U, and is contained within the domain of all $f_{i,s} \in F$. This will be the ε demanded by the lemma. Notice that more is true: for all $f_{i,s} \in F$ and all $\vec{x} \in B_\varepsilon(\vec{y})$, $f_{i,s}(\vec{x})$ and $f_{i,s}(\vec{y})$ must have the same sign, since otherwise there would be a path from $\vec{x}$ to $\vec{y}$ within $B_\varepsilon(\vec{y})$, along which $f_{i,s}$ would have to assume the value 0.

Now fix any $\vec{x} \in B_\varepsilon(\vec{y})$. We claim that in the run of M on input $\vec{x}$, at each stage $s \leq t$, the cells will contain precisely $\langle f_{0,s}(\vec{x}), \dots, f_{n_s,s}(\vec{x})\rangle$ and the machine will be in the same state in which it was at stage s on input $\vec{y}$. This is clear for stage 0, and we continue by induction, going from each stage $s < t$ to stage $s+1$. If the machine executed a copy instruction or a field operation in this step, then the result is clear, by inductive hypothesis. Otherwise, the machine executed a fork instruction, comparing some $f_{i,s}(\vec{x})$ with 0. But we saw above that $f_{i,s}(\vec{x})$ and $f_{i,s}(\vec{y})$ have the same sign (or else $f_{i,s}(y) = 0$, in which case $f_{i,s}$ is the constant function 0), so in both runs the machine entered the same state at stage $s+1$, leaving the contents of all cells intact. This completes the induction, and leaves us only to remark that therefore, at stage t, the run of M on input $\vec{x}$ must also have halted, with $\langle f_{0,t}(\vec{x}), \dots, f_{n,t}(\vec{x})\rangle$ in its cells as the output. ■

(If our BSS machines were allowed to compare the contents of two cells under $=$ or $<$, as is standard, then our set F would have to consist of all nonconstant differences $(f_{i,s} - f_{j,s})$. The proof would still work, but the method above is simpler.)

Lemma 2.1 provides quick proofs of several known results, including the undecidability of every proper subfield $F \subset \mathbb{R}$.

Corollary 2.2 *No BSS-decidable subset $S \subseteq \mathbb{R}^n$ can be both dense and co-dense in $\mathbb{R}^n$.*

Proof. If the characteristic function χ_S were BSS-computable, say by some machine M with parameters $\vec{z}$, then by Lemma 2.1, it would be constant in some neighborhood of every $\vec{y} \in \mathbb{R}^n$ with coordinates algebraically independent over $\vec{z}$. ■

Indeed, the same proof shows that any BSS-computable total function with discrete image must be constant on each of the ε-balls given by Lemma 2.1.

Corollary 2.3 *Define the boundary of a subset $S \subseteq \mathbb{R}^n$ to be the intersection of the closure of S with the closure of its complement. If S is BSS-decidable, then there is a finite tuple $\vec{z}$ such that every point on the boundary of S has coordinates algebraically dependent over $\vec{z}$. In particular, if M computes χ_S, then its parameters may serve as $\vec{z}$.*

Proof. This is immediate from Lemma 2.1. ■

Of course, Corollaries 2.2 and 2.3 have been deduced long since from other known results, in particular from the Path Decomposition Theorem described in [1]. We include them here because of the simplicity of these proofs, and because they introduce the methods to be used in the following section.

3 Countable Oracle Sets

It is natural to think of countability of a subset $S \subseteq \mathbb{R}^\infty$ as a bound on the amount of information which can be encoded into S. This intuition requires significant restating before it can be made into a coherent (let alone true) statement, but we will give a reasonable version in this section. In [8], it was asked whether there could exist a countable set $C \subseteq \mathbb{R}^\infty$ such that the halting problem $\mathbb{H}$ for BSS computation on $\mathbb{R}$ satisfies $\mathbb{H} \leq_{BSS} C$. We will show that the answer to this question is negative. For a formal definition of $\mathbb{H}$ in this context, we refer the reader to [1, §3.5]. Since it is equiconsistent with **ZFC** for the Continuum Hypothesis to be false, we will make our arguments applicable to all infinite cardinals $\kappa < 2^{\aleph_0}$, countable or otherwise.

First, of course, every subset of $\mathbb{R}^\infty$ is BSS-equivalent to its complement, and so countability and co-countability impose the same restriction on information content. Of course, many sets of size continuum, with equally large complements, are quite simple: the set of positive real numbers, for example, is BSS-decidable, hence less complex than the countable set $\mathbb{Q}$ (cf. Corollary 2.2). So it is not possible to prove absolute results relating cardinality and co-cardinality (within $\mathbb{R}^\infty$) to BSS reducibility, but nevertheless, we can produce theorems expressing the intuition that countable sets are not highly complex in the BSS model. This process will culminate in Theorem 3.4 below, but first we show that with a countable oracle, one cannot decide the BSS halting problem $\mathbb{H}$. We conjecture that $\mathbb{H}$ is not an upper bound on the degree of a countable set, i.e. that such a set can still be BSS-incomparable with $\mathbb{H}$, but no matter whether that conjecture holds or fails, it certainly constitutes progress just to know that the upper cone of sets above $\mathbb{H}$ contains no countable sets.

Theorem 3.1 *If $C \subseteq \mathbb{R}^\infty$ is a set such that $\mathbb{H} \leq_{BSS} C$, then $|C| = 2^{\aleph_0}$.*

We note that by BSS-equivalence, these conditions also ensure $|\mathbb{R}^\infty - C| = 2^{\aleph_0}$, and ensure $|\mathbb{R}^m - C| = 2^{\aleph_0}$ whenever $C \subseteq \mathbb{R}^m$.

Proof. Let $C \subseteq \mathbb{R}^\infty$ have cardinality $< 2^{\aleph_0}$, and suppose that M is an oracle BSS machine such that M^C computes the characteristic function of $\mathbb{H}$. We fix a program code number p for the program which accepts inputs $\langle x_1, x_2 \rangle \in \mathbb{R}^2$, searches through nonzero polynomials q in $\mathbb{Q}[Y_1, Y_2]$, and halts iff it finds one with $q(x_1, x_2) = 0$. Since the program coded by p uses no real parameters, p may be regarded as a natural number, but in our argument it can equally well be a tuple $\vec{p}$ from $\mathbb{R}^\infty$, with one or several real numbers coding program parameters. Then the elements of C, the finitely many parameters $\vec{z}$ of M, and the parameters, if any, in the program coded by $\vec{p}$ together generate a field $E \subseteq \mathbb{R}$ which also has cardinality $< 2^{\aleph_0}$, and so $\mathbb{R}$ is an extension of infinite transcendence degree (indeed of degree $2^{\aleph_0}$) over this E. (Since $C \subseteq \mathbb{R}^\infty$, we need to be precise: E is generated by the coordinates $p_1, \ldots, p_j$ and $z_1, \ldots, z_k$ of the tuples $\vec{p}$ and $\vec{z}$, and the coordinates of each tuple in C.)

Now fix a pair $\langle y_1, y_2\rangle$ of real numbers algebraically independent over E. Hence $\langle \vec{p}, y_1, y_2\rangle \notin \mathbb{H}$, so M^C on this input halts after finitely many steps and outputs 0. As in Lemma 2.1, we fix the finitely many functions $f_{i,s}(\vec{Y}) \in E(Y_1, Y_2)$ such that $f_{i,s}(y_1, y_2)$ appears in the i-th cell at stage s during this computation. (The program code $\vec{p} \in E^\infty$ will stay fixed throughout this proof, so we may treat it as part of the function $f_{i,s}$, rather than as a variable.) Let F be the set of those functions $f_{i,s}$ which are not constants in E, and fix an $\varepsilon > 0$ such that whenever $\langle x_1, x_2\rangle \in \mathbb{R}^2$ with $x_1 \in B_\varepsilon(y_1)$ and $x_2 \in B_\varepsilon(y_2)$, every $f \in F$ satisfies $f(x_1, x_2) \cdot f(y_1, y_2) > 0$. Write each $f \in F$ as a quotient $f = \frac{g}{h}$ with $g, h \in E[Y_1, Y_2]$ in lowest terms, and let n be the greatest degree of Y_2 in all of these finitely many polynomials g and h.

So far this mirrors the proof of Lemma 2.1, but an additional condition is needed. The oracle BSS machine M^C, running on input $\langle \vec{p}, x_1, x_2\rangle$, can ask its oracle, at any stage s and for any cell i, whether $f_{i,s}(x_1, x_2)$ lies in the oracle set C, and can fork according to the oracle answer. So, in addition to choosing $\langle x_1, x_2\rangle$ within ε of $\langle y_1, y_2\rangle$, we must ensure, for every i and s, that $[f_{i,s}(x_1, x_2) \in C \iff f_{i,s}(y_1, y_2) \in C]$. On input $\langle \vec{p}, y_1, y_2\rangle$, we know by algebraic independence over E that $f_{i,s}(y_1, y_2) \notin C$ unless $f_{i,s}$ is a constant function (in which case $f_{i,s}(x_1, x_2) = f_{i,s}(y_1, y_2)$, of course). So, for all of the finitely many $f \in F$, we need to ensure that $f_{i,s}(x_1, x_2) \notin C$ as well.

Now choose $x_1 \in \mathbb{R}$ to be transcendental over E and within ε of y_1, and pick x_2 within ε of y_2 such that x_2 is algebraic over $\mathbb{Q}(x_1)$ but has degree $> n$ over $E(x_1)$. For instance, let $x_1 = y_1$ and $x_2 = \sqrt[m]{x_1} + b$, where $m > n$ is prime and $b \in \mathbb{Q}$ is selected to place $x_2 \in B_\varepsilon(y_2)$. It follows from [6, Exercise 1, p. 256] that the polynomial $(Y^m - x_1)$ is irreducible in the one-variable polynomial ring $E(x_1)[Y]$, so this x_2 has degree m over $E(x_1)$.

Thus, for any $f \in F$, if $a = f(x_1, x_2) \in E$, then $0 = g(x_1, x_2) - ah(x_1, x_2)$. Since f is nonconstant, g is not a scalar multiple of h, and so $(g - ah)$ would then be a nonzero polynomial in $E[Y_1, Y_2]$ of degree $\leq n$, contradicting our choice of x_2. Hence $f(x_1, x_2) \notin E$ for every $f \in F$. But then the oracle computation $M^C(\vec{p}, x_1, x_2)$ must follow the same path as $M^C(\vec{p}, y_1, y_2)$ and give the same output, namely 0. Since $\langle \vec{p}, x_1, x_2\rangle \in \mathbb{H}$, this proves that M^C does not compute the characteristic function of $\mathbb{H}$. ■

Indeed the preceding proof shows more than was stated.

Corollary 3.2 *If $C \subseteq \mathbb{R}^\infty$ is a set such that $\mathbb{H} \leq_{BSS} C$, then $\mathbb{R}$ has finite transcendence degree over the field K generated by (the coordinates of the tuples in) C, and also has finite transcendence degree over the field generated by the complement of C.*

Proof. Given an oracle BSS machine M which computes $\mathbb{H}$ from oracle C, let E be the extension field $K(\vec{z}, \vec{p})$, with K as defined in the corollary. If $\mathbb{R}$ had transcendence degree ≥ 2 over this E, then the proof of Theorem 3.1 would go through: we could choose $y_1, y_2 \in \mathbb{R}$ algebraically independent over E, say with $y_1 > 0$, and again let $x_1 = y_1$ and $x_2 = b + \sqrt[m]{x_1}$, with m and b as in that proof. But this would show that M^C does not compute $\mathbb{H}$. So $\mathbb{R}$ has transcendence degree ≤ 1 over this E, and therefore is algebraic over $E(t) = K(t, \vec{z}, \vec{p})$ for some $t \in \mathbb{R}$.

Since C is BSS-equivalent to its complement, the same proof applies to $(\mathbb{R}^\infty - C)$, and also to $(\mathbb{R}^m - C)$ if $C \subseteq \mathbb{R}^m$. ■

As we consider the general case of a BSS computation of the characteristic function χ_S of a set $S \subseteq \mathbb{R}$ using an oracle C of infinite cardinality $\kappa < 2^{\aleph_0}$, the following definition will be useful. Here $\overline{S}$ denotes $(\mathbb{R} - S)$, the complement of S in $\mathbb{R}$ (as opposed to the topological closure).

Definition 3.3 A set $S \subseteq \mathbb{R}$ is *locally of bicardinality* $\leq \kappa$ if there exist two open subsets U and V of $\mathbb{R}$ with $|\mathbb{R} - (U \cup V)| \leq \kappa$ and $|U \cap S| \leq \kappa$ and $|V \cap \overline{S}| \leq \kappa$.

The *local bicardinality of S* is the least cardinal κ such that S is locally of bicardinality $\leq \kappa$.

If $\kappa < 2^{\aleph_0}$, then such U and V must be disjoint, since $(U \cap V)$ is open with $|U \cap V| \leq |U \cap S| + |V \cap \overline{S}| \leq \kappa$. So the definition roughly says that up to sets of size κ, each of S and $\overline{S}$ is equal to an open subset of $\mathbb{R}$. In Lemma 4.2 below, we will show that the Cantor middle-thirds set has local bicardinality $2^{\aleph_0}$.

The property of local bicardinality $\leq \kappa$ does not appear to us to be equivalent to any more easily stated property, and we are not aware of it having been used (or even stated) elsewhere in the literature. The same definition in higher dimensions completely loses its power: any connected component U_0 of U must have boundary ∂U_0 with $U_0 \cap \partial U = V \cap \partial U_0 = \emptyset$, since U and V are open and disjoint. But then $|\partial U_0| \leq |\mathbb{R}^n - (U \cup V)| \leq \kappa$, which is feasible in $\mathbb{R}^1$ but not in higher dimensions, unless U or V were empty or $\kappa = 2^{\aleph_0}$. Thus, in $\mathbb{R}^n$ with $n > 1$, every set of local bicardinality $< 2^{\aleph_0}$ has either cardinality $< 2^{\aleph_0}$ or co-cardinality $< 2^{\aleph_0}$. Nevertheless, within $\mathbb{R}^1$, this is exactly the condition needed in our general theorem on cardinalities.

Theorem 3.4 *If $C \subseteq \mathbb{R}^\infty$ is an oracle set of infinite cardinality $\kappa < 2^{\aleph_0}$, and $S \subseteq \mathbb{R}$ is a set with $S \leq_{BSS} C$, then S must be locally of bicardinality $\leq \kappa$. The same holds for oracles C of infinite co-cardinality $\kappa < 2^{\aleph_0}$.*

Proof. Again let $\vec{z}$ be the parameters used by the oracle BSS machine M which, given oracle C, computes χ_S. Then for any input $y \in \mathbb{R}$ transcendental over the subfield E of cardinality κ generated by $\vec{z}$ and the individual coordinates of all elements of C, there will again exist a finite set $F_y \subseteq E(X)$ as above, and an $\varepsilon > 0$ such that $f(x) \cdot f(y) > 0$ for all $x \in B_\varepsilon(y)$ and $f \in F_y$. For each such y, let $B(y)$ be an open interval of length less than the corresponding ε, such that $B(y)$ contains y and has rational end points. Now if $x \in B(y)$ is also transcendental over E, then the computation of $\chi_S(x)$ using this machine and the C-oracle proceeds along the same path as the computation for y, since $f(x) \notin E$ for all $f \in F_y$. (Indeed, this would hold whenever $x \in B(y)$ has degree $> n$ over E, where n is the maximum degree of all numerators and denominators of elements of F_y.) This shows that $\chi_S(x) = \chi_S(y)$ for all such x. Since only κ-many elements of $B(y)$ can be algebraic over the size-κ field E, it follows that either $|S \cap B(y)| \leq \kappa$ (if $y \notin S$) or $|\overline{S} \cap B(y)| \leq \kappa$ (if $y \in S$).

Now if $t \in B(y_0) \cap B(y_1)$ with t, y_0, and y_1 each transcendental over E, then t follows the same computation path as both y_0 and y_1, implying that $\chi_S(y_0) = \chi_S(y_1)$ whenever $B(y_0) \cap B(y_1) \neq \emptyset$, and therefore that either $B(y_0) \cap S$ and $B(y_1) \cap S$ both have size $\leq \kappa$, or else $B(y_0) \cap \overline{S}$ and $B(y_1) \cap \overline{S}$ both have size $\leq \kappa$. So when we set

$$U = \bigcup \{B(y) : |S \cap B(y)| \leq \kappa\} \quad \text{and} \quad V = \bigcup \{B(y) : |\overline{S} \cap B(y)| \leq \kappa\},$$

we will have $U \cap V = \emptyset$. Here the unions are over those $y \in \mathbb{R}$ transcendental over E (as $B(y)$ is not defined for any other y), and so the complement $\mathbb{R} - (U \cup V)$ is a subset of the algebraic closure of E, which has size κ. Moreover, being a union of open intervals $B(y)$ with rational end points, U in fact equals the union of countably many such intervals, say $U = \cup_{i \in \omega} B(y_i)$ for some sequence $y_0, y_1, \ldots$. Since each $B(y_i)$ has intersection of size $\leq \kappa$ with S (and since $\kappa \geq \aleph_0$), so does the entire union U. Likewise $|\overline{S} \cap V| \leq \kappa$, proving the theorem.

The claim about oracles of co-cardinality κ follows from applying the same argument to the oracle $(\mathbb{R}^\infty - C)$, which is BSS-equivalent to C. If $C \subseteq \mathbb{R}^m$ for some m, then the same holds of $(\mathbb{R}^m - C)$. ■

Notice that the set S of smaller complexity must be a subset of $\mathbb{R}$, whereas C is allowed to contain tuples from $\mathbb{R}^\infty$. We conjecture that to extend the theorem to sets $S \subseteq \mathbb{R}^\infty$, we would need to allow $\mathbb{R}^\infty - (U \cup V)$ to be a union of κ-many proper algebraic varieties defined over the field generated by C. It is an open question (of interest only under $\neg$**CH**) whether it is equivalent, for the purposes of this conjecture and Theorem 3.4, to replace $|\mathbb{R}^\infty - (U \cup V)| \leq \kappa$ by $|\mathbb{R}^\infty - (U \cup V)| \leq \aleph_0$ here or in Definition 3.3.

To understand that this theorem cannot readily be stated using a simpler property than Definition 3.3, consider the BSS-computable set

$$S = \{x \in (0,1) : (\exists m \in \omega)\ 2^{-(2m+1)} \leq x \leq 2^{-(2m)}\},$$

containing those $x \in (0,1)$ which have a binary expansion beginning with an even number of zeroes. Then clearly no open interval B which is locally of bicardinality $\leq \kappa < 2^{\aleph_0}$ can contain any of the countably many points 2^{-m}, so the theorem cannot require the complement $\mathbb{R}^\infty - (U \cup V)$ to be finite, let alone empty. Moreover, every open interval $B \subseteq \mathbb{R}$ which either contains 0 or has left end point 0 must have intersection of size $2^{\aleph_0}$ with both S and $\overline{S}$. One can make the same happen not only at 0, but at each rational in a sequence approaching 0, and with such tricks one can create examples defying most conceivable simplifications of Theorem 3.4.

4 The Cantor Set

As an example of a set of local bicardinality $2^{\aleph_0}$, we consider the Cantor set C, well known as a set of measure 0 within $\mathbb{R}$ which nevertheless has cardinality $2^{\aleph_0}$. By definition, C contains all real numbers $x \in [0,1]$ having ternary expansions in only 0's and 2's. One usually views C as the set of numbers in the unit interval $[0,1]$ which remain after ω-many iterations of deleting the open "middle third" of each interval (starting with the middle third $(\frac{1}{3}, \frac{2}{3})$ of $[0,1]$). It is clear from this description that C is co-semidecidable in the BSS model: even a Turing machine can enumerate all those middle-third intervals to be deleted. Hence $\overline{C} \leq_{BSS} \mathbb{H}$ (indeed via a 1-reduction), forcing $C \leq_{BSS} \mathbb{H}$ as well. The natural next question, whether $\mathbb{H} \leq_{BSS} C$, was settled in [10], as described below.

Lemma 4.1 *The Cantor set C is not BSS-semidecidable.*

Proof. Since $\overline{C}$ is semidecidable, semidecidability of C would show that C was BSS-decidable. However, for every BSS-machine with finite parameter tuple $\vec{z}$, C contains some y transcendental over $\mathbb{Q}(\vec{z})$, since otherwise C would be countable. Now no nonempty open interval within $\mathbb{R}$ is contained within C, and so every ε-ball around y contains elements of $\overline{C}$. Lemma 2.1 therefore shows that M does not compute the characteristic function χ_C. ∎

The next lemma, combined with Theorem 3.4, would also immediately prove Lemma 4.1. On the other hand, it dashes the hope that Theorem 3.4 might prove $\mathbb{H} \not\leq_{BSS} C$ the same way it proved $\mathbb{H} \not\leq_{BSS} \mathbb{A}$.

Lemma 4.2 *The Cantor set C has local bicardinality $2^{\aleph_0}$.*

Proof. Suppose C were locally of bicardinality $\leq \kappa < 2^{\aleph_0}$. Then we would have open disjoint sets U and V satisfying Definition 3.3, and C, having size $2^{\aleph_0}$, would have to intersect V in some point x, since

$$C - V \subseteq (U \cap C) \cup \overline{(U \cup V)}$$

and the right-hand side has size $\leq \kappa$. The open set V would then contain an ε-ball around x. However, every open interval around x intersects each of C and $\overline{C}$ in $2^{\aleph_0}$-many points. (To see this, just consider all y whose ternary expansions match that of x for sufficiently many places to lie within that interval.) Therefore $|V \cap \overline{C}| = 2^{\aleph_0}$, yielding a contradiction. ∎

Corollary 4.3 *The Cantor set C is not BSS-semidecidable below $\mathbb{A}$, or below any other oracle of cardinality $< 2^{\aleph_0}$.*

Proof. This simply means that no function which is BSS-computable in the oracle $\mathbb{A}$ can have C as its domain. Indeed, if it did, then $C \leq_{BSS} \mathbb{A}$, since C and $\overline{C}$ would both be $\mathbb{A}$-semidecidable. Lemma 4.2 and Theorem 3.4 together rule out this possibility. The same holds for any other oracle of size $< 2^{\aleph_0}$. ■

Corollary 3.2, our other natural hope for proving $\mathbb{H} \not\leq_{BSS} C$, also fails to do so, for the field generated by C does not satisfy the hypothesis there. It seems counterintuitive that a set of measure 0 could generate such a large field, so we prove it here. (The authors assume that this fact has been proven long since, and would appreciate a reference for it.)

Lemma 4.4 (Folklore) *The Cantor set C generates the entire field $\mathbb{R}$. Indeed, it generates $\mathbb{R}$ as a ring.*

Proof. The argument is best understood by seeing an example. Here we begin with an element of $[0,1]$, chosen arbitrarily, in ternary form:

$$\begin{aligned} & 0.2201020001211\ldots \\ = \; & 0.2200020000200\ldots \\ & +0.0001000001011\ldots \\ = \; & 0.2200020000200\ldots \\ & +(0.0002000002022\ldots)\cdot\frac{1}{2} \end{aligned}$$

Since $\frac{1}{2}$ lies in every subfield of $\mathbb{R}$, this shows that this number is generated from C by field operations. Indeed, since $\frac{1}{2} = 2\cdot\frac{1}{4} = 2\cdot(0.020202\ldots)$, the number is generated from elements of C by ring operations. The same process can be applied to any element of $[0,1]$, so C generates the entire unit interval, and hence all of $\mathbb{R}$. ■

At this point the authors abandoned their search for a proof that $\mathbb{H} \not\leq_{BSS} C$. Fortunately, an anonymous referee familiar with Yonezawa's paper [10] pointed out the necessary result there.

Theorem 4.5 (Corollary 2.5 in [10]) *The sets $\mathbb{Q}$ and C are BSS-incomparable.* ■

Since the BSS-semidecidable set $\mathbb{Q}$ must be $\leq_{BSS} \mathbb{H}$, this immediately answers the question.

Corollary 4.6 $\mathbb{H} \not\leq_{BSS} C$. ■

5 Other Results

In addition to the theorems on cardinality described above, the authors have proven a selection of results on BSS-reducibility among the different sets $\mathbb{A}_{=d}$, where

$$\mathbb{A}_{=d} = \{x \in \mathbb{R} : x \text{ is algebraic over } \mathbb{Q} \text{ with minimal polynomial of degree } d\}.$$

For reasons of space, we omit most discussion of these theorems here, as well as their proofs. (They were presented by the third author in a short talk at the meeting *Logical Approaches to Computational Barriers* in Greifswald, Germany in February 2010.) However, we do state the main theorems here. The basic result simply concerns $\mathbb{A}_{=d-1}$ and $\mathbb{A}_{=d}$, and a proof appears in [3].

Theorem 5.1 *For every $d > 0$, $\mathbb{A}_{=d} \not\leq_{BSS} \mathbb{A}_{d-1}$.*

This is generalized to a pair of arbitrary degrees. Neither direction is trivial, but when p is prime to $\frac{r}{p}$, the backwards direction is implicit in [8], by Meer and Ziegler, and one particular case is explicitly shown by them.

Theorem 5.2 *Let p and r be any nonnegative integers. Then $\mathbb{A}_{=p} \leq_{BSS} \mathbb{A}_{=r}$ if and only if p divides r.*

Of course, $\mathbb{A}_{=0}$ is just the empty set, and $\emptyset \lneq_{BSS} \mathbb{A}_{=d}$ for all $d > 0$, since Meer and Ziegler showed in [8] that no $\mathbb{A}_{=d}$ with $d > 0$ is BSS-decidable. So the theorem also holds when $p = 0$, but not when $p > 0 = r$.

To extend these results further, we define, for all $S \subseteq \omega$, $\mathbb{A}_S = \cup_{d \in S} \mathbb{A}_{=d}$, the set of all algebraic real numbers whose degrees over $\mathbb{Q}$ lie in S. The proof of Theorem 5.1 is readily adjusted to yield the following.

Theorem 5.3 *For every $d > 0$ in ω and every set $S \subset \omega$ with $S \cap d\mathbb{Z} = \emptyset$, $\mathbb{A}_{=d} \not\leq_{BSS} \mathbb{A}_S$.*

Corollary 5.4 *Let P be the set of all prime numbers in ω. Then for all S and T in the power set $\mathscr{P}(P)$, $\mathbb{A}_S \leq_{BSS} \mathbb{A}_T$ if and only if $S \subseteq T$.*

An immediate further corollary imparts substantial richness to the partial order of the BSS-semidecidable degrees.

Corollary 5.5 *There is a subset $\mathscr{L}$ of the BSS-semidecidable degrees such that $(\mathscr{L}, \leq_{BSS}) \cong (\mathscr{P}(\omega), \subseteq)$.*

Proof. We have $(\mathscr{P}(\omega), \subseteq) \cong (\mathscr{P}(P), \subseteq)$, and Corollary 5.4 shows that the latter partial order embeds into the BSS-semidecidable degrees via the map $S \mapsto \mathbb{A}_S$. ■

We emphasize that Corollary 5.5 only states that there exists an isomorphism between the two partial orders. It is unknown whether this map is also an isomorphism of the two structures as lattices, or indeed whether an arbitrary $\mathbb{A}_S$ and $\mathbb{A}_T$ must have a greatest lower bound under $\leq_{BSS}$. Of course, for $S, T \subseteq P$, $\mathbb{A}_{S \cap T}$ is the obvious candidate, and if it really were the greatest lower bound, we would have many *minimal pairs* of BSS-semidecidable degrees. (Recall that in Turing computability, a *minimal pair* consists of two degrees **c** and **d** whose infimum is the computable degree **0**. The existence of a minimal pair of nonzero computably enumerable degrees was a significant result in Turing computability.)

Finally, we consider reducibility among the sets $\mathbb{A}_S$ and $\mathbb{A}_T$, for arbitrary $S, T \subseteq \omega$. Certain questions here remain open. First, we have a negative result.

Theorem 5.6 *For sets $S, T \subseteq \omega$, if $\mathbb{A}_S \leq_{BSS} \mathbb{A}_T$, then there exists $N \in \omega$ such that all $p \in S$ satisfy $\{p, 2p, 3p, \ldots, Np\} \cap T \neq \emptyset$.*

The next two propositions are both positive results (showing that reducibilities do exist). Proposition 5.7 uses a nonuniform construction, and therefore only applies when the set-theoretic difference $(S - T)$ is finite. Proposition 5.8 has a nonuniform construction, but requires a stronger hypothesis involving relative primality.

Proposition 5.7 *For any subsets S and T of ω, if $(S - T)$ is finite and for every $p \in S - T$, there exists an integer $q > 0$ such that $pq \in T$, then $\mathbb{A}_S \leq_{BSS} \mathbb{A}_T$.*

Proposition 5.8 *Let S and T be subsets of the positive integers. Suppose that for some absolute constant N and each $d \in S$, there is a positive integer $n_d \leq N$ and prime to d such that $dn_d \in T$. Then $\mathbb{A}_S \leq_{\mathrm{BSS}} \mathbb{A}_T$.*

Of course, n_d is allowed to equal 1, since 1 is prime to d. Thus every element of $S \cap T$ is immediately accounted for, and only elements of $(S - T)$ can pose problems. When $(S - T)$ is finite, Proposition 5.7 handles those problems, showing how to prove the result even in the absence of relative primality. When $(S - T)$ is infinite, the proof of Proposition 5.7 no longer applies. One would hope to be able to remove from Proposition 5.8 the assumption that n_d must be prime to d, or else to extend Theorem 5.6 to yield a nonreducibility result for this case, but for now this problem remains open.

References

[1] L. Blum, F. Cucker, M. Shub, and S. Smale; *Complexity and real computation* (Berlin: Springer-Verlag, 1997).

[2] L. Blum, M. Shub, and S. Smale; On a theory of computation and complexity over the real numbers, *Bulletin of the American Mathematical Society (New Series)* **21** (1989), 1–46.

[3] W. Calvert, K. Kramer, & R. Miller; Noncomputable functions in the Blum-Shub-Smale model, in the abstract booklet for the conference *Logical Approaches to Barriers in Computing and Complexity* (17-20 February 2010, Greifswald, Germany). Available at `qcpages.qc.cuny.edu/~rmiller/BSSabstract.pdf`.

[4] M.D. Fried & M. Jarden; *Field Arithmetic* (Berlin: Springer-Verlag, 1986).

[5] C. Gassner; A hierarchy below the halting problem for additive machines, *Theory of Computing Systems* **43** (2008) 3–4, 464–470.

[6] N. Jacobson; *Basic Algebra I* (New York: W.H. Freeman & Co., 1985).

[7] W. Koolen & M. Ziegler; Kolmogorov complexity theory over the reals, in *Proceedings of the Fifth International Conference on Computability and Complexity in Analysis, CCA '08*, *Electronic Notes in Theoretical Computer Science* **221** (Elsevier, 2008), 153-169.

[8] K. Meer and M. Ziegler; An explicit solution to Post's Problem over the reals, *Journal of Complexity* **24** (2008) 3–15.

[9] B.L. van der Waerden; *Algebra*, volume I, trans. F. Blum & J.R. Schulenberger (New York: Springer-Verlag, 1970 hardcover, 2003 softcover).

[10] Y. Yonezawa; The Turing degrees for some computation model with the real parameter, *J. Math. Soc. Japan* **60** 2 (2008), 311-324.