HALIN'S INFINITE RAY THEOREMS: COMPLEXITY AND REVERSE MATHEMATICS

JAMES S. BARNES, JUN LE GOH, AND RICHARD A. SHORE

Abstract. Halin [1965] proved that if a graph has (a set of) n many disjoint rays for each n then it has (a set of) infinitely many disjoint rays. We analyze the complexity of this and other similar results in terms of computable and proof theoretic complexity. The statement of Halin's theorem and the construction proving it seem very much like standard versions of compactness arguments such as König's Lemma. Those results while not computable are relatively simple. They only use arithmetic procedures or, equivalently, finitely many iterations of the Turing jump. We show that several Halin type theorems are much more complicated. They are among the theorems of hyperarithmetic analysis. Such theorems imply the ability to iterate the Turing jump along any computable well ordering. Several important logical principles in this class have been extensively studied beginning with work of Kreisel, H. Friedman, Steel and others in the 1960s and 1970s. Until now, only one purely mathematical example was known. Our work provides many more and so answers Question 30 of Montalbán's Open Questions in Reverse Mathematics [2011]. Some of these theorems including ones in Halin [1965] are also shown to have unusual proof theoretic strength as well.

1. Introduction

In this paper we analyze the complexity of several results in infinite graph theory. These theorems are said to be ones of Halin type or, more generally, of ubiquity theory. The classical example is a theorem of Halin [11]: If a countable graph G contains, for each n, a sequence $\langle R_0, \ldots R_{n-1}\rangle$ of disjoint rays (a ray is a sequence $\langle x_i | i \in \mathbb{N}\rangle$ of distinct vertices such that there is an edge between each x_i and x_{i+1}) then it contains an infinite such sequence of rays. (Halin actually deals with arbitrary graphs and formulates the result differently. The uncountable cases, however, are essentially just counting arguments. We deal only with countable structures but discuss his formulation in §6.) This standard formulation of his theorem seems like a typical compactness theorem going from arbitrarily many finitely many objects to an infinite collection. The archetypical example here is König's Lemma: If a finitely branching tree has paths of length n for every n then it has a branch, i.e. an infinite path. In outline, a modern proof of Halin's theorem for countable graphs (due to Andreae but see Diestel [5, Theorem 8.2.5(i)]) seems much like that of König's Lemma (and many others in infinite graph theory). The construction of the desired sequence of rays proceeds by a recursion through the natural numbers in which each step is a relatively simple procedure. While the procedure is much more delicate than for König's Lemma, it is basically of the same complexity. It uses Menger's theorem for finite graphs at each step but this represents a computable procedure (for finite graphs) and the other parts of the

Date: January 22, 2021.

2020 *Mathematics Subject Classification.* Primary 05C63, 03D55, 03B30; Secondary 03D80, 03F35, 05C38, 05C69, 05C70.

All the authors were partially supported by NSF Grant DMS-1161175.

step depend on the same type of information as in König's Lemma which ask, for example, if various sets (computable in the given graph) are nonempty or infinite. Nonetheless, we prove that the complexity of this construction and theorem are much higher than that for König's Lemma or other applications of compactness. (The definitions from graph theory, computability theory and proof theory/reverse mathematics that we need for our analysis are given in §3. Basic references for other terminology, background and standard results not otherwise attributed are Diestel [5] for graph theory; Rogers [23] and Sacks [25] for computability theory; and for reverse mathematics Simpson [29] with a view from proof theory and Hirschfeldt [13] with one from computability.)

We follow two well established procedures for measuring the complexity of constructions and theorems. The first is basically computability theoretic. It has its formal beginnings in the 1950s but has much earlier roots in constructive or computable mathematics reaching back to antiquity. (See Ershov et al. [6] for history and surveys of the approach in several areas of combinatorics, algebra and analysis.) The measuring rod here is relative computability. We say a set A is (Turing) computable from one B, $A \leq_T B$, if there is an algorithm (say on a Turing machine or any other reasonable model of general computation) that, when given access to all membership facts about B (an oracle for B) computes membership in A. The standard hierarchies of complexity here are based on iterations of the Turing jump. This operator takes B to B', the halting problem relativized to B, i.e. the set of programs with oracle for B, Φ_e^B, such that Φ_e^B halts on input e. For example, if the tree of König's Lemma is computable in B then there is a branch computable in the double jump B'' of B.

The second approach is proof theoretic. It measures the complexity of a theorem by the logical strength of the axioms needed to prove it. This approach also has a long history but the formal subject, now called reverse mathematics, starts with H. Friedman's work in the 1970s (e.g. [7, 8]). (Simpson [29] is now the basic reference.) One compares axiomatic systems S and T by saying that T is stronger than S, $T \vdash S$ (T proves S) if one can prove every sentence $\Theta \in S$ from the axioms of T. Of course, we know what it means for Θ to be provable in S. The goal here is to characterize to the extent possible the axioms needed to prove a given mathematical theorem Θ. To this end, one begins with a weak base theory. Then one wants to find a system S such that not only does $S \vdash \Theta$ but also Θ (with the weak base theory) proves all the axioms of S. Hence the name reverse mathematics as we seek to prove the "axioms" of S from the theorem Θ. Typically, the systems here are formalized in arithmetic with quantification over sets as well as numbers. The standard base theory (RCA_0) corresponds to the axioms needed to do computable constructions. Stronger systems are then usually generated by adding comprehension axioms which assert the existence of specific families of sets. For example, a very important system is ACA_0. It is equivalent in the sense of reverse mathematics just described to König's Lemma. Formally, it asserts that every subset of $\mathbb{N}$ defined by a formula that quantifies only over $\mathbb{N}$ (and not its subsets) exists. This is also equivalent to asserting that for every set B, the set B' exists and so each iteration $B^{(n)}$ for $n \in \mathbb{N}$ exists.

The early decades of reverse mathematics were marked by a large variety of results characterizing a wide array of theorems and constructions as being one of five or so specific levels of complexity including RCA_0 and ACA_0. Each of these systems (Simpson's "big five") have corresponding specific recursion theoretic construction principles. In more recent decades, there has been a proliferation of results placing theorems and constructions outside the big five. Sometimes inserted linearly and sometimes with incomparabilities. They are now collectively often called the "zoo" of reverse mathematics. (See https://rmzoo.math.uconn.edu/diagrams/ for pictures.).

Theorems and constructions in combinatorics in general, and graph theory in particular, have been a rich source of such denizens of this zoo. Almost all of them have fallen below ACA_0 (König's Lemma) and so have the objects they seek constructible computably in finitely many iterations of the Turing jump. Ramsey theory, in particular, has provided a very large class of constructions and theorems of distinct complexity. One example of the infinite version of a classical theorem of finite graph theory that is computationally and reverse mathematically strictly stronger than ACA_0 is König's Duality Theorem (KDT) for infinite or even countable graphs. (Every bipartite graph has a matching and a cover consisting of one vertex from each edge of the matching.) The proofs of this theorem for infinite graphs (Podewski and Steffens [22] for countable and Aharoni [1] for arbitrary ones) are not just technically difficult but explicitly used both transfinite recursions and well orderings of all subsets of the given graph). These techniques lie far beyond ACA_0. Aharoni, Magidor and Shore [2] proved that this theorem is of great computational strength in that there are computable graphs for which the required matching and cover compute all the iterations of the Turing jump through all computable well-orderings. They also showed that it was strong reverse mathematically as it implied ATR_0, the standard system above ACA_0 used to deal with such transfinite recursions. Some of the lemmas used in each of the then known proofs were shown to be equivalent to the next and final of the big five systems, Π^1_1-CA_0 and of corresponding computational strength. (Simpson [28] later provided a new proof of the theorem using logical methods that avoided these lemmas and showed that the theorem itself is equivalent to ATR_0 and so strictly weaker than the lemmas both computationally and in terms of reverse mathematics.)

The situation for the theorems of Halin type that we study here is quite different. The standard proofs do not seem to use such strong methods. Nonetheless, as we mentioned above the theorems are much stronger than ACA_0 with some versions not even provable in ATR_0. We prove that these theorems occupy a few houses in the area of the reverse mathematics zoo devoted to what are called theorems (or theories) of hyperarithmetic analysis, THA (Definition 3.13). Computationally, for each computable well ordering α, there is a computable instance of any THA which has all of its required objects Turing above $0^{(\alpha)}$, the αth iteration of the Turing jump. On the other hand, they are computationally and proof theoretically much weaker than ATR_0 and so KDT. (Remember, there is a single computable graph such that the matching and cover required by KDT lies above $0^{(\alpha)}$ for all the computable well-orderings α. While for each computable instance of a THA there is a computable well-ordering α such that $0^{(\alpha)}$ computes the desired object. In our cases, the instances are graphs with arbitrarily many disjoint rays and the desired object is an infinite sequence of disjoint rays.)

Beginning with work of Kreisel [15], H. Friedman [9], Steel [30] and others in the 1960s and 1970s and continuing into the last decade (by Montalbán [16, 17, 19], Neeman [20, 21] and others), several axiomatic systems and logical theorems were found to be THA and proven to lie in a number of distinct classes in terms of proof theoretic complexity. Until now, however, there has been only one mathematical but not logical example, i.e. one not mentioning classes of first order formulas or their syntactic complexity. This was a result

(INDEC) about indecomposability of linear orderings in Jullien's thesis [14] (see Rosenstein [24, Lemma 10.3]). It was shown to be a THA by Montalbán [16].

The natural quest then became to find out if there are any other THA in the standard mathematical literature. The issue was raised explicitly in Montalbán's "Open Questions in Reverse Mathematics" [18, Q30]. As our answer, we provide many examples. Most of them are provable in a well known system above ACA_0 gotten by adding on a weak form of the axiom of choice (Σ^1_1-AC_0).

Several of the basic Halin type theorems (Definition 6.1) have versions like those appearing in the original papers that show that there are always families of disjoint rays of maximal cardinality which are of the same computational strength as the versions described above (Proposition 6.3 and Corollary 6.4). On the other hand, they are strictly stronger proof theoretically because they imply more induction than is available in Σ^1_1-AC_0 (Theorem 6.8 and Corollary 6.9). Two of the variations we consider are as yet open problems of graph theory. We show that if we restrict the class of graphs (to directed forests) the principles are not only provable but reverse mathematically equivalent to Σ^1_1-$\mathsf{AC}_0 + I\Sigma^1_1$. Note that as $\mathsf{ATR}_0 \nvdash I\Sigma^1_1$, these theorems are not provable even in ATR_0 or from KDT (Corollary 6.14). We do not know of other mathematical but nonlogical theorems of this strength. Other versions that require maximal sets of rays are much stronger and, in fact, equivalent to Π^1_1-CA_0 (Theorem 6.18).

2. Outline of Paper

Section 3 contains basic definitions and background information. The first subsection (3.1) provides the concepts that we need from graph theory. Almost all the definitions are standard. At times we give slight variations that are equivalent to the standard ones but make dealing with the computability and proof theoretic analysis easier. We also state the theorems of Halin and some variants that are the main targets of our analysis.

The second subsection (3.2), assumes an intuitive view of computability of functions $f : \mathbb{N} \to \mathbb{N}$ such as having an algorithm given by a program in any standard computer language. It then supplies all the definitions and a few standard theorems needed to follow our analysis of the computational complexity of the graph theoretic theorems we study. In particular, it defines the Turing jump operator and its iterations along countable well orderings. These are our primary computational measuring rods. The final subsection (3.3) provides the syntax and semantics for the formal systems of arithmetic that are used to measure proof theoretic complexity. It also describes the standard basic axiomatic systems and their connections to the computational measures of the previous subsection. It includes the formal definition of the class of theorems which includes most of our graph theoretic examples, the THA, Theorems (or Theories) of Hyperarithmetic Analysis. These are defined in terms of the transfinite iterations of the Turing jump and the hyperarithmetic sets of the previous subsection. In addition it defines Σ^1_1-AC_0 a weak version of the axiom of choice that is an early well known example of such theories which plays a crucial role in our analysis.

Section 4 provides the proof that Halin's original theorem IRT (Definition 3.4) is computationally very complicated. For example, given any iteration $0^{(\alpha)}$ of the Turing jump, there is a computable graph satisfying the hypotheses of IRT such that any instance of its conclusion computes $0^{(\alpha)}$. Indeed, IRT is a THA. At times, theorems or lemmas are stated in terms of the formal systems of §3.3, but the proofs rely only on the computational notions of §3.2.

Section 5 studies several variations $\mathsf{IRT}_{\mathrm{XYZ}}$ of Halin's IRT where we consider directed as well as undirected graphs, edge rather than vertex disjointness for the rays and double as well as single rays. (See Definitions 3.1 and 3.2 and the discussion after Definition 3.4.) We provide reductions over RCA_0 between many of the pairs of the eight possible variants. The proofs of these reductions proceed purely combinatorially by providing one computational process that takes an instance of some $\mathsf{IRT}_{\mathrm{XYZ}}$, i.e. a graph satisfying its hypotheses, and produces a graph satisfying the hypotheses of another $\mathsf{IRT}_{\mathrm{X'Y'Z'}}$ and another computable process that takes any solution to the $\mathsf{IRT}_{\mathrm{X'Y'Z'}}$ instance, i.e. any sequence of rays satisfying the conclusion of $\mathsf{IRT}_{\mathrm{X'Y'Z'}}$, and produces a solution to the original instance of $\mathsf{IRT}_{\mathrm{XYZ}}$. (See Propositions 5.3, 5.5 and 5.7 and the associated Lemmas. An additional reduction using a stronger base theory is given in the next section (Theorem 6.15).)

We then show that five of the eight possible variants of IRT are THA (Theorem 5.1). Of the remaining three, two are still open problems in graph theory but we do have an analysis of their restrictions to special classes of graphs in Theorem 5.16 and §6.1. The last of the variations, $\mathsf{IRT}_{\mathrm{UED}}$, has been proven more recently by Bowler, Carmesin, Pott [3] using more sophisticated methods than the other results. We have some lower bounds (Theorem 5.9) but we have yet to fully analyze the complexity of their construction.

In the next section (§6) we study some variations of IRT that ask for different types of maximality for the solutions. The first sort actually follow the original formulation of IRT in Halin [11]: IRT^*: In any graph there is a set of disjoint rays of maximum cardinality. For uncountable graphs this amounts to a basic counting argument on uncountable cardinals as all rays are countable. When restricted to countable graphs this is easily seen to be equivalent to our more modern formulation by induction. Technically, the induction used is for Σ^1_1 formulas ($I\Sigma^1_1$) which is not available in RCA_0. More specifically we show (Proposition 6.3) that $\mathsf{IRT}_{\mathrm{XYZ}} + I\Sigma^1_1$ and $\mathsf{IRT}^*_{\mathrm{XYZ}} + I\Sigma^1_1$ are equivalent (over RCA_0). As the definition of THA only depends on standard models where full induction holds, if $\mathsf{IRT}_{\mathrm{XYZ}}$ is a THA then so is $\mathsf{IRT}^*_{\mathrm{XYZ}}$.

We then prove that these maximal cardinality variants $\mathsf{IRT}^*_{\mathrm{XYZ}}$ are strictly stronger proof theoretically than the basic $\mathsf{IRT}_{\mathrm{XYZ}}$ (when they are known to be provable in Σ^1_1-AC_0). This is done by showing (Theorem 6.8 and the Remark that follows it) that the relevant $\mathsf{IRT}^*_{\mathrm{XYZ}}$ all imply weaker versions of $I\Sigma^1_1$ that are analogous to the restrictions of Σ^1_1-AC_0 embodied in weak (or unique)-Σ^1_1-AC_0 and finite-Σ^1_1-AC_0 (Definitions 7.1 and 7.2). In all the cases, it is enough induction to prove (with the apparatus of the basic $\mathsf{IRT}_{\mathrm{XYZ}}$) the consistency of Σ^1_1-AC_0 and so by Gödel's second incompleteness theorem they cannot be proved in Σ^1_1-AC_0 (Corollary 6.9).

As for proving full Σ^1_1 induction from an $\mathsf{IRT}^*_{\mathrm{XYZ}}$ we are in much the same situation mentioned above for Σ^1_1-AC_0 and $\mathsf{IRT}_{\mathrm{XYZ}}$. In particular, $\mathsf{IRT}^*_{\mathrm{DVD}}$ and $\mathsf{IRT}^*_{\mathrm{DED}}$ for directed forests each proves $I\Sigma^1_1$ as well as Σ^1_1-AC_0 (Theorems 6.12 and 6.13) and so are equivalent to $I\Sigma^1_1 + \Sigma^1_1$-AC_0. As before, this shows that they are strictly stronger than Σ^1_1-AC_0 (Corollary 6.14). Indeed, as mentioned at the end of §1, they are not even provable in ATR_0. We do not know of any other mathematical theorems with this level of reverse mathematical strength.

The second variation of maximality, $\mathsf{MIRT}_{\mathrm{XYZ}}$, studied in §6.2 is also mentioned in the original Halin paper [11]. It asks for a set of disjoint rays which is maximal in the sense of set containment. Of course, this follows immediately from Zorn's Lemma for all graphs. For countable graphs we provide a reverse mathematical analysis, showing that each of the $\mathsf{MIRT}_{\mathrm{XYZ}}$ is equivalent to Π^1_1-CA_0 (Theorem 6.18).

In §7, we discuss the reverse mathematical relationships between the THAs associated with variations of Halin's theorem and previously studied THAs as well as one new logical one (finite-Σ^1_1-AC_0 of Definition 7.1). Basically, all the $\mathsf{IRT}^*_{\mathrm{XYZ}}$ (and so $\mathsf{IRT}_{\mathrm{XYZ}} + I\Sigma^1_1$) imply H. Friedman's ABW_0 (Definition 7.4) by Theorem 7.7 and finite-Σ^1_1-AC_0 (Theorem 7.3). On the other hand, none of them are implied by it (Theorem 7.10) or by Δ^1_1-CA_0 (Definition 7.8 and Theorem 7.9). $\mathsf{ABW}_0 + I\Sigma^1_1$ does, however, imply finite-Σ^1_1-AC_0 which is not implied by weak (unique)-Σ^1_1-AC_0 (Goh [10]). Figure 4 here summarizes the known relations with references.

In the penultimate section (§8) we study the only use of Σ^1_1-AC_0 in each of our proofs of $\mathsf{IRT}_{\mathrm{XYZ}}$. It consists of $\mathsf{SCR}_{\mathrm{XYZ}}$ which says we can go from the hypothesis that there are arbitrarily many disjoint rays to a sequence $(X_k)_k$ each of which is a sequence of k many disjoint rays. We analyze the strength of the $\mathsf{SCR}_{\mathrm{XYZ}}$ and the weakenings $\mathsf{WIRT}_{\mathrm{XYZ}}$ of $\mathsf{IRT}_{\mathrm{XYZ}}$ which each take the existence of such a sequence $(X_k)_k$ as its hypothesis in place of there being arbitrarily many disjoint rays. For example, for all the $\mathsf{IRT}_{\mathrm{XYZ}}$ which are consequences of Σ^1_1-AC_0 and so are THAs, $\mathsf{IRT}_{\mathrm{XYZ}}$ is equivalent to $\mathsf{SCR}_{\mathrm{XYZ}}$ over RCA_0 (Corollary 8.5) and so all of them are also THAs. For the same choices of XYZ, ACA_0 proves $\mathsf{WIRT}_{\mathrm{XYZ}}$ over RCA_0. While a natural strengthening of $\mathsf{WIRT}_{\mathrm{XYZ}}$ does imply ACA_0 and indeed is equivalent to it (Theorem 8.9), we do not know if $\mathsf{WIRT}_{\mathrm{XYZ}}$ itself implies ACA_0. All we can prove is that it is not a consequence of RCA_0 (Theorem 8.10).

In the last section (§9), we mention some open problems.

3. Basic Notions and Background

We begin with basic notions and terminology from graph theory. At times we use formalizations that are clearly equivalent to more standard ones but are easier to work with computationally or proof-theoretically. The following two subsections supply background and basic information about the standard computational and logical/proof theoretic notions that we use here to measure the complexity of the graph theorems and constructions that we analyze in the rest of this paper.

3.1. **Graph Theoretic Notions.**

Definition 3.1. A *graph* H is a pair $\langle V, E\rangle$ consisting of a set V (of *vertices*) and a set E of unordered pairs $\{u, v\}$ with $u \neq v$ from V (called *edges*). These structures are also called *undirected graphs (or here U-graphs)*. A structure H of the form $\langle V, E\rangle$ as above is a *directed graph (or here D-graph)* if E consists of ordered pairs $\langle u, v\rangle$ of vertices with $u \neq v$. To handle both cases simultaneously, we often use X to stand for undirected (U) or directed (D). We then use (u, v) to stand for the appropriate kind of edge, i.e. $\{u, v\}$ or $\langle u, v\rangle$.

An *X-subgraph* of the X-graph H is an X-graph $H' = \langle V', E'\rangle$ such that $V' \subseteq V$ and $E' \subseteq E$. It is an *induced X-subgraph* if $E' = \{(u, v) | u, v \in V' \ \&\ (u, v) \in E\}$.

Definition 3.2. An *X-ray in H* is a pair consisting of an X-subgraph $H' = \langle V', E'\rangle$ of H and an isomorphism $f_{H'}$ from $\mathbb{N}$ with edges $(n, n+1)$ for $n \in \mathbb{N}$ to H'. We say that the *ray begins at* $f(0)$. We also describe this situation by saying that H contains the X-ray $\langle H', f_{H'}\rangle$. We sometimes abuse notation by saying that the sequence $\langle f(n)\rangle$ of vertices is an X-ray in H. Similarly we consider *double X-rays* where the isomorphism $f_{H'}$ is from $\mathbb{Z} = \{-n, n | n \in N\}$ with edges $(z, z+1)$ for $z \in \mathbb{Z}$. We use Z-ray to stand for either a (single) ray ($Z = S$) or double ray ($Z = D$) and so we have, in general, Z-X-rays or just Z-rays if the type of graph (U or D) is already established. For brevity, when we describe rays we will often only list their vertices in order instead of defining H' and f explicitly. However the reader should be aware that we always have H' and f in the background.

H contains k many Z-X-rays for $k \in N$ if there is a sequence $\langle H_i, f_i \rangle_{i<k}$ such that each $\langle H_i, f_i \rangle$ is a Z-X-ray in H (with $H_i = \langle V_i, E_i \rangle$).

H contains k many disjoint (or vertex-disjoint) Z-X-rays if the V_i are pairwise disjoint. *H contains k many edge-disjoint Z-X-rays* if the E_i are pairwise disjoint. We often use Y to stand for either vertex (V) or edge (E) as in the following definitions.

An X-graph H *contains arbitrarily many Y-disjoint Z-X-rays* if it contains k many such rays for every $k \in \mathbb{N}$.

An X-graph H *contains infinitely many Y-disjoint Z-X-rays* if there is an X-subgraph $H' = \langle V', E' \rangle$ of H and a sequence $\langle H_i, f_i \rangle_{i \in N}$ such that each $\langle H_i, f_i \rangle$ is a Z-X-ray in H (with $H_i = \langle V_i, E_i \rangle$) such that the V_i or E_i, respectively for $Y = V, E$, are pairwise disjoint and $V' = \cup V_i$ and $E' = \cup E_i$.

Definition 3.3. An *X-path P* in an X-graph H is defined similarly to single rays except that the domain of f is a proper initial segment of $\mathbb{N}$ instead of $\mathbb{N}$ itself. Thus they are finite sequences of distinct vertices with edges between successive vertices in the sequence. If $P = \langle x_0, \ldots, x_n \rangle$ is a path, we say it is a *path of length n* between x_0 and x_n. Our *notation for truncating and combining paths* $P = \langle x_0, \ldots, x_n \rangle$, $Q = \langle y_0, \ldots, y_m \rangle$ and $R = \langle z_0, \ldots, z_l \rangle$ is as follows: $x_i P = \langle x_i, \ldots, x_n \rangle$, $P x_i = \langle x_0, \ldots, x_i \rangle$, and we use concatenation in the natural way, e.g., if the union of Px, xQy and yR is a path, we denote it by $PxQyR$. We treat rays as we do paths in this notation, as long as it makes sense, writing, for example, $x_i R$ for the ray which is gotten by starting R at an element x_i of R; Rx_i is the path which is the initial segment of R ending in x_i and we use concatenation as for paths as well.

The starting point of the work in this paper is a theorem of Halin [11] that we call the infinite ray theorem as expressed in Diestel [5, Theorem 8.2.5(i)].

Definition 3.4 (Halin's Theorem)**.** IRT*, the infinite ray theorem,* is the principle that every graph H which contains arbitrarily many disjoint rays contains infinitely many.

The versions of Halin's theorem which we consider in this paper allow for H to be an undirected or a directed graph and for the disjointness requirement to be vertex or edge. We also allow the rays to be single or double. The corresponding versions of Halin's Theorem are labeled as $\mathsf{IRT}_{\mathrm{XYZ}}$ for appropriate values of X, Y and Z to indicate whether the graphs are undirected or directed ($X = U$ or D); whether the disjointness refers to the vertices or edges ($Y = V$ or E) and whether the rays are single or double ($Z = S$ or D), respectively, in the obvious way. We often state a theorem for several or all XYZ and then in the proof use "graph", "edge" and "disjoint" unmodified with the intention that the proof can be read for any of the cases. This is convenient for minimizing repetition in some of our arguments.

We will also consider restrictions of these theorems to specific families of graphs. We need a few more notions to define them.

Definition 3.5. A *tree* is a graph T with a designated element r called its *root* such that for each vertex $v \neq r$ there is a unique path from r to v. A *branch* in T is a ray that begins at its root. We denote the set of its branches by $[T]$ and say that T is *well-founded* if $[T] = \emptyset$ and otherwise it is *ill-founded.* A *forest* is an *effective disjoint union* of trees, or more formally, a graph with a designated set R (of vertices called roots) such that for each vertex v there is a unique $r \in R$ such that there is a path from r to v and, moreover, there is only one such path. (Here we conflate the set of roots R with the usual indexing set for a typical effective disjoint union so that we continue to have a designated set of the roots.) In general, the *effectiveness* we assume when we take *disjoint unions of graphs* means that we can effectively (i.e. computably) identify each vertex in the union with the original vertex (and the graph to which it belongs) which it represents in the disjoint union.

A *directed tree* is a directed graph $T = \langle V, E\rangle$ such that its *underlying graph* $\hat{T} = \langle V, \hat{E}\rangle$ where $\hat{E} = \{\{u, v\} | \langle u, v\rangle \in E \vee \langle v, u\rangle \in E\}$ is a tree. A *directed forest* is a directed graph whose underlying graph is a forest.

Definition 3.6. An X-graph H is *locally finite* if, for each $u \in V$, the set $\{v \in E | (u, v) \in E \vee (v, u) \in E\}$ of *neighbors of* u is finite. A locally finite X-tree is also called *finitely branching*. (Note this does not mean there are finitely many branches in the tree.)

Of course, there are many well known equivalent definitions of trees and associated notions. We have given one possible set of graph theoretic ones. In the case of undirected graphs our definition is equivalent to all the standard ones. Readers are welcome to think in terms of their favorite definition. Note, however, we are restricting ourselves to what would (in set theory) be called countable trees with all nodes of finite rank. Thus, we typically think of trees as *subtrees of* $\mathbb{N}^{<\mathbb{N}}$, i.e. the sets of finite strings of numbers (as vertices) with an edge between σ and τ if and only if they differ by one being an extension of the other by one element, e.g. $\sigma^\frown k = \tau$.

It does not seem as if there is a single standard definition for directed graphs being directed trees. We have picked one that seems to be at least fairly common and works for the only situations for which we consider them in Theorems 5.16, 6.12, and 6.13 and Corollary 6.14.

3.2. **Computability Hierarchies.** While we may cite results about uncountable graphs, all sets and structures actually studied in this paper will be countable. Thus for purposes of defining their complexity, we can think of all of them as being subsets of, or relations or functions on, $\mathbb{N}$.

We do not give a formal stand alone definition of computability for sets or functions but assume an at least intuitive grasp of some model of computation such as by a Turing or Register machine that has unbounded memory and is allowed to run for unboundedly many steps. (We do provide in §3.3 a definition via definability in arithmetic that is equivalent to the formal versions of machine model definitions.) Thus we say a function $f : \mathbb{N} \to \mathbb{N}$ is computable if there is a program for one of these machines that computes $f(n)$ as output when given input n. A set X is computable if its characteristic function $\mathbb{N} \to \{0, 1\}$ is computable. Note that as the alphabets or our languages are finite, there are only countably many programs and as our formation rules are effective, we have a computable list of the programs and hence one, Φ_e, of the partial functions they compute. (They are only partial as, of course, some programs fail to halt on some inputs.)

Fundamental to measuring the relative computational complexity of sets or functions is the notion of machines with oracles and Turing reduction. Given a set X or function f we consider machines augmented by the ability to produce $X(n)$ or $f(n)$ if it has already produced n. We say that such a machine is one with an oracle for X or f. We then say that X is *computable from (or Turing reducible to)* Y if there is a machine with oracle Y which computes X via some *reduction* Φ_e^X. We write this as $X \leq_T Y$. We say X is of the same (Turing) degree as Y, $X \equiv_T Y$, if $X \leq_T Y$ and $Y \leq_T X$. We use all the same terminology and notations for functions.

The first level beyond the computable in our basic hierarchy of computable complexity is given by the halting problem $H = \{e | \Phi_e(e) \text{ converges}\}$ that is H is the set of e such that the computation of the eth machine Φ_e on input e eventually halts. We then define an operator on sets $X \longmapsto X' = \{e | \Phi_e^X(e) \text{ converges}\}$ that is X' is the set of e such that the computation of the eth machine with oracle X, Φ_e^X on input e eventually halts. (It is easy to see that $H \equiv_T \emptyset'$.) The crucial fact here is the undecidability of the halting problem (for every oracle), i.e. for every X, X' is strictly above X in terms of Turing computability. The other basic fact that we need about $\emptyset'$ is that it is *computably enumerable,* i.e. there is a computable function f whose range is $\emptyset'$. If $f(s) = x$ we say that x is *enumerated in*, or *enters*, $0'$ at (stage) s. If we view H as defined by using the empty oracle $\emptyset$, the procedure that takes us from the halting problem to the Turing jump by replacing $\emptyset$ as oracle by X is an instance of a general procedure called *relativization.* It takes any computable function or proof about computable functions or degrees (i.e. ones with oracle $\emptyset$) to the same function, or proof about functions, computable in X (or degrees above that of X). Almost always this procedure trivially transforms correct proofs with oracle $\emptyset$ to ones with arbitrary oracle X. Typically, this transformation keeps the same programs doing the required work with any oracle. For example, X' is *computably enumerable in* X (or relative to X), i.e. there is a function Φ_e^X whose range is X' and this can be taken to be the same e such that $\Phi_e^\emptyset$ enumerates $\emptyset'$. We also use X'_s to denote *the set of numbers enumerated in (or that have entered)* X' *by stage* s. This phenomena of the procedure or result not depending on the particular oracle or depending in a fixed computable way on some other parameters is described as it being *uniform* in the oracle or other parameters. We describe an important example of uniformity in Remark 3.9.

We can now generate a hierarchy of computational complexity by iterating the jump operator beginning with any set X: $X^{(0)} = X$; $X^{(n+1)} = (X^n)'$. While the finite iterations of the jump capture most construction techniques and theorems in graph theory (and most other areas of classical countable/separable mathematics), we will be interested in ones that go beyond such techniques and proofs. The basic idea is that we continue the hierarchy by iteration into the transfinite while still tying the iteration to computable procedures.

Definition 3.7. We represent well-orderings or *ordinals* α as well-ordered relations on $\mathbb{N}$. Typically such *ordinal notations* are endowed with various additional structure such as identifying 0, successor and limit ordinals and specifying cofinal ω-sequences for the limit ordinals. If we have a representation of α then restricting the well-ordering to numbers in its domain provides representations of each ordinal $\beta < \alpha$. We generally simply work with ordinals and omit concerns about translating standard relations and procedures to the representation. An ordinal is computable (in a set X) if it has a computable (in X) representation. For a set X and ordinal (notation) α computable from X, we define the transfinite iterations $X^{(\beta)}$ of the Turing jump of X by transfinite induction on $\beta \leq \alpha$: $X^{(0)} = X$; $X^{(\beta+1)} = (X^\beta)'$ and for a limit ordinal λ, $X^{(\lambda)} = \oplus\{X^{(\beta)} | \beta < \lambda\} = \cup\{\beta \times X^{(\beta)} | \beta < \lambda\}$ (or as the effective disjoint sum over the $X^{(\beta)}$ in the specified cofinal sequence in λ).

Definition 3.8. HYP(X), the collection of all sets *hyperarithmetic in* X consists of those sets computable in some $X^{(\alpha)}$ for α an ordinal computable in X. We say that Y is *hyperarithmetic in* X or *hyperarithmetically reducible to* X, $Y \leq_h X$ if $Y \in \text{HYP}(X)$.

These sets too, will be characterized by a definability class in arithmetic in §3.3. For now we just note that they clearly go far beyond the sets computable from the finite iterations of the jump.

The computational strength of our graph theoretic theorems such as IRT is measured by this hierarchy as we will show that, for every set X and every set Y hyperarithmetic in X, there is a graph G computable from X which satisfies the hypotheses of IRT but for which any collection of rays satisfying its conclusion computes Y. On the other hand, placing an upper bound on the strength of IRT requires analyzing its proof and the principles used in it. The relevant one is a form of the axiom of choice. We define it in the next subsection along with a general class of such principles, the theorems/theories of hyperarithmetic analysis which are, computationally, the primary objects of our analysis in this paper.

We note one important well known basic fact relating the jumps of X to trees computable from X. We will need it for our proofs that IRT and its variants are computationally complex enough to compute all the sets hyperarithmetic in any given set X (as the instances of the graphs range over graphs computable from X).

Remark 3.9. For any set X and any ordinal α computable from X, there is a sequence $\langle T_\beta | \beta < \alpha \rangle$ computable from X of trees (necessarily) computable from X such that each tree has exactly one branch P_β and P_β is of the same complexity as $X^{(\beta)}$, i.e. $P_\beta \equiv_T X^{(\beta)}$. The procedure for computing this sequence is uniform in X and the index for the program computing the well ordering α from X, i.e. there is one computable function that when given an oracle for X, an *index for* α (i.e. the i such that Φ_i^X is the well ordering α) and a β in the ordering, computes the whole sequence $\langle T_\beta | \beta < \alpha \rangle$ and the indices for the reductions between P_β and $X^{(\beta)}$. (See, e.g. [27, Theorem 2.3]). We may also easily assure that the T_β are effectively disjoint so that their union is a forest.

Some versions of the variations on IRT (see §6.2) that call for types of maximality for the infinite set of disjoint rays are stronger both computationally and proof theoretically than the $\mathsf{IRT}_{\mathrm{XYZ}}$ described above. Their computational strength is captured by a kind of jump operator that goes beyond all the hyperarithmetic ones. It captures the ability to tell if a computable ordering is a well-ordering.

Definition 3.10. The *hyperjump of* X, $\mathcal{O}^X$, is the set $\{e | \Phi_e^X$ is (the characteristic function of) a subtree of $\mathbb{N}^{<\mathbb{N}}$ which is well-founded$\}$.

This operator also corresponds to a syntactically defined level of comprehension as we note in §3.3.

3.3. **Logical and Axiomatic Hierarchies.** The basic notions from logic that we need here are those of languages, structures and axiomatic systems and proofs. As we will deal only with countable sets and structures, we can assume that we are dealing just with the natural numbers with a way to define and use sets and functions on them. Thus, at the beginning, we have in mind the natural numbers $\mathbb{N}$ along with the usual apparatus of the *language of (first order) arithmetic*, say $+, \times, <, 0$ and 1 along with the syntax of standard first order logic (the Boolean connectives $\vee, \wedge$ and $\neg$; the variables such as x and y ranging over the numbers with the usual quantifiers $\forall x$ and $\exists y$ as well as the equality relation $=$). A *structure* for this language is a set N along with elements for 0 and 1, binary functions for $+$ and $\times$ and a binary relation for $<$. We also need a way of talking about subsets of (or functions on) the numbers. We follow the standard practice in reverse mathematics of using sets and defining functions in terms of their graphs. So we expand our language by adding on new classes of (second order) variables such as X and Y and the associated quantifiers $\forall X$ and $\exists Y$ along with a new relation symbol $\in$ between numbers and sets.

A structure for this language is one of the form $\mathcal{N} = \langle N, S, +, \times, <, 0, 1, \in \rangle$ where $\langle N, +, \times, <, 0, 1\rangle$ is a structure for first order arithmetic; $S \subseteq 2^N$ is a specified collection of subsets of N, the set of "numbers" of $\mathcal{N}$, over which the second order quantifiers and variables of our language range. It is called the *second order part of* $\mathcal{N}$. The usual membership symbol $\in$ always denotes the standard membership relation between elements of N and subsets of N that are in S. So a sentence Θ *is true in* N, $\mathcal{N} \vDash \Theta$, if first order quantification is interpreted as ranging over N, second order quantification ranges over S and the relations and functions of the language are as described. This specifies the semantics for second order arithmetic.

Proof theoretic notions deal with all possible structures for the language and axiom systems to specify what we need in any particular argument. For most of our purposes and all of the computational ones, one can restrict attention to *standard models of arithmetic*, i.e. ones $\mathcal{N}$ with $N = \mathbb{N}$ and some $S \subseteq 2^{\mathbb{N}}$ with the usual interpretations of the functions and relations. We generally abbreviate these structures as $\langle \mathbb{N}, S\rangle$ with $S \subseteq 2^{\mathbb{N}}$ as all the functions and relations are then fixed.

We view the use of a second kind of variable as a short hand for a typical first order language with predicates $R^{\mathcal{N}}$ and $S^{\mathcal{N}}$ for the two parts (N and S) so that, e.g. $\exists x\phi(x)$ means $\exists x(\phi(x) \wedge R^{\mathcal{N}}(x))$ and $\exists X\phi(X)$ means $\exists X(\phi(X) \wedge S^{\mathcal{N}}(X))$. We can thus assume any standard proof theoretic system for basic first order logic. This generates the provability notion $\vdash$ used above to define our notion of logical strength and equivalences of theories (sets of sentences often called axioms) as above. We now define the standard weak base theory RCA_0 used to define the logical strength of mathematical theorems as described above. We then define a few other common systems that will be used later.

Each axiomatic subsystem of second order arithmetic that we consider contains the standard basic axioms for $+$, $\times$, and $<$ (which say that N is an ordered semiring). In addition, they all include the usual extension axiom for sets and a form of induction that applies only to sets (that belong to the model):

$$\text{(EXT)}\quad (\forall X, Y)(\forall n(n \in X \leftrightarrow n \in Y) \rightarrow X = Y).$$
$$(\mathsf{I}_0)\quad (\forall X)((0 \in X \,\&\, \forall n\,(n \in X \rightarrow n+1 \in X)) \rightarrow \forall n\,(n \in X)).$$

We call the system consisting of EXT, I_0 and the basic axioms of ordered semirings P_0. Typically axiom systems for second order arithmetic are defined by adding various types of set existence axioms to P_0 although at times additional induction axioms are used as well. In order to define them we need to specify various standard syntactic classes of formulas determined by quantifier complexity.

Definition 3.11. The Σ^0_0 and Π^0_0 formulas of second order arithmetic are just the ones with no quantifiers. Proceeding inductively, a formula Φ is Σ^0_{n+1} (Π^0_{n+1}) if it is of the form $\exists x\Psi$ ($\forall x\Psi$) where Ψ is Π^0_n (Σ^0_n). We assume some computable coding of all these formulas (viewed as strings of symbols from our language) by natural numbers. We say Φ is *arithmetic* if it is Σ^0_n or Π^0_n for some $n \in \mathbb{N}$. It is Σ^1_1 (Π^1_1) if it is of the form $\exists X\Psi$ ($\forall X\Psi$) where Ψ is arithmetic. (One can continue to define Σ^1_n and Π^1_n in the natural way but we will not need to consider such formulas here.) We say a set X is in one of these classes Γ (in or relative to Y, i.e. with Y as a parameter) if there is a formula $\Psi(n, Y) \in \Gamma$ such that $n \in X \Leftrightarrow \Psi(n, Y)$. If X is both Σ^i_n (in Y) and Π^i_n (in Y) it is called Δ^i_n (in Y).

We mention a few additional standard connections between the syntactic complexity of the definition of a set X and X's properties in terms of computability and graph theoretic notions.

Remark 3.12. The sets $Y^{(n)}$ are Σ^0_n in Y. A set X is computable (in Y) if and only if it is Δ^0_1 (in Y). More generally, it is computable in $Y^{(n)}$ iff and only if it is Δ^0_{n+1} (in Y). It is hyperarithmetic (in Y) if and only if it is Δ^1_1 (in Y). There is a computable function $f(e,n)$ such that if X is Σ^1_1 (in Y) via the Σ^1_1 formula with code e then for every n, $\Phi^Y_{f(e,n)}$ is (the characteristic function of) a tree T such that $n \in X \Leftrightarrow T$ has a branch.

The first system for analyzing the proof theoretic strength of theorems and theories in reverse mathematics is just strong enough to prove the existence of the computable sets and so supplies us with all the usual computable functions such as pairing $\langle n, m\rangle$ or more generally those coding finite sequences as numbers. In particular, it provides the predicates defining the (codes e of) the partial computable functions Φ_e and the relations saying the computation $\Phi^X_e(n)$ halts in s many steps with output y. Thus we have the basic tools to define and discuss Turing reducibility and the Turing jump. It is our weak base theory and is assumed to be included in every system we consider.

(RCA_0) Recursive Comprehension Axioms: In addition to P_0 its axioms include the schemes of recursive (generally called Δ^0_1) comprehension and Σ^0_1 induction:

$$(\Delta^1_1\text{-}\mathsf{CA}_0)\quad \forall n\,(\Phi(n) \leftrightarrow \Psi(n)) \to \exists X\,\forall n\,(n \in X \leftrightarrow \Phi(n)) \text{ for all } \Sigma^0_1 \text{ formulas } \Phi \text{ and } \Pi^0_1 \text{ formulas } \Psi \text{ in which } X \text{ is not free.}$$

$$(I\Sigma^0_1)\quad (\varphi(0)\,\&\,\forall n\,(\varphi(n) \to \varphi(n+1))) \to \forall n\,\varphi(n) \text{ for all } \Sigma^0_1 \text{ formulas } \varphi.$$

Note that these formulas may have free set or number variables. As usual, the existence assertion $\exists X....$ of the axiom is taken to mean that for each instantiation of the free variables (by numbers or sets, as appropriate, called *parameters*) there is an X as described. We take this for granted as well as the restriction that the X is not free in the rest of the formula in all of the set existence axioms of any or our systems. The standard models of this theory are just those whose second order part is closed under Turing reduction and disjoint union ($X \oplus Y = \{\langle 0, x\rangle\,|x \in X\} \cup \{\langle 1, y\rangle\,|y \in Y\}$). Historically the computable (in Y) sets which are, as mentioned above, the Δ^0_1 (in Y) sets, have also been called the recursive (in Y) sets. Hence the terminology in RCA_0.

We next move up to arithmetic comprehension.

(ACA_0) Arithmetic Comprehension Axioms: $\exists X\,\forall n\,(n \in X \leftrightarrow \Phi(n))$ for every arithmetic formula Φ.

As mentioned above the $X^{(n)}$ are defined by a Σ^0_n formula with X as a parameter. So one can show that this system is equivalent (over RCA_0) to the totality of the Turing jump operator, i.e. for every X, X' exists. Its standard models are those of RCA_0 whose second order part is also closed under Turing jump. It is also equivalent (in the sense of reverse mathematics) to König's Lemma that every finitely branching tree with paths of arbitrarily long length has a branch.

In general, we say one system of axioms S is *logically or reverse mathematically reducible* to another T over one R if $R \cup T \vdash \psi$ for every sentence $\psi \in S$. Note that S and/or T may be a single sentence or theorem. We say that S and T are equivalent over R if each is reducible to the other. If no system R is specified we assume that RCA_0 is intended.

As we will not deal with it, we have omitted the formal definition of the usual system (WKL_0) which falls strictly between RCA_0 and ACA_0. It is characterized by the restriction of König's Lemma to trees that are subsets of $2^{<N}$, the tree of finite binary strings under extension. Similarly, we omit the usual system (ATR_0) which comes after ACA_0. It says that arithmetic comprehension can be iterated along any countable well-order and so implies the existence of the sets hyperarithmetic in X for each X but is computationally stronger than this assumption.

Instead, we describe the computationally defined class of theorems/theories that are the main focus of this paper and include several variations of IRT. The definition is semantic, not axiomatic and involves only standard models. (Indeed by Van Wesep [31], there can be no axiomatic characterization of this class in second order arithmetic.)

Definition 3.13. A sentence (theory) T is a *theorem (theory) of hyperarithmetic analysis (THA)* if

(1) For every $X \subseteq \mathbb{N}$, $(\mathbb{N}, \mathrm{HYP}(X)) \vDash T$ and
(2) For every $S \subseteq 2^{\mathbb{N}}$, if $(\mathbb{N}, S) \vDash T$ and $X \in S$ then $\mathrm{HYP}(X) \subseteq S$.

The last of the standard axiomatic systems is characterized by the comprehension axiom for Π^1_1 formulas:

(Π^1_1-CA_0) The Π^1_1 comprehension axioms: $\exists X \, \forall k \, (k \in X \leftrightarrow \Phi(k))$ for every Π^1_1 formula $\Phi(k)$.

Remark 3.14. The hyperjump, $\mathcal{O}^X$, is clearly a Π^1_1 set with parameter X. In fact, every Π^1_1 set with parameter X is reducible to $\mathcal{O}^X$: There is a computable function $f(e,n)$ such that for every index e for a Π^1_1 formula $\Psi(n)$ with parameter X and every n, $\Psi(n) \Leftrightarrow f(e,n) \in \mathcal{O}$. Thus Π^1_1-CA_0 corresponds to closure under the hyperjump. We will see it appear as equivalent to a version of IRT where we ask for a maximal set of disjoint rays in Theorem 6.18.

For this paper, the most important other existence axiom is a restricted form of the axiom of choice.

(Σ^1_1-AC_0) $\forall n \exists X \, \Phi(n,X) \to \exists X \forall n \Phi(n, X^{[n]})$ where Φ is arithmetic and $X^{[n]} = \{m | \langle n,m \rangle \in X\}$ is the *nth column of* X.

This axiom falls strictly between ACA_0 and ATR_0 and is well known to be a THA (essentially in Kreisel [15]). It plays a crucial role in our analysis because we provide the upper bound on the strength of most of our theorems by showing that they follow from Σ^1_1-AC_0. This provides the computational upper bound for being a THA as any consequence of a THA must satisfy Definition 3.13(1) as Σ^1_1-AC_0 is true in $\mathrm{HYP}(X)$. Thus the bulk of our proofs for the computational complexity of the theorems we study consist of showing that they imply Definition 3.13(2), i.e. closure under "hyperarithmetic in".

Over the past fifty years, several other logical axioms have been shown to be THA. We will discuss some of them in §7. However, as we discussed in §1, only one somewhat obscure purely mathematical theorem was previously known to be a THA. We provide several more in this paper (Theorem 5.1, Corollary 6.4 and Theorem 6.13). We also introduce a new logical axiom, finite-Σ^1_1-AC_0 (Definition 7.2) which is a THA as well.

For those interested in the proof theory and so nonstandard models, we also at times explicitly consider the induction axiom at the same Σ^1_1 level.

($I\Sigma^1_1$) $(\Phi(0) \ \wedge \ \forall n(\Phi(n) \to \Phi(n+1))) \ \to \forall n \Phi(n)$ for every Σ^1_1 formula Φ.

This axiom does not imply the existence of any infinite sets and is, of course, true in every standard model. Thus the readers interested only in the computational complexity of the Halin type theorems can safely ignore these considerations.

4. IRT and Hyperarithmetic Analysis

We devote this section to the proof of

Theorem 4.1. IRT *is a theorem of hyperarithmetic analysis.*

In this section we consider only vertex-disjoint single rays in undirected graphs, as in the statement of IRT. The proof of Theorem 4.1 will be split into two parts. The first part is

Theorem 4.2. *Every standard model of* $\mathsf{RCA}_0 + \mathsf{IRT}$ *is closed under hyperarithmetic reduction.*

Proof. Fix a standard model $\mathcal{M}$ of $\mathsf{RCA}_0 + \mathsf{IRT}$. First, we show that $\mathcal{M}$ contains $\emptyset'$. By relativizing the proof, it follows that $\mathcal{M}$ is closed under Turing jump.

For each n, consider the tree $T_n \subseteq \mathbb{N}^{<\mathbb{N}}$ consisting of all strings of the form $s^\frown 0^t$ such that some number below n is enumerated into $\emptyset'$ at stage s, and either $t \leq s$ or $\emptyset'_s \upharpoonright n = \emptyset'_t \upharpoonright n$. Observe that T_n has a unique computable branch $s^\frown 0^\infty$, where s is the smallest number such that $\emptyset' \upharpoonright n = \emptyset'_s \upharpoonright n$.

Consider the disjoint union $\bigsqcup_n T_n$. Observe that $\bigsqcup_n T_n$ satisfies the premise of IRT (in $\mathcal{M}$), because each T_n has a computable branch. Apply IRT to $\bigsqcup_n T_n$ to obtain a sequence $(R_i)_i$ of disjoint rays in $\bigsqcup_n T_n$. Each R_i is contained in some T_n. We can, uniformly in i, extend or truncate R_i to the unique branch P_n of T_n. Hence $(R_i)_i$ computes a sequence of infinitely many distinct branches P_n, which in turn computes longer and longer initial segments of $\emptyset'$. This proves that $\mathcal{M}$ contains $\emptyset'$.

Next, we show that if $\mathcal{M}$ contains $\emptyset^{(\alpha)}$ for each $\alpha < \lambda$, then $\mathcal{M}$ contains $\emptyset^{(\lambda)}$. (Again the desired result follows by relativization.) By Remark 3.9, there is a computable sequence $(T_\beta)_{\beta<\lambda}$ of trees such that each tree has exactly one branch $P_\beta \equiv_T \emptyset^{(\beta)}$. Fix an increasing sequence $(\alpha_n)_n$ which is cofinal in λ and consider the disjoint union $\bigsqcup_n T_{\alpha_n}$. Observe that $\bigsqcup_n T_{\alpha_n}$ satisfies the premise of IRT (in $\mathcal{M}$): for each n, $\emptyset^{(\alpha_n)}$ computes the branches P_{α_m} for $m \leq n$. Apply IRT to $\bigsqcup_n T_{\alpha_n}$ to obtain a sequence $(R_i)_i$ of disjoint rays in $\bigsqcup_n T_{\alpha_n}$. As before, $(R_i)_i$ computes a sequence of infinitely many distinct branches P_{α_n}, and hence a sequence of infinitely many distinct $\emptyset^{(\alpha_n)}$. Each $\emptyset^{(\alpha_n)}$ uniformly computes $\emptyset^{(\alpha_m)}$ for $m \leq n$, so we conclude that $(R_i)_i$ computes $\bigoplus_m \emptyset^{(\alpha_m)}$ as desired. $\square$

It follows that IRT is not provable in ACA_0, despite the apparent similarity between IRT and a compactness result. Later we will show that IRT does imply ACA_0 (Proposition 4.6) but is not even provable in the theory of Δ^1_1-comprehension (Theorem 7.9).

Next, we present the proof of IRT attributed to Andreae (see Diestel [5, Theorem 8.2.5 and bottom of pg. 275]), with emphasis on the axioms which can be used to formalize it. We will then use this analysis to complete the proof of Theorem 4.1.

Given a graph which has arbitrarily many disjoint rays, we will build approximations $R^n_0, \dots, R^n_{n-1}$ to the desired infinite sequence of disjoint rays by induction. For each n, $R^n_0, \dots, R^n_{n-1}$ is a sequence of disjoint rays. As n grows, R^n_i and R^{n+1}_i will agree on a growing initial segment. Hence for each i, $\lim_n R^n_i$ exists, and $\langle \lim_n R^n_i : i < \mathbb{N} \rangle$ will be an infinite sequence of disjoint rays.

The key combinatorial lemma implicit in Andreae's proof is:

Lemma 4.3. *Given n disjoint rays $R_0, \dots, R_{n-1}$ and n^2+1 disjoint rays $S_0, \dots, S_{n^2}$, there are $n+1$ disjoint rays $R'_0, \dots, R'_n$ such that for each $i < n$, R_i and R'_i start at the same vertex.*

On the face of it, constructing such $R'_0, \dots, R'_n$ could be difficult; perhaps as difficult as constructing a solution to a Σ^1_1 predicate. However, Andreae's proof actually constructs $R'_0, \dots, R'_n$ such that for each $i \leq n$, R'_i shares a tail with some R_j or S_j. We will see that this lowers the complexity of constructing such rays considerably.

Before proving Lemma 4.3, let us use it to prove IRT.

Proof of IRT *assuming Lemma 4.3.* Suppose that we have constructed disjoint rays $R^n_0, \dots, R^n_{n-1}$. We want to construct disjoint rays $R^{n+1}_0, \dots, R^{n+1}_n$ such that for each $i < n$, R^n_i and R^{n+1}_i agree on their first $n - i$ vertices. By assumption, let $S_0, \dots, S_{f(n)}$ be a finite sequence of disjoint rays, where $f(n) = n^2 + \sum_{i<n}(n-i)$.

Next, discard all rays S_j which intersect the first $n-i$ vertices of some R^n_i, if any. (By the way that we have formalized rays, this can be done in RCA_0.) There are at most $\sum_{i<n}(n-i)$ many of them, so by discarding and renumbering if necessary we are left with $S_0, \dots, S_{n^2}$.

We are ready to apply Lemma 4.3. For each $i < n$, let x_i denote the $(n-i)^{\text{th}}$ vertex on R^n_i. Apply Lemma 4.3 to $x_0R^n_0, \dots, x_{n-1}R^n_{n-1}$ and $S_0, \dots, S_{n^2}$. We obtain $n+1$ disjoint rays $R'_0, \dots, R'_n$ such that for each $i < n$, R'_i begins at vertex x_i. Then, define $R^{n+1}_i = R^n_i x_i R'_i$ for $i < n$ and define $R^{n+1}_n = R'_n$. This completes the inductive step of the proof of IRT. □

It remains to prove Lemma 4.3. The key ingredient is Menger's theorem for finite graphs (see [5, Theorem 3.3.1]). If A and B are disjoint sets of vertices in a graph, we say that P is an *A-B path* if P starts with some vertex in A and ends with some vertex in B. A set of vertices S *separates A and B* if any A-B path passes through at least one vertex in S.

Theorem 4.4 (Menger)**.** *Let G be a finite graph. If A and B are disjoint sets of vertices in G, then the minimum size of a set of vertices which separate A and B is equal to the maximum size of a set of disjoint A-B paths.*

It is easy to see that Menger's theorem is provable in ACA_0 (certainly even less is needed). We now show that Lemma 4.3 is provable in ACA_0.

Proof of Lemma 4.3. Suppose we are given n disjoint rays $R_0, \dots, R_{n-1}$ and n^2+1 disjoint rays $S_0, \dots, S_{n^2}$. First, use ACA_0 to define the set

$$\{\langle i, q\rangle : R_i \text{ intersects } S_q\}.$$

Then we perform the following recursive procedure. At each step, check if there is some $i < n$ such that R'_i has not been defined and R_i intersects at most n many rays S_q which have not been discarded. If there is no such i, we end the procedure. Otherwise, find the least such i and do the following:

(1) discard all rays S_q which intersect R_i;
(2) define $R'_i = R_i$.

After the procedure is complete, let I be the set of $i < n$ for which R'_i has not been defined. Let $\mathcal{R}$ be the set of rays S_q which have not been discarded. Let $m = |I|$. We observe that $|\mathcal{R}| \geq m^2 + 1$, because

$$(n^2+1) - (n-m)n = mn + 1 \geq m^2 + 1.$$

Next, for each $i \in I$, let z_i be the first vertex on R_i such that $R_i z_i$ meets exactly m many rays in $\mathcal{R}$. (Each z_i exists by construction of I. By RCA_0, $\{z_i : i \in I\}$ exists.)

Observe that the finite set $\bigcup_{i\in I} R_i z_i$ meets at most m^2 many rays in $\mathcal{R}$. Since $|\mathcal{R}| \geq m^2+1$, we may pick some ray in $\mathcal{R}$ which does not meet $\bigcup_{i\in I} R_i z_i$. We define R'_n to be said ray. Then, discard all rays in $\mathcal{R}$ which do not meet $\bigcup_{i\in I} R_i z_i$.

Finally, we use Menger's theorem (Theorem 4.4) to define R'_i for each $i \in I$. For each $i \in I$, let x_i denote the first vertex of R_i. For each q such that S_q remains in $\mathcal{R}$, let y_q be the first vertex on S_q such that $y_q S_q$ and $\bigcup_{i\in I} R_i z_i$ are disjoint. Then consider the following finite sets of vertices (which exist, by ACA_0):

$$\begin{aligned} X &= \{x_i : i \in I\} \\ Y &= \{y_q : S_q \in \mathcal{R}\} \\ H &= \bigcup_{i\in I} R_i z_i \cup \bigcup_{S_q \in \mathcal{R}} S_q y_q. \end{aligned}$$

We want to apply Menger's theorem to $X, Y \subseteq H$. Towards that end, we claim that X cannot be separated from Y in H by fewer than m vertices.

Suppose that $A \subseteq H$ and $|A| < m$. Since $|I| = m$ and $\{R_i : i \in I\}$ is disjoint, there is some $i \in I$ such that R_i does not meet A. Next, since $R_i z_i$ meets m many disjoint rays in $\mathcal{R}$, there is some q such that $S_q \in \mathcal{R}$ and $R_i z_i$ meets S_q, but S_q does not meet A. Let z be any vertex in both $R_i z_i$ and S_q. Then $R_i z S_q y_q$ is a path in H from x_i to y_q which does not meet A. This proves our claim.

By Menger's theorem, there are m many disjoint X-Y paths in H. Then, for each $i \in I$, define R'_i by starting from x_i, then following the X-Y path given by Menger's theorem to some y_q, and finally following S_q.

We have constructed a collection $R'_0, \dots, R'_n$ of rays. It is straightforward to check that they are disjoint, and that for each $i < n$, R_i and R'_i start at the same vertex. □

We can now conclude that IRT follows from Σ^1_1-AC_0, which is known to be a theory of hyperarithmetic analysis:

Theorem 4.5. Σ^1_1-AC_0 *implies* IRT*. Hence for every* $Y \subseteq \mathbb{N}$*,* $\mathrm{HYP}(Y)$ *satisfies* IRT*.*

Proof. Note that Σ^1_1-AC_0 proves ACA_0. The proof of IRT presented above can be carried out in ACA_0, except for the following fact we used in the inductive step:

> Let $S_0, \dots, S_{f(n)}$ be a finite sequence of disjoint rays, where $f(n) = n^2 + \sum_{i<n}(n-i)$.

In order to carry out the proof, it is not sufficient to know that for each n, there is some collection of disjoint rays of size $f(n)+1$. Rather, we require that there is a single sequence such that the n^{th} entry of the sequence is a collection of disjoint rays of size $f(n)+1$. Such a sequence can be obtained using the axiom of choice. In this case, since the predicate "there exists $f(n)+1$ many disjoint rays" is Σ^1_1, such a sequence can be obtained using Σ^1_1-AC_0. Therefore we can prove IRT using Σ^1_1-AC_0. □

Theorem 4.1 follows from Theorems 4.2 and 4.5.

We will establish some implications and nonimplications between IRT and other theorems of hyperarithmetic analysis in §7. For now, we have

Proposition 4.6. IRT *implies* ACA_0*.*

Proof. We indicate how to formalize the first part of the proof of Theorem 4.2 in RCA_0. For each n, consider the tree $T_n \subseteq N^{<N}$ consisting of all strings of the form $s^\frown 0^t$ such that some number below n is enumerated into $\emptyset'$ at stage s, and either $t \leq s$ or $\emptyset'_s \upharpoonright n = \emptyset'_t \upharpoonright n$. First note that RCA_0 proves that for each n, there is some s such that $\emptyset' \upharpoonright n = \emptyset'_s \upharpoonright n$. This implies that for each n, T_m has a branch for all $m \leq n$. Therefore the forest $\bigsqcup_n T_n$ contains arbitrarily many disjoint rays.

By IRT, $\bigsqcup_n T_n$ contains infinitely many disjoint rays. In order to compute $\emptyset'$ from an infinite sequence of disjoint rays, we want to show that for each n, the infinite sequence contains some ray in some T_m for $m \geq n$. It suffices to show that each T_n has at most one branch (and hence contains at most one ray from the given sequence). Fix n. Any branch on T_n must be of the form $s^\frown 0^\infty$. If $s^\frown 0^\infty$ lies on T_n, then no number below n is enumerated into $\emptyset'$ after stage s. Hence if $s' > s$, then $s'^\frown 0^\infty$ cannot be a branch on T_n. This shows that T_n has at most one branch. ☐

5. Variants of IRT and Hyperarithmetic Analysis

In this section, we show that at least five of the eight principles $\mathsf{IRT}_{\mathrm{XYZ}}$ are theorems of hyperarithmetic analysis:

Theorem 5.1. *All single-ray variants of* IRT *(i.e.,* $\mathsf{IRT}_{\mathrm{XYS}}$*) and* $\mathsf{IRT}_{\mathrm{UVD}}$ *are theorems of hyperarithmetic analysis.*

$\mathsf{IRT}_{\mathrm{UVS}}$ and $\mathsf{IRT}_{\mathrm{UVD}}$ were proved by Halin [11, 12]. $\mathsf{IRT}_{\mathrm{UES}}$ is an exercise in [5, 8.2.5(ii)]. $\mathsf{IRT}_{\mathrm{DVS}}$ and $\mathsf{IRT}_{\mathrm{DES}}$ appear to be folklore.

Of the other three variants, $\mathsf{IRT}_{\mathrm{DED}}$ and $\mathsf{IRT}_{\mathrm{DVD}}$ are open problems of graph theory ([3] and Bowler, personal communication). We do, however, have interesting results about these principles when restricted to directed forests (Theorem 6.13, Corollary 6.14). The other one, $\mathsf{IRT}_{\mathrm{UED}}$, was proved by Bowler, Carmesin, Pott [3] using structural results about ends. We hope to analyze its strength in future work.

The proof of Theorem 5.1 consists of several variations of the proof of Theorem 4.1. One of which ($\mathsf{IRT}_{\mathrm{DES}}$) requires some additional ideas.

In order to minimize repetition, we establish some implications between some variants of IRT over RCA_0. The proofs of each of these reductions follow the same basic plan. To deduce $\mathsf{IRT}_{\mathrm{XYZ}}$ from $\mathsf{IRT}_{\mathrm{X'Y'Z'}}$ we provide computable maps g, h and k which take X-graphs G to X′-graphs G', Y-disjoint Z-rays or sets of Y-disjoint Z-rays in G to Y′-disjoint Z′-rays or sets of Y′-disjoint Z′-rays in G', and Y′-disjoint Z′-rays or sets of Y′-disjoint Z′-rays in G' to Y-disjoint Z-rays or sets of Y-disjoint Z-rays in G, respectively. These functions are designed to take witnesses of the hypothesis of $\mathsf{IRT}_{\mathrm{XYZ}}$ in G to witnesses of the hypothesis of $\mathsf{IRT}_{\mathrm{X'Y'Z'}}$ in G' and witnesses to the conclusion of $\mathsf{IRT}_{\mathrm{X'Y'Z'}}$ in G' to witnesses to the conclusion of $\mathsf{IRT}_{\mathrm{XYZ}}$ in G. Clearly it suffices to provide such computable maps to establish the desired reduction in RCA_0.

Lemma 5.2. *Given an undirected graph G, we can uniformly compute a directed graph G' and mappings between Z-rays in G and Z-rays in G' which preserve Y-disjointness.*

Proof. We define a computable map g from undirected graphs G to directed graphs G' as follows. The set of vertices of G' consists of the vertices of G, together with two new vertices $x = x(u, v)$ and $y = y(u, v)$ for each edge $\{u, v\}$ in G. The set of edges of G' consists of five edges $\langle u, x\rangle$, $\langle v, x\rangle$, $\langle x, y\rangle$, $\langle y, u\rangle$, $\langle y, v\rangle$ for each edge $\{u, v\}$ in G.

Next we define a computable map h_{S}: given a ray $u_0, u_1, \ldots$ in G, h_{S} maps it to the ray $u_0, x(u_0, u_1), y(u_0, u_1), u_1, \ldots$ in G'. Conversely, we define a computable map k_{S} from rays R' in G' into rays R in G as follows. Observe that exactly one of the first three vertices in R' is a vertex in G, because the only outgoing edges from a vertex $y(u, v)$ lead to u or to v, and the only outgoing edge from a vertex $x(u, v)$ leads to $y(u, v)$. We take this vertex (say u_0) to be the first vertex of R. Every outgoing edge from u_0 leads to some $x(u_0, v)$. Combining the above observations, we deduce that the tail $u_0 R'$ has the form $u_0, x(u_0, u_1), y(u_0, u_1), u_1, \ldots$. Then k_{S} maps R' to the ray $R = u_0, u_1, \ldots$ in G.

Similarly, given a double ray $\ldots, u_{-1}, u_0, u_1, \ldots$ in G, h_{D} maps it to the double ray $\ldots, u_{-1}, x(u_{-1}, u_0), y(u_{-1}, u_0), u_0, x(u_0, u_1), y(u_0, u_1), u_1, \ldots$ in G'. We can show that every double ray in G' has this form by considering the incoming edges to each vertex in G'. Therefore we can define a computable map from double rays in G' to double rays in G by $k_{\mathrm{D}} = h_{\mathrm{D}}^{-1}$.

It is straightforward to check that h_{S}, k_{S}, h_{D} and k_{D} preserve Y-disjointness. □

Therefore we have

Proposition 5.3. *The directed variants of* IRT *imply their corresponding undirected variants, i.e.,* $\mathsf{IRT}_{\mathrm{DYZ}}$ *implies* $\mathsf{IRT}_{\mathrm{UYZ}}$ *for each value of Y and Z.*

Lemma 5.4. *Given a directed graph G, we can uniformly compute a directed graph G' and mappings between Z-rays in G and Z-rays in G' which satisfy the following properties: if two Z-rays in G are vertex-disjoint, then the corresponding Z-rays in G' are edge-disjoint, and if two Z-rays in G' are edge-disjoint, then the corresponding Z-rays in G are vertex-disjoint.*

Proof. We define a computable map g from directed graphs $G = \langle V, E\rangle$ to directed graphs G' as follows. The set of vertices of G' is $\{x_i, x_o : x \in V\}$, where i and o stand for incoming and outgoing respectively. The set of edges of G' consists of $\langle u_o, v_i\rangle$ for each $\langle u, v\rangle \in E$, and $\langle x_i, x_o\rangle$ for each $x \in V$.

Next we define a computable map h_{S}: given a ray $x^0, x^1, \ldots$ in G, h_{S} maps it to the ray $x_i^0, x_o^0, x_i^1, x_o^1, \ldots$ in G'. Conversely, we define a computable map k_{S} from rays R' in G' to rays in G as follows. Given R', the ray R visits the vertex x in G whenever $\langle x_i, x_o\rangle$ appears in R'. (For example, we map $x_i^0, x_o^0, x_i^1, x_o^1, \ldots$ to $x^0, x^1, \ldots$ and we map $x_o^0, x_i^1, x_o^1, \ldots$ to $x^1, x^2, \ldots$.)

Similarly, h_{D} maps a given double ray $\ldots, x^{-1}, x^0, x^1, \ldots$ in G to the double ray

$$\ldots, x_i^{-1}, x_o^{-1}, x_i^0, x_o^0, x_i^1, x_o^1, \ldots$$

in G'. Every double ray in G' has this form, so we may define $k_{\mathrm{D}} = h_{\mathrm{D}}^{-1}$.

It is straightforward to check that the above mappings have the desired properties. □

Therefore we have

Proposition 5.5. *The directed edge-disjoint variants of* IRT *imply their corresponding directed vertex-disjoint variants, i.e.,* $\mathsf{IRT}_{\mathrm{DEZ}}$ *implies* $\mathsf{IRT}_{\mathrm{DVZ}}$ *for each value of Z.*

Lemma 5.6. *Given a directed graph G, we can uniformly compute a directed graph G' and mappings between sets of rays in G and sets of double rays in G' which preserve cardinality, vertex-disjointness, and edge-disjointness.*

Proof. We define a computable map g from directed graphs G to directed graphs G' as follows. The vertex set of G' contains the vertex set of G, and for each vertex x of G, additional vertices x_n for each $n < 0$. In G', we also have an edge $\langle x_{n-1}, x_n\rangle$ for all $n < 0$.

$$\begin{tikzcd}
 & \mathrm{DED} \arrow[dl] \arrow[d] \arrow[dr] & \\
\mathrm{DVD} \arrow[d] \arrow[dr] & \mathrm{UED} & \mathrm{DES} \arrow[dl] \arrow[d] \\
\mathrm{UVD} \arrow[dr, "I\Sigma^1_1"'] & \mathrm{DVS} \arrow[d] & \mathrm{UES} \\
 & \mathrm{UVS} &
\end{tikzcd}$$

FIGURE 1. Known implications between variants of IRT. All implications are over RCA_0, except for the implication from UVD to UVS (Theorem 6.15).

Next we define a computable map h from sets of Y-disjoint rays in G to sets of Y-disjoint double rays in G' as follows. Given a set of Y-disjoint rays in G, we first ensure that each ray begins at a different vertex, by replacing it with a tail if necessary. (This is only relevant if the rays are edge-disjoint rather than vertex-disjoint.) Then for each ray $x^0, x^1, \ldots,$ we consider the double ray $\ldots, x^0_{-2}, x^0_{-1}, x^0, x^1, \ldots$ in G'. This yields a set of double rays in G' of the same cardinality.

Finally we define a computable map k from sets of Y-disjoint double rays in G' to sets of Y-disjoint rays in G. Given a double ray $\ldots, x^{-1}, x^0, x^1, \ldots$ in G', we search for the least $n \geq 0$ such that x^n is a vertex in G. Then we map the given double ray to the ray $x^n, x^{n+1}, \ldots$ in G. This induces a cardinality-preserving map k from sets of double rays in G' to sets of rays in G.

It is straightforward to check that the above mappings preserve the desired disjointness properties. $\square$

Therefore we have

Proposition 5.7. *The directed double ray variants of* IRT *imply their corresponding directed single ray variants, i.e.,* $\mathsf{IRT}_{\mathrm{DYD}}$ *implies* $\mathsf{IRT}_{\mathrm{DYS}}$ *for each value of* Y.

Figure 1 summarizes the known implications between variants of IRT over RCA_0. We will show that $\mathsf{IRT}_{\mathrm{UVD}}$ implies $\mathsf{IRT}_{\mathrm{UVS}}$ over $\mathsf{RCA}_0 + I\Sigma^1_1$ (Theorem 6.15).

Remark 5.8. Bowler, Carmesin, Pott [3, pg. 2 l. 3–7] describe an implication from $\mathsf{IRT}_{\mathrm{UVS}}$ to $\mathsf{IRT}_{\mathrm{UES}}$ which appears to use Σ^1_1-AC_0. It turns out that the graph-theoretic principle used to carry out the implication does not imply even ACA_0 over RCA_0 (and is hence much weaker than Σ^1_1-AC_0), but when combined with ACA_0, yields a theory of hyperarithmetic analysis. It and several other principles with the same property (almost theorems/theories of hyperarithmetic analysis) are analyzed in Shore [26].

We return to the goal of proving Theorem 5.1.

Theorem 5.9. *For each choice of XYZ,* $\mathsf{IRT}_{\mathrm{XYZ}}$ *implies* ACA_0. *Furthermore, every standard model of* RCA_0 *and* $\mathsf{IRT}_{\mathrm{XYZ}}$ *is closed under hyperarithmetic reduction.*

Proof. By Proposition 5.3, it suffices to prove the desired result for the undirected variants of IRT. Theorem 4.2 and Proposition 4.6 together assert the desired result for $\mathsf{IRT}_{\mathrm{UVS}}$. We describe how to modify the proofs of Theorem 4.2 and Proposition 4.6 to prove the desired result for the other variants of IRT.

Observe that in the aforementioned proofs, we only applied IRT to forests such that each of the constituent trees has a unique branch. In such graphs, none of the constituent trees can contain two rays which are edge-disjoint. Hence the aforementioned proofs establish the desired result for $\mathsf{IRT}_{\mathrm{UES}}$ as well.

In order to prove the desired result for $\mathsf{IRT}_{\mathrm{UVD}}$ and $\mathsf{IRT}_{\mathrm{UED}}$, we modify the aforementioned proofs as follows. For each tree, we may modify it by adding a computable branch which is not already on the tree. (For concreteness: increment every entry of every string in the tree by 1, then add the branch 0^∞ to the tree). The resulting tree satisfies the following properties:

- it contains some double ray which is Turing equivalent to the branch on the original tree;
- no two double rays in the tree can be vertex-disjoint or edge-disjoint;
- given any double ray in the tree, we can uniformly compute the branch on the original tree.

It is straightforward to check that the modified proofs establish the desired result for $\mathsf{IRT}_{\mathrm{UVD}}$ and $\mathsf{IRT}_{\mathrm{UED}}$. $\square$

Henceforth we will not explicitly mention uses of ACA_0 whenever we are assuming any $\mathsf{IRT}_{\mathrm{XYZ}}$.

Next, we show that $\mathsf{IRT}_{\mathrm{UVD}}$ and $\mathsf{IRT}_{\mathrm{DES}}$ are provable in Σ^1_1-AC_0 (Theorems 5.10, 5.15). It then follows from Propositions 5.3 and 5.5 that $\mathsf{IRT}_{\mathrm{UES}}$, $\mathsf{IRT}_{\mathrm{DVS}}$ and $\mathsf{IRT}_{\mathrm{UVS}}$ are also provable in Σ^1_1-AC_0, completing the proof of Theorem 5.1. Note that as Σ^1_1-AC_0 implies ACA_0, we will use ACA_0 without explicit mention whenever we are assuming Σ^1_1-AC_0.

Theorem 5.10. Σ^1_1-AC_0 *implies* $\mathsf{IRT}_{\mathrm{UVD}}$.

Proof. This proof is very similar to that of Theorem 4.5, except we need to grow our family in “two directions”. We build the infinite family by induction. At stage s we will start with a family of s many finite disjoint paths $P^s_1, P^s_2, \dots, P^s_s$ and disjoint double rays $R^s_1, R^s_2, \dots, R^s_s$ such that P^s_i is a subpath of R^s_i for $i = 1, \dots, s$. At the end of the stage, we will have extended each of the paths P^s_i at both ends, begun a new path, and maintained that each of our paths can be extended to a family of disjoint double rays. Our infinite family will have members $\bigcup_s P^s_i$ for each $i \in N$.

If $s = 0$, then we start with nothing. Fix a double ray R^1_1 in G and pick any two adjacent vertices in R^1_1 to be P^1_1. If $s > 0$, then we have, by induction, disjoint paths $P^s_1, \dots, P^s_s$, all of length at least two, and disjoint double rays $R^s_1, \dots, R^s_s$ extending them, respectively. For ease of reading, we omit the superscript s and write $P^s_i = P_i$ and $R^s_i = R_i$.

Pick a set $\mathcal{R}$ of $|P_1 \cup \dots \cup P_s| + 4n^2 + 1$ disjoint double rays. From $\mathcal{R}$ delete any double rays that intersect any of the paths $P_1, \dots, P_s$. By the pairwise disjointness of the members of $\mathcal{R}$, we delete at most $|P_1 \cup \dots \cup P_s|$ many members from $\mathcal{R}$. Hence, $\mathcal{R}$ has at least $4n^2 + 1$ many double rays remaining.

Now, choose and delete an edge in each path P_i (possible as each path is of length at least two), which breaks P_i into two paths P^f_i and P^b_i (f for forward and b for backward).

Remove the corresponding edge from R_i to obtain single-rays R_i^f and R_i^b extending P_i^f and P_i^b, respectively.

Repeat the following step as often as possible: If there exists an $i \in \{1, \dots, s\}$ and a direction $d \in \{f, b\}$ such that $(R_i^d)'$ is undefined and the tail of R_i^d following P_i^d meets at most $2s$ many double rays currently in $\mathcal{R}$, then delete those double rays from $\mathcal{R}$, set $(R_i^d)' = R_i^d$, and extend P_i^d one step further down R_i^d to get $(P_i^d)'$.

Let I be the subset of $\{1, \dots, s\} \times \{f, b\}$ for which we have not defined $(R_i^d)'$, and let $m = |I|$. Then $\mathcal{R}$ contains at least

$$4s^2 + 1 - (2s - m)2s = 2sm + 1 \geq m^2 + 1$$

many double rays.

From the above, for every $\langle i, d \rangle \in I$ there are at least $2s$ many elements of $\mathcal{R}$ that intersects R_i^d. Let z_i^d be the first vertex on the mth ray $R \in \mathcal{R}$ that R_i^d meets (as R_i^d is a single ray this definition makes sense). Also, let Q_i^d be the portion of R_i^d between P_i^d and z_i^d (but including the endpoints), and let $Z = \bigcup_{(i,d) \in I} Q_i^d$. The set Z meets at most m^2 many members of $\mathcal{R}$, meaning that Z does not intersect at least one member of $\mathcal{R}$. Pick such a double ray R, and let $R_{s+1}^{s+1} = R$ and let P_{s+1}^{s+1} be any two adjacent vertices of R_{s+1}^{s+1}. Now, discard all such rays from $\mathcal{R}$, so that it only consists of those double rays that meet Z.

For each double ray $R \in \mathcal{R}$ remaining, choose a tail R' of R that contains $R \cap Z$. This is possible as Z is finite. Now we are finally in the same situation as in the proof of Theorem 4.5, and we can thus use Menger's theorem to define new rays $(R_i^d)'$ that are pairwise disjoint, extend the paths P_i^d, and do not intersect R_{s+1}^{s+1}. Finally, we turn everything back into double rays by replacing the removed edge between P_i^b and P_i^d and that same edge between $(R_i^b)'$ and $(R_i^f)'$, which completes the induction step.

Analysis. As in with the proof of Theorem 4.5, one can use Σ_1^1-AC_0 to produce a single set X such that the nth column of X is a set of n many pairwise disjoint double rays, as the predicate "there exists a set of n many pairwise disjoint double rays in G" is arithmetical. Other than this, all of the construction can be carried out in ACA_0, as detecting intersections between (single or double) rays is arithmetical in those rays. □

As for $\mathsf{IRT}_{\mathrm{DES}}$, instead of following the proof of Theorem 4.5, we will reduce $\mathsf{IRT}_{\mathrm{DES}}$ to the problem of finding an infinite sequence of *vertex-disjoint* rays in a certain locally finite graph (see [3, pg. 2 l. 3–7]). To carry out this reduction, we define the line graph:

Definition 5.11 (RCA_0)**.** The *line graph* $L(G)$ of an X-graph G is the X-graph whose vertices are the edges of G and whose edges are the $((u, v), (v, w))$, where (u, v) and (v, w) are edges in G.

Lemma 5.12. *Let G be an X-graph. There is a computable mapping from rays in G to rays in $L(G)$ such that if two rays in G are edge-disjoint, then their images are vertex-disjoint.*

Proof. Map $x_0, x_1, x_2, \dots$ to $(x_0, x_1), (x_1, x_2), \dots$. □

Vertex-disjoint rays in $L(G)$ do not always yield edge-disjoint rays in G, however. An extreme counterexample is what is called the *(undirected) star graph* which consists of a single vertex with infinitely many neighbors: It does not contain any rays yet its line graph is isomorphic to the complete graph on $\mathbb{N}$, which contains infinitely many vertex-disjoint rays. Nonetheless, if G is locally finite, then vertex-disjoint rays in $L(G)$ do correspond to edge-disjoint rays in G:

Lemma 5.13 (ACA_0). *Let G be a locally finite X-graph. There is a mapping from rays in $L(G)$ to rays in G such that if two rays in $L(G)$ are vertex-disjoint, then their images are edge-disjoint rays in G.*

Proof. Given a ray $R = e_0, e_1, \ldots$ in $L(G)$, we construct a ray $S = y_0, y_1, \ldots$ in G by recursion. Start by defining y_0 to be the least vertex in e_0. Having defined y_n, we define y_{n+1} as follows. Let k be the largest index such that y_n is an endpoint of e_k. Such k exists because G is locally finite and R is a ray. We can find such k by ACA_0. Then define y_{n+1} to be the endpoint of e_k other than y_n. This completes the recursion. By construction S is infinite and contains no repeated vertices, hence it is a ray. Observe that every edge in S is a vertex in R, so the above mapping maps vertex-disjoint rays in $L(G)$ to edge-disjoint rays in G. □

It remains to show that we can restrict our attention to locally finite graphs. We accomplish this with the help of Σ^1_1-AC_0. Given a directed graph G with arbitrarily many edge-disjoint rays, we can use Σ^1_1-AC_0 to choose a family $(R^k_j)_{j\leq k}$ where for each k, the rays $R^k_1, R^k_2, \ldots, R^k_k$ are edge-disjoint. From this family, we may construct an appropriate locally finite subgraph of G:

Lemma 5.14 (RCA_0). *Suppose that G is an X-graph and there is some family $(R^k_j)_{j\leq k}$ such that for each k, $R^k_1, R^k_2, \ldots, R^k_k$ are edge-disjoint rays in G. Then there is some locally finite X-subgraph G' of G and some family $(S^k_j)_{j\leq k}$ such that for each k, $S^k_1, S^k_2, \ldots, S^k_k$ are edge-disjoint rays in G'.*

Proof. Define the vertices of G' to be the vertices of G, say $\{v_i : i \in N\}$. We specify the set of edges E' of G' by providing a recursive construction of sets E_i of edges putting in a set of edges at each step. We guarantee that each E_i is a union of finitely many sets of edge-disjoint rays in G and that after stage k no edge with a vertex v_i for $i < k$ as an endpoint is ever put into E' after stage k.

Begin at stage 0 by putting all the edges in R^1_1 into E_1. Proceeding recursively at stage k we have E_k and consider the edge-disjoint rays $R^k_1, R^k_2, \ldots, R^k_k$. For each $i < k$, say $R^k_i = x^k_{i,0}, x^k_{i,1}, \ldots$. Each v_j for $j < k$ appears at most once in R^k_i as R^k_i is a ray. For each $j < k$, since we have access to the set of vertices of R^k_i, we can decide whether R^k_i contains v_j and find the largest index n such that $v_j = x^k_{i,n}$. Call it $n^k_{i,j}$. If there is no such n, set $n^k_{i,j} = 0$. Define S^k_i to be the tail of R^k_i after $x^k_{i,\max_{j<k} n^k_{i,j}}$. We put all the edges in S^k_i into E_{k+1}. Let $E' = \bigcup_k E_k$.

It remains to show that G' is locally finite. Consider any vertex v_k. No edge containing v_k as an endpoint is put in after stage k. On the other hand, E_k is the union of finitely many finite sets of edge-disjoint rays (all of which have been computed uniformly). Each set of edge-disjoint rays in this union has v_k appearing at most once in each of its rays. Thus at most two edges containing v_k appear in each of the finitely many rays in this set. Therefore there are only finitely many edges containing v_k in each of the finite sets of edge-disjoint rays making up E_k. All in all, only finitely many edges in G' contain v_k. □

We are ready to prove

Theorem 5.15. Σ^1_1-AC_0 *implies* $\mathsf{IRT}_{\mathrm{XES}}$ *for each value of X.*

Proof. Given an X-graph G with arbitrarily many edge-disjoint rays, we can use Σ^1_1-AC_0 to choose a family $((Q^k_j)_{j\leq k})_{k\in N}$ such that for each k, the rays $Q^k_1, Q^k_2, \ldots, Q^k_k$ are edge-disjoint.

By Lemma 5.14, there is a locally finite subgraph H of G and a family $(R_j^k)_{j\leq k}$ such that for each k, $R_1^k, R_2^k, \dots, R_k^k$ are edge-disjoint rays in H. By Lemma 5.12, there is a family $((S_j^k)_{j\leq k})_{k\in N}$ such that for each k, $S_1^k, S_2^k, \dots, S_k^k$ are vertex-disjoint rays in $L(H)$. By the second part of the proof of Theorem 4.5 (which can be carried out in ACA_0), $L(H)$ has infinitely many vertex-disjoint rays. Finally by Lemma 5.13, H has infinitely many edge-disjoint rays. Hence G has infinitely many edge-disjoint rays. □

By the discussion before Theorem 5.10, this completes the proof of Theorem 5.1.

Finally, we give a proof of $\mathsf{IRT}_{\mathrm{DED}}$ for directed forests using Σ^1_1-AC_0 (recall that $\mathsf{IRT}_{\mathrm{DED}}$ remains open). We will see that Σ^1_1-AC_0 and $\mathsf{IRT}_{\mathrm{DED}}$ for directed forests are equivalent over $\mathsf{RCA}_0 + I\Sigma^1_1$ (Theorem 6.13).

Theorem 5.16. *Σ^1_1-AC_0 implies $\mathsf{IRT}_{\mathrm{DED}}$ for directed forests.*

We first prove two lemmas.

Lemma 5.17. *Let G be a directed forest and let $R_0 = \langle x_{0,i} | i \in \mathbb{Z}\rangle, R_1 = \langle x_{1,i} | i \in \mathbb{Z}\rangle$ be directed double rays in G. If R_0 and R_1 have an edge in common, then there are $[i,j]$ and $[k,l]$ with $-\infty \leq i < j \leq +\infty$ and $-\infty \leq k < l \leq +\infty$ such that $R_0 \restriction [i,j] = R_1 \restriction [k,l]$ and otherwise R_0 and R_1 have no vertex in common. We call $R_0 \restriction [i,j] = R_1 \restriction [k,l]$ the intersection of R_0 and R_1.*

Proof. Suppose R_0, R_1 provide a counterexample. As their intersection contains an edge they lie in the same directed tree T in G and can be viewed as (undirected) double rays in $\hat{T}$ (the underlying graph for T). As they form a counterexample to the lemma, there are vertices $x_{0,i_0} = x_{1,i_1}; x_{0,j_0} = x_{1,j_1}$ in both R_0 and R_1 with $[x_{0,i_0}, x_{0,j_0}] \neq [x_{1,i_1}, x_{1,j_1}]$. Thus there are two different paths in $\hat{T}$ from $x_{0,i_0} = x_{1,i_1}$ to $x_{0,j_0} = x_{1,j_1}$ contradicting $\hat{T}$'s being a tree. □

Lemma 5.18. *There is a computable function f such that given any sequence $\langle S_i | i < n\rangle$ of DED rays in a directed tree T and sequence $\langle R_j | j < f(n)\rangle$ of DED rays in T, we can construct a sequence $\langle S'_i | i \leq n\rangle$ of DED rays in T such that $S_i \restriction [-n,n] = S'_i \restriction [-n,n]$ for $i < n$. Indeed, we may take $f(n) = 2n^2 + 2^{2n}n! + 1$.*

Proof. First we remove all the R_j that contain an edge in any $S_i \restriction [-n,n]$ at cost of at most $2n^2$ many j. Consider any remaining R_j in the second sequence. By Lemma 5.17, its intersections with the S_i are intervals $Q_{j,i}$ of edges in R_j which are disjoint as the S_i are. By our first thinning of the R_j list, none of the $Q_{j,i}$ intersect any of the $S_i \restriction [-n,n]$ so each $Q_{j,i}$ must lie entirely above or entirely below $S_i \restriction [-n,n]$. We associate to each R_j a label consisting of the set $C_j = \{i < n | Q_{j,i} \neq \emptyset\}$; the elements i of C_j in the order in which the $Q_{j,i}$ (for $i \in C_j$) appear in R_j (in terms of the ordering of $\mathbb{Z}$) along with a $+$ or $-$ depending on which side of $[-n,n]$ it falls in S_i. We write $Q^s_{j,i}$ for the starting vertex of $Q_{j,i}$ and $Q^e_{j,i}$ for the ending one. Now there are, of course, at most finitely many such labels. In particular, there are at most $2^n n! 2^n$ such labels. Thus if we have $2^{2n}n! + 1$ many R_j left at least two of them, say R_a and R_b have the same label say with set C.

Claim: $|C| < 2$.

For the sake of a contradiction, assume we have $k \neq l$ in C with k preceding l in the ordering of C in the label. Say R_a is the ray such that $Q_{a,k}$ is before $Q_{b,k}$ in S_k. We consider two cases: (1) $Q_{a,l}$ is before $Q_{b,l}$ in S_l and (2) $Q_{b,l}$ is before $Q_{a,l}$ in S_l. We now produce, for each case, two vertices with two distinct sequences (i) and (ii) of adjacent edges in T connecting them. These sequences are illustrated in Figures 2 and 3.

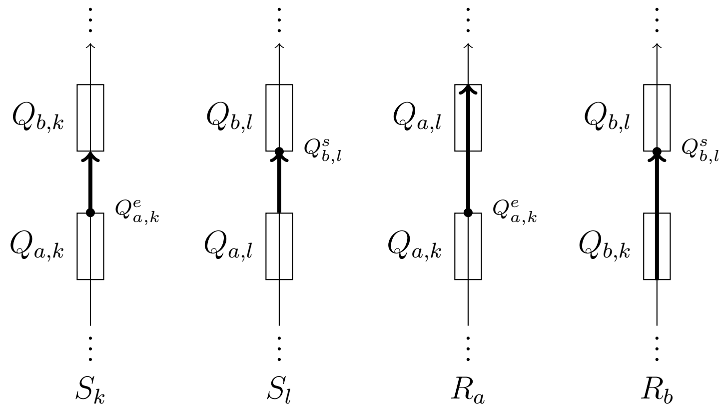

FIGURE 2. (1i) follows the thick arrows in R_a and S_l. (1ii) follows the thick arrows in S_k and R_b.

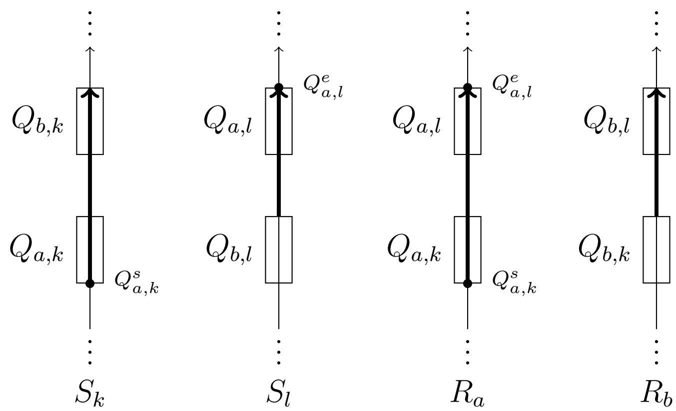

FIGURE 3. (2i) follows the thick arrows in S_k, R_b and S_l. (2ii) follows the thick arrow in R_a.

(1i) Start at $Q^e_{a,k}$ in R_a and go to $Q^e_{a,l}$ then in S_l go to $Q^s_{b,l}$.

(1ii) Start at $Q^e_{a,k}$ in S_k and go to $Q^s_{b,k}$ then go in R_b to $Q^s_{b,l}$.

(2i) Start at $Q^s_{a,k}$ in S_k and go to $Q^e_{b,k}$ then in R_b go to $Q^e_{b,l}$ then in S_l go to $Q^e_{a,l}$.

(2ii) Start at $Q^s_{a,k}$ in R_a and go to $Q^e_{a,l}$.

To see that the two sequences of vertices are different, note for (1) that (1i) contains an edge in $Q_{a,l}$ but (1ii) does not. For (2) note that (2i) contains an edge in $Q_{b,k}$ but (2ii) does not. We now, in each case, view the associated two distinct sequences of vertices with the same endpoints in the underlying tree $\hat{T}$. The only way one can have such sequences in a tree is for one of the sequences to contain vertices uvu in order. However, any three successive vertices in any of these sequences lie within one of the four rays being considered but no ray can have a single vertex appear in two locations:

Case 1: (i) If u is any vertex in the sequence before $Q^s_{a,l}$ then vu is also in R_a. If u is in $Q_{a,l}$ (and so in S_l) or later in the sequence then vu is also in S_l. (ii) If u is any vertex in the sequence before $Q^s_{b,k}$ then vu is also in S_k. If u is in $Q_{b,k}$ or later in the sequence then it and

vu are all in R_b.

Case 2: (i) If u is any vertex in the sequence before $Q^s_{b,k}$ then vu is also in S_k. If u is in $Q_{b,k}$ (and so in R_b) or later in the sequence but before $Q^s_{b,l}$ then vu is also in R_b. If u is at or after $Q^s_{b,l}$ then it and vu are in S_l. (ii) All of this route is in R_a and so has no repetitions of vertices.

Knowing now that $|C|$ is 0 or 1, we complete the proof of the Lemma. If $|C| = 0$, then both R_a and R_b are disjoint from all the S_i and so we may add on either one of them as S'_n while keeping $S_i = S'_i$ for all $i < n$. Otherwise say $C = \{i\}$. Let c be the one of a or b such that $Q_{c,i}$ is closer to $S_i \restriction [-n, n]$. (Remember that they are both on the same side of this interval in S_i by our fixing the label.) Now replace the tail of S_i starting with $Q_{c,i}$ and going away from $S_i \restriction [-n, n]$ by the tail of R_c starting with $Q_{c,i}$ and going in the same direction. Let this ray be S'_i. Note that it is disjoint from all the S_j, $j \neq i$ as it contains only edges that are in S_i or R_c neither of which share any edges with such S_j. It is also disjoint from R_d where d is the one of a, b which is not c since all of its edges are either in R_c or in S_i outside of $Q_{d,i}$ by our choice of $Q_{c,i}$ as closer to $S_i \restriction [-n, n]$. As R_d is also disjoint from all the S_j for $j \neq i$ by our fixing the label, we may define $S'_j = S_j$ for $j < n$, $j \neq i$ and $S'_n = R_d$ to get the sequence required in the Lemma. □

Lemma 5.18 provides the inductive step for the following proof:

Proof of Theorem 5.16. Assume we are given a directed forest G with arbitrarily many DED rays. By Σ^1_1-AC_0 we may take a sequence $\langle R_{k,i} | k \in N\rangle$ such that, for each k, $\langle R_{k,i} | i < k\rangle$ is a sequence of k many DED rays in G. If there are infinitely many of the trees making up $\hat{G}$ each of which contains some $R_{k,i}$ then we are done. So we may assume that all of them are in one directed tree T. We now wish to define $\langle S_{k,s} | k \leq s \in N\rangle$ by recursion such that, for each s, $\langle S_{k,s} | k < s\rangle$ is a sequence of s many DED rays and moreover, for each $s \in N$, $S_{k,s+1} \restriction [-s, s] = S_{k,s} \restriction [-s, s]$ so the $S_{k,s}$ reach limits S_k which form an infinite sequence of DED rays in T as required. Lemma 5.18 provides precisely the required inductive step for the construction since we have the required sequences of DED rays $\left\langle R_{f(n),i} | i < f(n)\right\rangle$ at each step n of the construction. □

6. Variations on Maximality

In this section, we consider variants of IRT whose solutions are required to be maximal in terms of cardinality or maximal in terms of set inclusion.

6.1. Maximum Cardinality Variants of IRT.

Definition 6.1. Let $\mathsf{IRT}^*_{\mathrm{XYZ}}$ be the statement that every X-graph G has a set of Y-disjoint Z-rays of maximum cardinality.

$\mathsf{IRT}^*_{\mathrm{UVS}}$ was proved by Halin [11], who also proved the corresponding statement for uncountable graphs.

Remark 6.2. The notation in Definition 6.1 is inspired by the well known version ACA^*_0 of ACA_0 which asserts the existence of $A^{(n)}$, the nth jump of A, for every n:

$$\mathsf{ACA}^*_0: \quad (\forall A)(\forall n)(\exists W)(W^{[0]} = A \wedge (\forall i < n)(W^{[i+1]} = W^{[i]'})).$$

This asserts (in addition to ACA_0) particular instances of $I\Sigma^1_1$. So too (in addition to $\mathsf{IRT}_{\mathrm{XYZ}}$) do the $\mathsf{IRT}^*_{\mathrm{XYZ}}$ as we are about to see.

Proposition 6.3. *For each choice of XYZ,* $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *implies* $\mathsf{IRT}_{\mathrm{XYZ}}$ *over* RCA_0 *and* $\mathsf{IRT}_{\mathrm{XYZ}}$ *implies* $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *over* $\mathsf{RCA}_0 + I\Sigma^1_1$*. Therefore* $\mathsf{IRT}_{\mathrm{XYZ}}$ *and* $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *are equivalent over* $\mathsf{RCA}_0 + I\Sigma^1_1$*. In particular, they have the same standard models.*

Proof. The first implication holds because if an X-graph has arbitrarily many Y-disjoint Z-rays, then any sequence of Y-disjoint Z-rays in the graph of maximum cardinality must be infinite. To prove the second implication, let I be the set of m such that there is a sequence of Y-disjoint Z-rays which is indexed by m in the given graph. Note that I is Σ^1_1 in the given graph. If I is closed under successor, then by $I\Sigma^1_1$, the given graph satisfies the premise of $\mathsf{IRT}_{\mathrm{XYZ}}$. Then we may apply $\mathsf{IRT}_{\mathrm{XYZ}}$ to show that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ is satisfied. If I is not closed under successor, then $\mathsf{IRT}^*_{\mathrm{XYZ}}$ is trivially satisfied. □

It follows from Proposition 6.3 and Theorem 5.1 that

Corollary 6.4. $\mathsf{IRT}^*_{\mathrm{XYS}}$ *and* $\mathsf{IRT}^*_{\mathrm{UVD}}$ *are theorems of hyperarithmetic analysis.*

It follows from Proposition 6.3 and Theorem 5.9 that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ implies ACA_0, so we will not explicitly mention uses of ACA_0 whenever we are assuming any $\mathsf{IRT}^*_{\mathrm{XYZ}}$.

Using Lemmas 5.2, 5.4 and 5.6, we can prove

Proposition 6.5. $\mathsf{IRT}^*_{\mathrm{DYZ}}$ *implies* $\mathsf{IRT}^*_{\mathrm{UYZ}}$*,* $\mathsf{IRT}^*_{\mathrm{DEZ}}$ *implies* $\mathsf{IRT}^*_{\mathrm{DVZ}}$ *and* $\mathsf{IRT}^*_{\mathrm{DYD}}$ *implies* $\mathsf{IRT}^*_{\mathrm{DYS}}$*.*

Next, we show that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ proves sufficient induction in order to transcend Σ^1_1-AC_0. This implies that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ is strictly stronger than $\mathsf{IRT}_{\mathrm{XYZ}}$ for certain choices of XYZ (Corollary 6.10). The connection between Σ^1_1-AC_0 and graphs is obtained by viewing the set of solutions of an arithmetic predicate as the set of (projections of) branches on a subtree of $N^{<N}$. In detail:

Lemma 6.6 (Simpson [29, V.5.4])**.** *If $A(X)$ is an arithmetical formula,* ACA_0 *proves that there is a tree $T \subseteq N^{<N}$ such that*

$$\forall X(A(X) \leftrightarrow \exists f(\langle X, f\rangle \in [T]) \\ \textit{and}\, \forall X(\exists \textit{ at most one } f)(\langle X, f\rangle \in [T]).$$

The following easy corollary will be useful.

Lemma 6.7. *If $A(n, X)$ is an arithmetical formula,* ACA_0 *proves that there is a sequence of subtrees $(T_n)_n$ of $N^{<N}$ such that for each $n \in N$,*

$$\forall X(A(n, X) \leftrightarrow \exists f(\langle X, f\rangle \in [T_n]) \\ \textit{and}\, \forall X(\exists \textit{ at most one } f)(\langle X, f\rangle \in [T_n]).$$

Proof. Say that $B(Y)$ holds if and only if $A(Y(0), X)$ holds, where X is such that $Y = Y(0)^\frown X$. Apply Lemma 6.6 to the arithmetical formula $B(Y)$ to obtain a tree $T \subseteq N^{<N}$. For each $n \in N$, define T_n to be the set of all σ such that $n^\frown\sigma \in T$. It is straightforward to check that $(T_n)_n$ satisfies the desired properties. □

Theorem 6.8. $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *proves* ACA^*_0*.*

Before proving the above theorem, we derive some corollaries:

Corollary 6.9. $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *proves the consistency of* Σ^1_1-AC_0*. Therefore it is not provable in* Σ^1_1-AC_0*.*

Proof. Simpson [29, IX.4.6] proves that $\mathsf{ACA}_0 + I\Sigma^1_1$ implies the consistency of Σ^1_1-AC_0. The only use of $I\Sigma^1_1$ in Simpson's proof is to establish ACA^*_0, so Simpson's proof shows that ACA^*_0 implies the consistency of Σ^1_1-AC_0. The desired result then follows from Theorem 6.8 and Gödel's second incompleteness theorem. □

Corollary 6.10. $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *is strictly stronger than* $\mathsf{IRT}_{\mathrm{XYZ}}$ *for the following choices of XYZ: XYS and UVD.*

Proof. We showed in §5 that the specified variants of IRT are provable in Σ^1_1-AC_0. On the other hand, none of the IRT^* are not provable in Σ^1_1-AC_0 (Corollary 6.9). □

We now prove Theorem 6.8:

Proof that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *implies* ACA^*_0. By Proposition 6.5, it suffices to prove the desired result for $\mathsf{IRT}^*_{\mathrm{UYZ}}$. To prove ACA^*_0 from $\mathsf{IRT}^*_{\mathrm{UYS}}$, begin by using Lemma 6.7 to define a sequence of trees $(T_n)_n$ such that for each n,

$$\begin{aligned}\forall W((W^{[0]} = A \wedge (\forall i < n)((W^{[i]})' = W^{[i+1]}) \wedge (\forall i > n)(W^{[i]} = \emptyset))\\ \leftrightarrow \exists f(\langle W, f\rangle \in [T_n]))\end{aligned}$$

and $\forall W(\exists \text{ at most one } f)(\langle W, f\rangle \in [T_n])$.

We want to show that each T_n is ill-founded. Note that if $m < n$ and T_n is ill-founded, then so is T_m. Therefore it suffices to show that for cofinally many n, T_n is ill-founded.

Apply $\mathsf{IRT}^*_{\mathrm{UYS}}$ to the disjoint union $\bigsqcup_n T_n$ to obtain a collection C of Y-disjoint rays of maximum cardinality. We prove that C is infinite. Suppose not. Then there is some maximum m such that C contains a ray in T_m. A ray in T_m can be computably truncated or extended to a branch on T_m, so T_m is ill-founded. Hence T_{m+1} is ill-founded as well (by ACA_0). But then there is a collection of Y-disjoint rays in $\bigsqcup_n T_n$ which has cardinality greater than that of C, contradiction.

We have proved that C is infinite. Next we prove that each T_n has at most one branch. That would imply that each T_n contains at most one ray in C, so C contains rays in cofinally many T_n, as desired.

If T_n has two distinct branches $\langle W_0, f_0\rangle$ and $\langle W_1, f_1\rangle$, then $W_0 \neq W_1$ by construction of T_n. Consider the least i such that $W_0^{[i]} \neq W_1^{[i]}$. Such i exists by ACA_0. Note that $0 < i \leq n$ because $W_0^{[0]} = A = W_1^{[0]}$ and $W_0^{[i]} = \emptyset = W_1^{[i]}$ for $i > n$. But then $W_0^{[i-1]} = W_1^{[i-1]}$ and $(W_0^{[i-1]})' \neq (W_1^{[i-1]})'$, contradiction.

This proves that $\mathsf{IRT}^*_{\mathrm{UYS}}$ implies ACA^*_0. In order to prove that $\mathsf{IRT}^*_{\mathrm{UYD}}$ implies ACA^*_0, we modify the above proof by adding a computable branch to each T_n to form a tree S_n. Apply $\mathsf{IRT}^*_{\mathrm{UYD}}$ to $\bigsqcup_n S_n$ to obtain a collection C of Y-disjoint double rays of maximum cardinality. Following the above proof, we may prove that C is infinite and each S_n contains at most one double ray in C. So C contains double rays in cofinally many S_n, as desired. □

Remark 6.11. The same proof shows that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ implies the following induction scheme: Suppose $(T_n)_n$ is a sequence of trees such that

(1) T_0 has a unique branch;
(2) for all n, the number of branches on T_{n+1} is the same as the number of branches on T_n.

Then for all n, there is a sequence $(P_m)_{m<n}$ such that for each $m < n$, P_m is the unique branch on T_m. We will use similar ideas to prove that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ implies unique-Σ^1_1-AC_0 (Definition 7.1) in §7. $\mathsf{IRT}^*_{\mathrm{XYZ}}$ also implies a similar induction scheme analogous to finite-Σ^1_1-AC_0 (Definition 7.2).

We can prove a stronger result for $\mathsf{IRT}^*_{\mathrm{DVD}}$:

Theorem 6.12. $\mathsf{IRT}^*_{\mathrm{DVD}}$ *(even for directed forests) implies* $I\Sigma^1_1$ *over* RCA_0*.*

Proof. Suppose $\Psi(n)$ is a Σ^1_1 formula such that $\Psi(0)$ and $\forall n(\Psi(n) \to \Psi(n+1))$ hold. Let $(T_n)_n$ be a sequence of subtrees of $N^{<N}$ such that T_n is ill-founded if and only if $(\forall m \leq n)\Psi(m)$ holds. For each n, orient each edge in T_n towards its root and add a computable ray which starts at its root. This forms a directed tree G_n which contains a double ray if and only if $(\forall m \leq n)\Psi(m)$ holds. Furthermore, no two disjoint double rays can lie in the same G_n.

Let G be the directed forest $\bigsqcup_n G_n$. By $\mathsf{IRT}^*_{\mathrm{DVD}}$, there is a sequence $(R_i)_i$ of disjoint double rays in G of maximum cardinality. Since $\Psi(0)$ holds, $(R_i)_i$ is nonempty. If $(R_i)_i$ is finite, let n be maximum such that G_n contains some R_i. Then $\Psi(n)$ holds, so $\Psi(n+1)$ holds as well. It follows that G_{n+1} contains some double ray, which we can then add to $(R_i)_i$ to obtain a larger sequence of disjoint double rays in G. Contradiction. Therefore $(R_i)_i$ is infinite. Since each G_n contains at most one R_i, it follows that infinitely many G_n contain some R_i. Therefore $\Psi(n)$ holds for all n. □

In fact, we have the following equivalences:

Theorem 6.13. *The following are equivalent (over* RCA_0*):*

(1) Σ^1_1-$\mathsf{AC}_0 + I\Sigma^1_1$;
(2) $\mathsf{IRT}_{\mathrm{DED}}$ *for directed forests* $+ I\Sigma^1_1$;
(3) $\mathsf{IRT}^*_{\mathrm{DED}}$ *for directed forests;*
(4) $\mathsf{IRT}^*_{\mathrm{DVD}}$ *for directed forests;*
(5) $\mathsf{IRT}_{\mathrm{DVD}}$ *for directed forests* $+ I\Sigma^1_1$.

Proof. (1) → (2) follows from Theorem 5.16. (2) → (3) follows from the proof of Proposition 6.3. (3) → (4) follows from the observation that the mapping of graphs defined in Lemma 5.4 sends a directed forest to a directed forest. (4) → (5) follows from Theorem 6.12 and the proof of Proposition 6.3.

To prove (5) → (1), suppose $A(n, X)$ is an arithmetical formula such that $\forall n \exists X A(n, X)$. By Lemma 6.7, there is a sequence $(T_n)_n$ of subtrees of $N^{<N}$ such that

$$\forall n \forall X (A(n, X) \leftrightarrow \exists f(\langle X, f\rangle \in [T_n])$$
$$\text{and } \forall X(\exists \text{ at most one } f)(\langle X, f\rangle \in [T_n]).$$

By assumption on $A(n, X)$, each T_n is ill-founded. We use $(T_n)_n$ to construct a sequence $(G_n)_n$ of directed trees as follows. First, for each n, we may construct a tree S_n whose branches are joins of branches on T_m for $m \leq n$, i.e., $P_0 \oplus \cdots \oplus P_m$ is a branch on S_n if and only if for each $m \leq n$, P_m is a branch on T_m. Second, orient each edge in S_n towards its root. Third, add to S_n a computable ray which starts at its root. This forms a directed tree G_n.

By $I\Sigma^1_1$, the directed forest $\bigsqcup_n G_n$ contains arbitrarily many disjoint double rays. Therefore $\bigsqcup_n G_n$ contains infinitely many disjoint double rays $(R_k)_k$, by $\mathsf{IRT}_{\mathrm{DVD}}$. Note that any double ray in any G_n must contain the computable ray we added, so any two double rays in the same G_n must intersect. This implies that each R_k belongs to some distinct G_n. Therefore for every m, there is some k and some $n > m$ such that R_k is a double ray in G_n. This allows us to construct a sequence $(X_n)_n$ where each X_n is a branch on T_n. $\square$

Since Σ^1_1-AC_0 (ATR_0, even) does not prove $I\Sigma^1_1$ (Simpson [29, IX.4.7]), it follows that

Corollary 6.14. $\mathsf{IRT}^*_{\mathrm{DYD}}$ *(even for directed forests) is not provable in* ATR_0*, and strictly implies* Σ^1_1-AC_0 *over* RCA_0*.*

Next, we show that $\mathsf{IRT}^*_{\mathrm{UVD}}$ implies $\mathsf{IRT}_{\mathrm{UVS}}$ over RCA_0 (see Figure 1).

Theorem 6.15. $\mathsf{IRT}^*_{\mathrm{UVD}}$ *implies* $\mathsf{IRT}_{\mathrm{UVS}}$ *over* RCA_0*. Therefore (1)* $\mathsf{IRT}_{\mathrm{UVD}}$ *implies* $\mathsf{IRT}_{\mathrm{UVS}}$ *over* $\mathsf{RCA}_0 + I\Sigma^1_1$*; (2) if any standard model of* RCA_0 *satisfies* $\mathsf{IRT}_{\mathrm{UVD}}$*, then it satisfies* $\mathsf{IRT}_{\mathrm{UVS}}$ *as well.*

Proof. Let G be a graph which contains arbitrarily many disjoint single rays. By $\mathsf{IRT}^*_{\mathrm{UVD}}$, there is a sequence of disjoint double rays in G of maximum cardinality. If this sequence is infinite, then there are infinitely many disjoint single rays in G as desired. Otherwise, suppose that $(R_i)_{i<j}$ is a sequence of disjoint double rays in G of maximum cardinality j. Let $\mathcal{R}$ be the subgraph of G consisting of the union of all R_i. Let H be the induced subgraph of G consisting of all vertices which do not lie in $\mathcal{R}$. Note that H does not contain any double ray, otherwise G would contain $j+1$ many disjoint double rays. Next, we expand H to the graph H', defined below.

Decompose H into its connected components $(H_i)_i$ (there may only be finitely many). Any two single rays in the same H_i must intersect, because if S_0 and S_1 are disjoint single rays in the same H_i, then we can construct a double ray in H_i by connecting them (start with a path between S_0 and S_1 of minimum length, then connect it to the tails of S_0 and S_1 which begin at the endpoints of the path).

For each i, define H'_i by adding a computable ray to H_i, which begins at the $<_N$-least vertex in H_i. Define H' to be the disjoint union $\bigsqcup_i H'_i$.

By $\mathsf{IRT}^*_{\mathrm{UVD}}$, there is a sequence of disjoint double rays in H' of maximum cardinality.

<u>Case 1.</u> If this sequence is infinite, then G contains infinitely many disjoint single rays because each double ray in the sequence has a tail which lies in H. In this case we are done.

<u>Case 2.</u> Otherwise, H' does not contain arbitrarily many disjoint double rays. Since any two single rays in the same H_i must intersect, we can transform any collection of disjoint single rays in H into a collection of disjoint double rays in H' of equal cardinality, by connecting each single ray to the $<_N$-least vertex in its connected component H_i and then following the computable ray we added. It follows that H does not contain arbitrarily many disjoint single rays. Fix l such that H does not contain $l+1$ many disjoint single rays.

Towards a contradiction, we construct a collection of $(j+1)$-many disjoint double rays in G as follows. Fix a collection $\mathcal{S}$ of $l + 2j + 4j(j+1)$ many disjoint single rays in G. First, at most $2j$ of these single rays lie in $\mathcal{R}$. In fact at most $2j$ of these single rays can have finite intersection with H, because given a collection of disjoint single rays each of which have finite intersection with H, we can obtain a collection of disjoint single rays in $\mathcal{R}$ of the same cardinality by replacing each ray with an appropriate tail. Second, by reasoning analogous to the above, at most l of these single rays can have finite intersection with $\mathcal{R}$. Therefore, there are at least $4j(j+1)$ many disjoint single rays in $\mathcal{S}$ each of which have infinite intersection with both $\mathcal{R}$ and H.

Next, split each of the j double rays R_i into two single rays R^b_i and R^f_i. By the pigeonhole principle, there is some single ray R of the form R^b_i or R^f_i, and at least $2(j+1)$ many disjoint single rays in $\mathcal{S}$, each of which have infinite intersection with both $\mathcal{R}$ and H. Call these rays $S_0, S_1, \ldots, S_{2(j+1)-1}$. Discard all the other rays in $\mathcal{S}$. Below we describe how to connect pairs of single rays S_k using segments of R, in order to form a collection of $(j+1)$-many disjoint double rays in G.

Let $x_0, x_1, \ldots$ denote the vertices of R. Since each single ray S_k has infinite intersection with R, by the pigeonhole principle, there is a pair of disjoint rays S_{k_0} and S_{l_0} such that for each tail R' of R, there is a vertex in $S_{k_0} \cap R'$ and a vertex in $S_{l_0} \cap R'$ such that no S_k intersects R between these two vertices. (Formally, we justify this by defining the following coloring recursively. Start from the first vertex on R which intersects some S_k. Search for the next vertex on R which intersects some S_l, $l \neq k$. Then we color 0 using the unordered pair $\{k, l\}$. Next, we search for the next vertex on R which intersects some S_m, $m \neq l$ and color 1 using $\{l, m\}$, and so on. Some color $\{k_0, l_0\}$ must appear infinitely often.) Then we commit to connecting S_{k_0} and S_{l_0} (but we do not do so just yet). Applying the pigeonhole principle again, there is a pair of disjoint rays S_{k_1} and S_{l_1} $(k_1, l_1 \neq k_0, l_0)$ such that for each tail R' of R, there is a vertex in $S_{k_1} \cap R'$ and a vertex in $S_{l_1} \cap R'$ such that no S_k, except perhaps S_{k_0} or S_{l_0}, intersects R between these two vertices. Again we commit to connecting S_{k_1} and S_{l_1}. Repeat this process until we have obtained $j+1$ pairs of single rays.

Finally, we connect these pairs of single rays in the opposite order in which we defined them: Start by picking some $x^j \in S_{k_j} \cap R$ and some $y^j \in S_{l_j} \cap R$. Then we define a double ray D_j by following S_{k_j} until x^j, then following R until y^j, and finally following S_{l_j}, i.e., $D_j := S_{k_j} x^j R y^j S_{l_j}$. Having defined $D_j, D_{j-1}, \ldots, D_{i+1}$, define $D_i := S_{k_i} x^i R y^i S_{l_i}$, where $x^i \in S_{k_i} \cap R$ and $y^i \in S_{l_i} \cap R$ are chosen as follows. Consider a tail R' of R such that the union of $x^j R y^j, \ldots, x^{i+1} R y^{i+1}$ is disjoint from (1) R'; (2) $S_{k_i} x$ for each $x \in S_{k_i} \cap R'$; (3) $y S_{l_i}$ for each $y \in S_{l_i} \cap R'$. By choice of k_i and l_i, there are vertices $x^i \in S_{k_i} \cap R'$ and $y^i \in S_{l_i} \cap R'$ such that none of $S_{k_j}, \ldots, S_{k_{i+1}}$ or $S_{l_j}, \ldots, S_{l_{i+1}}$ intersect $x^i R y^i$. It is straightforward to check that D_i is disjoint from $D_j, D_{j-1}, \ldots, D_{i+1}$. This process yields disjoint double rays $D_j, D_{j-1}, \ldots, D_0$ in G, contradicting maximality of j. □

Using some of the ideas in the previous proof, we can prove

Theorem 6.16. $\mathsf{IRT}^*_{\mathrm{UYS}}$ *for forests follows from* $\mathsf{IRT}^*_{\mathrm{UYD}}$ *for forests over* RCA_0*. Therefore* $\mathsf{IRT}_{\mathrm{UYS}}$ *for forests follows from* $\mathsf{IRT}_{\mathrm{UYD}}$ *for forests over* $\mathsf{RCA}_0 + I\Sigma^1_1$.

This result will be used in the proofs of Theorems 7.3, 7.7 and 7.10.

Proof. Let G be a forest. If G happens to have arbitrarily many disjoint double rays, then by $\mathsf{IRT}^*_{\mathrm{UYD}}$, G has infinitely many disjoint double rays. Therefore there is an infinite sequence of disjoint single rays in G. Such a sequence has maximum cardinality, so we are done in this case.

Suppose G does not have arbitrarily many disjoint double rays. By $\mathsf{IRT}^*_{\mathrm{UYD}}$ for forests, there is a sequence $(R_i)_{i<j}$ of disjoint double rays in G of maximum cardinality. Following the proof of Theorem 6.15, define the forests $\mathcal{R}$, H, and H'. There, we proved that no two single rays in the same connected component H_i of H can be disjoint.

By $\mathsf{IRT}^*_{\mathrm{UYD}}$ for forests, there is a sequence of disjoint double rays in H' of maximum cardinality. If this sequence is infinite, then there is an infinite sequence of disjoint single rays in H because each double ray in the sequence has a tail which lies in H. This is a sequence of disjoint single rays of maximum cardinality in G, so we are done in this case.

Otherwise, suppose $(S_k)_{k<l}$ is a disjoint sequence of double rays in H' of maximum cardinality. Consider the following disjoint sequence of single rays in G. First, for each $k < l$, consider the single ray formed by intersecting H and the double ray S_k. Second, for each $i < j$, we can split the double ray R_i into a pair of disjoint single rays in G. This yields a finite sequence $(Q_m)_{m<n}$ of disjoint single rays in G.

We claim that $(Q_m)_{m<n}$ is a sequence of disjoint single rays in G of maximum cardinality. Suppose there is a larger sequence of disjoint single rays in G. Since G is a forest, any two single rays in G which share infinitely many (at least two, even) edges or vertices must share a tail. Therefore there is a single ray Q in this larger sequence which only shares finitely many edges and vertices with each Q_m. Then some tail of Q, say xQ, is vertex-disjoint from each Q_m. In particular, xQ is vertex-disjoint from each R_i, i.e., xQ lies in H. Extend xQ to a double ray in H' by first connecting x to the $<_N$-least vertex in its connected component H_i, then following the computable ray which we added. The resulting double ray is disjoint from every S_k, because no S_k can lie in the same H'_i as xQ (for xQ is vertex-disjoint from $S_k \cap H$ by construction). This contradicts maximality of l. □

6.2. **Maximal Variants of IRT.** Instead of sets of disjoint rays of maximum cardinality, we could consider sets of disjoint rays which are maximal. For uncountable graphs, Halin [11] observed that any uncountable maximal set of disjoint rays is in fact of maximum cardinality (because rays are countable). This suggests another variant of IRT, which we call *maximal* IRT:

Definition 6.17. Let $\mathsf{MIRT}_{\mathrm{XYZ}}$ be the statement that every X-graph G has a (possibly finite) sequence $(R_i)_i$ of Y-disjoint Z-rays which is maximal, i.e., for any Z-ray R in G, there is some i such that R and R_i are not Y-disjoint.

$\mathsf{MIRT}_{\mathrm{XYZ}}$ immediately follows from Zorn's Lemma. It is straightforward to show that $\mathsf{MIRT}_{\mathrm{XYZ}}$ implies Π^1_1-CA_0 (see the proof of Theorem 6.18 below), hence $\mathsf{MIRT}_{\mathrm{XYZ}}$ is much stronger than $\mathsf{IRT}_{\mathrm{XYZ}}$ or even $\mathsf{IRT}^*_{\mathrm{XYZ}}$. We show below that $\mathsf{MIRT}_{\mathrm{XYZ}}$ is equivalent to Π^1_1-CA_0. This situation is reminiscent of König's duality theorem: Even though ATR_0 suffices to construct a so-called König cover (Simpson [28]), the existence of a König cover with a certain maximality property is equivalent to Π^1_1-CA_0 (Aharoni, Magidor, Shore [2]).

Theorem 6.18. $\mathsf{MIRT}_{\mathrm{XYZ}}$ *is equivalent to* Π^1_1-CA_0.

After proving the forward direction of the above theorem, we will present two proofs for the backward direction. The first proof is for those familiar with hyperarithmetic theory. The second proof is a standard reverse mathematical proof. Using $\mathsf{MIRT}_{\mathrm{XYZ}}$, we will prove a strong form of Σ^1_1-AC_0 which is known to be equivalent to Π^1_1-CA_0.

Proof that $\mathsf{MIRT}_{\mathrm{XYZ}}$ *implies* Π^1_1-CA_0. We prove that $\mathsf{MIRT}_{\mathrm{XYZ}}$ implies ACA_0 by adapting the proof of Theorem 5.9: If we apply $\mathsf{MIRT}_{\mathrm{XYZ}}$ instead of $\mathsf{IRT}_{\mathrm{XYZ}}$ to any of the forests constructed in the proof, we obtain a sequence containing a Z-ray in each tree which constitutes the forest. This is more than sufficient for carrying out the remainder of the proof of Theorem 5.9.

To prove that $\mathsf{MIRT}_{\mathrm{UVS}}$ implies Π^1_1-CA_0, suppose we are given a set A. Consider the disjoint union of all A-computable trees (this exists, by ACA_0). Any maximal sequence of Y-disjoint rays in this forest must contain a ray in each ill-founded X-computable tree. Hence it computes $\mathcal{O}^A$. This shows that $\mathsf{MIRT}_{\mathrm{UYS}}$ implies Π^1_1-CA_0. To prove that the other $\mathsf{MIRT}_{\mathrm{XYZ}}$ imply Π^1_1-CA_0, it suffices to exhibit a computable procedure which takes trees $T \subseteq N^{<N}$ to X-graphs T' such that T is ill-founded if and only if T' contains a Z-ray. For $\mathsf{MIRT}_{\mathrm{UYD}}$, it suffices to modify each tree by adding a computable branch which is not already on the tree

(as we did in the proof of Theorem 5.9). For $\mathsf{MIRT}_{\mathrm{DYZ}}$, it suffices to orient each of the graphs we constructed above in the obvious way. □

First proof that Π^1_1-CA_0 *implies* $\mathsf{MIRT}_{\mathrm{XYZ}}$. First, we will prove $\mathsf{MIRT}_{\mathrm{XVZ}}$ using Π^1_1-CA_0. Then we will describe how to modify the proof to prove $\mathsf{MIRT}_{\mathrm{XEZ}}$. Suppose we are given an X-graph G. We assume that G is computable; the proof for general graphs G follows by relativization. We will use $\mathcal{O}$ to compute a maximal sequence of disjoint Z-rays in G, as follows. At stage n, we attempt to add a Z-ray which begins at the vertex n and is disjoint from all of the rays we have constructed thus far. We will maintain the fact that the hyperjump of the finite sequence of Z-rays we have constructed is Turing equivalent to $\mathcal{O}$.

Suppose we have constructed disjoint Z-rays $(R_{i_j})_{i_j<n}$ such that for each i_j, R_{i_j} begins at i_j, and $\mathcal{O}^{(R_{i_j})_{i_j<n}} \equiv_T \mathcal{O}$. First, $\mathcal{O}^{(R_{i_j})_{i_j<n}} \equiv_T \mathcal{O}$ can tell us whether there is a Z-ray in G which begins at n and is disjoint from $(R_{i_j})_{i_j<n}$. If not, then we end stage n without constructing any ray. Otherwise, by the Gandy basis theorem, there is some such Z-ray R such that $\mathcal{O}^{(R_{i_j})_{i_j<n}\oplus R} \equiv_T \mathcal{O}$. We add R to our sequence of disjoint Z-rays.

This construction produces a (possibly finite) sequence $R_{i_0}, R_{i_1}, \ldots$ of disjoint Z-rays in G. We show that this sequence is maximal. If R is a Z-ray which is disjoint from every R_{i_j}, then go to stage n of the construction, where n is the first vertex of R. If we did construct some R_{i_j} during stage n, then R_{i_j} would not be disjoint from R. Hence we did not construct any Z-ray during stage n. But R is a Z-ray that begins at n and is disjoint from $(R_{i_j})_{i_j<n}$, contradiction.

To prove $\mathsf{MIRT}_{\mathrm{XEZ}}$, we modify the above construction as follows. At stage (u, v), we search for a Z-ray $\langle H, f\rangle$ such that $(f(0), f(1)) = (u, v)$ and $\langle H, f\rangle$ is disjoint from all Z-rays constructed thus far. The rest of the proof proceeds as above. □

Second proof that Π^1_1-CA_0 *implies* $\mathsf{MIRT}_{\mathrm{XYZ}}$. We will prove that $\mathsf{MIRT}_{\mathrm{XYZ}}$ implies *strong* Σ^1_1-DC_0, which consists of the scheme

$$(\exists Z)(\forall n)(\forall Y)\left(\eta\left(n, \bigoplus_{i<n} Z^{[i]}, Y\right) \to \eta\left(n, \bigoplus_{i<n} Z^{[i]}, Z^{[n]}\right)\right),$$

for any Σ^1_1 formula $\eta(n, X, Y)$. It is known that strong Σ^1_1-DC_0 and Π^1_1-CA_0 are equivalent (Simpson [29, VII.6.9]). Note that in contrast to Σ^1_1-DC_0, the premise of strong Σ^1_1-DC_0 does not assume that for all n and X, there is some Y such that $\eta(n, X, Y)$ holds. Furthermore, the conclusion of strong Σ^1_1-DC_0 does not place any restriction on $Z^{[n]}$, if there is no Y such that $\eta(n, \bigoplus_{i<n} Z^{[i]}, Y)$ holds.

We prove $\mathsf{MIRT}_{\mathrm{XVZ}}$ using strong Σ^1_1-DC_0 as follows. Given an X-graph G, we will define an arithmetical formula $\eta(n, X, Y)$ with parameter G. First, we inductively define a finite sequence $i_0, \ldots, i_k < n$. If $i_0, \ldots, i_{j-1}$ have been defined, define i_j to be the least number (if any) above i_{j-1} and below n such that $X^{[i_j]}$ is a Z-ray in G which is disjoint from $X^{[i_0]}, \ldots, X^{[i_{j-1}]}$. It is clear that there is an arithmetical formula with parameter G which defines $i_0, \ldots, i_k$ from n and X. Next, we say that $\eta(n, X, Y)$ holds if Y is a Z-ray in G which begins with n, and Y is disjoint from $X^{[i_0]}, \ldots, X^{[i_k]}$.

Apply strong Σ^1_1-DC_0 for the formula η to obtain some set Z. By Σ^0_1-comprehension with parameter $G \oplus Z$, we may inductively define a (possibly finite) sequence $i_0, i_1, \ldots$, just as we did in the definition of η. Clearly $(Z^{[i_j]})_j$ is a sequence of disjoint Z-rays in G. We claim that it is maximal.

Suppose that R is a Z-ray in G which is disjoint from every $Z^{[i_j]}$. Suppose that R begins with vertex n. Then R is disjoint from $Z^{[i_0]}, \dots, Z^{[i_k]}$, where i_k is the largest i_j below n. It follows that $\eta(n, \bigoplus_{i<n} Z^{[i]}, R)$ holds. So $\eta(n, \bigoplus_{i<n} Z^{[i]}, Z^{[n]})$ holds, i.e., $Z^{[n]}$ is a Z-ray which begins with n and $Z^{[n]}$ is disjoint from $Z^{[i_0]}, \dots, Z^{[i_k]}$. By definition of i_{k+1}, that means that $n = i_{k+1}$. But then R and $Z^{[i_{k+1}]}$ are not disjoint, contradiction.

To prove $\mathsf{MIRT}_{\mathrm{XEZ}}$, we modify the above construction as follows. We say that $\eta((u, v), X, Y)$ holds if Y is a Z-ray $\langle H, f\rangle$ in G such that $(f(0), f(1)) = (u, v)$ and Y is disjoint from $X^{[i_0]}, \dots, X^{[i_k]}$. The rest of the proof proceeds as above. □

7. Relationships Between IRT and Other Theories of Hyperarithmetic Analysis

In this section, we establish implications and nonimplications between variants of IRT and theories of hyperarithmetic analysis other than Σ^1_1-AC_0. One such theory is as follows:

Definition 7.1. The theory *unique*-Σ^1_1-AC_0 consists of RCA_0 and

$$(\forall n)(\exists! X)A(n, X) \to (\exists (X_n)_n)(\forall n)A(n, X_n)$$

for each arithmetical formula $A(n, X)$.

The above theory is typically known as weak-Σ^1_1-AC_0 (e.g., [29, VIII.4.12]). Our reason for deviating from this terminology will soon be clear.

Let us sketch a proof that $\mathsf{IRT}^*_{\mathrm{UVS}}$ implies unique-Σ^1_1-AC_0. (We will modify this sketch to prove a stronger result in Theorem 7.3.) By Lemma 6.7, it suffices to prove that for any sequence $(T_n)_n$ of subtrees of $N^{<N}$ such that each T_n has a unique branch P_n, the sequence $(P_n)_n$ exists. Analogously to the proof of (5) → (1) in Theorem 6.13, we may construct a sequence of trees $(S_n)_n$ such that for each n, the unique branch on S_n is $P_0 \oplus \cdots \oplus P_n$. By $\mathsf{IRT}^*_{\mathrm{UVS}}$, there is a sequence $(R_k)_k$ of disjoint rays in $\bigsqcup_n S_n$ of maximum cardinality. We claim that $(R_k)_k$ is infinite. If not, only finitely many S_n contain any R_k. Then we can increase the cardinality of $(R_k)_k$ by adding any ray from any S_m which does not contain any R_k, contradiction. Therefore $(R_k)_k$ is infinite. Since each S_n contains at most one R_k, it follows that infinitely many S_n contain some R_k. Thus we may construct the sequence $(P_n)_n$ (analogously to the proof of (5) → (1) in Theorem 6.13).

In the above sketch, we used the property that any two rays in some S_n must intersect, which follows from the assumption that each T_n has a unique branch. We could carry out the above sketch even if each T_n has finitely many branches, rather than a unique branch. This motivates the following definition:

Definition 7.2. The theory *finite*-Σ^1_1-AC_0 consists of RCA_0 and

$$(\forall n)(\exists \text{ nonzero finitely many } X)A(n, X) \to (\exists (X_n)_n)(\forall n)A(n, X_n)$$

for each arithmetic formula $A(n, X)$. Formally, "$(\exists$ nonzero finitely many $X)A(n, X)$" means that there is a nonempty sequence $(X_i)_{i<j}$ such that for each X, $A(n, X)$ holds if and only if $X = X_i$ for some $i < j$.

Since Σ^1_1-AC_0 implies finite-Σ^1_1-AC_0 which in turn implies unique-Σ^1_1-AC_0, it follows that finite-Σ^1_1-AC_0 is a THA. Goh [10] shows that finite-Σ^1_1-AC_0 is strictly stronger than unique-Σ^1_1-AC_0.

Theorem 7.3. $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *implies finite-*Σ^1_1-AC_0 *over* RCA_0. *(It follows that* $\mathsf{IRT}_{\mathrm{XYZ}}$ *implies finite-*Σ^1_1-AC_0 *over* $\mathsf{RCA}_0 + I\Sigma^1_1$*, but this is superseded by Theorem 7.7 below.)*

Proof. To prove the implication from $\mathsf{IRT}^*_{\mathrm{UYS}}$, modify the sketch above by replacing "unique" with "finitely many". This shows that $\mathsf{IRT}^*_{\mathrm{UYS}}$ for forests implies finite-Σ^1_1-AC_0. By Theorem 6.16, it follows that $\mathsf{IRT}^*_{\mathrm{UYD}}$ for forests implies finite-Σ^1_1-AC_0. By Proposition 6.5, it follows that $\mathsf{IRT}^*_{\mathrm{DYZ}}$ implies finite-Σ^1_1-AC_0 as well. □

Another theory of hyperarithmetic analysis which follows from $\mathsf{IRT}^*_{\mathrm{XYZ}}$ is *arithmetic Bolzano-Weierstrass* (ABW_0):

Definition 7.4. The theory ABW_0 consists of RCA_0 and the following statement: If $A(X)$ is an arithmetic predicate on 2^N, either there are finitely many X such that $A(X)$ holds, or the set $\{X : A(X)\}$ has an accumulation point.

Friedman [8] introduced ABW_0 and asserted that it follows from Σ^1_1-AC_0 (with unrestricted induction). Conidis [4] proved Friedman's assertion and established relationships between ABW_0 and most known theories of hyperarithmetic analysis. Goh [10] shows that $\mathsf{ABW}_0+I\Sigma^1_1$ implies finite-Σ^1_1-AC_0. We do not know if ABW_0 is strictly stronger than finite-Σ^1_1-AC_0.

The following two lemmas will be useful in deriving ABW_0 from $\mathsf{IRT}^*_{\mathrm{XYZ}}$. The first lemma describes a connection between sets of solutions of arithmetic predicates and disjoint rays in trees.

Lemma 7.5 (ACA_0)**.** *Suppose $A(X)$ is an arithmetic predicate. Then there is a tree $T \subseteq N^{<N}$ such that if there is a sequence of distinct solutions of $A(X)$, then there is a sequence of Y-disjoint single rays in T of the same cardinality, and vice versa.*

Proof. By Lemma 6.6, there is a tree $T \subseteq N^{<N}$ such that

$$\forall X(A(X) \leftrightarrow \exists f(\langle X, f\rangle \in [T])$$
$$\text{and } \forall X(\exists \text{ at most one } f)(\langle X, f\rangle \in [T]).$$

If $(X_i)_i$ is a sequence of distinct solutions of $A(X)$, then there is a sequence of distinct branches $(\langle X_i, f_i\rangle)_i$ on T of the same cardinality. By taking an appropriate tail of each branch, we obtain a sequence of vertex-disjoint (hence edge-disjoint) single rays in T of the same cardinality.

Conversely, suppose there is a sequence $(R_i)_i$ of Y-disjoint single rays in T. For each R_i, we define a branch on T which corresponds to it as follows. Let x be the vertex in R_i which is closest to the root of T. Then we can extend xR_i to the root to obtain a branch $\langle X_i, f_i\rangle$ on T. We claim that $(X_i)_i$ is a sequence of distinct solutions of $A(X)$. For each $i \neq j$, since R_i and R_j are Y-disjoint, they cannot share any tail. So $\langle X_i, f_i\rangle$ and $\langle X_j, f_j\rangle$ must be distinct. Since for each X, there is at most one f such that $\langle X, f\rangle$ is a branch on T, it follows that $X_i \neq X_j$ as desired. □

The second lemma is essentially the well-known fact that the Bolzano-Weierstrass theorem is provable in ACA_0:

Lemma 7.6 (see Conidis [4, pg. 4476])**.** ACA_0 *proves that if $(X_n)_n$ is a sequence of distinct elements of 2^N, then there is some Z which is an accumulation point of $\{X_n : n \in N\}$.*

Theorem 7.7. $\mathsf{IRT}^*_{\mathrm{XYZ}}$ *implies* ABW_0 *over* RCA_0. *Therefore* $\mathsf{IRT}_{\mathrm{XYZ}}$ *implies* ABW_0 *over* $\mathsf{RCA}_0 + I\Sigma^1_1$.

Proof. By Proposition 6.5, it suffices to show that the undirected variants of IRT^* imply ABW_0.

Suppose $A(X)$ is an arithmetic predicate on 2^N which does not have finitely many solutions. By Lemma 7.5, there is a tree $T \subseteq N^{<N}$ such that for any sequence of distinct solutions of $A(X)$, there is a sequence of Y-disjoint single rays in T of the same cardinality, and vice versa.

By $\mathsf{IRT}^*_{\mathrm{UYS}}$, or by $\mathsf{IRT}^*_{\mathrm{UYD}}$ and Theorem 6.16, there is a sequence of Y-disjoint single rays in T of maximum cardinality. This yields a sequence of distinct solutions of $A(X)$ of the same cardinality.

If this sequence is finite, then there is a solution Y of $A(X)$ not in the sequence, because $A(X)$ does not have finitely many solutions. Hence there is a sequence of distinct solutions of $A(X)$ of larger cardinality, which yields a sequence of Y-disjoint single rays in T of larger cardinality. Contradiction.

Therefore there is an infinite sequence $(X_n)_n$ of distinct solutions of A. By Lemma 7.6, there is an accumulation point of $\{X_n : n \in N\}$, which is of course an accumulation point of $\{X : A(X)\}$, as desired. □

We turn our attention to nonimplications. One prominent theory of hyperarithmetic analysis is the scheme of Δ^1_1-*comprehension* (studied by Kreisel [15]):

Definition 7.8. The theory Δ^1_1-CA_0 consists of RCA_0 and the statement

$$(\forall n)(\Phi(n) \leftrightarrow \neg\Psi(n)) \rightarrow \exists X(n \in X \leftrightarrow \Phi(n))$$

for all Σ^1_1 formulas Φ and Ψ.

Theorem 7.9. Δ^1_1-$\mathsf{CA}_0 \nvdash \mathsf{IRT}_{\mathrm{XYZ}}, \mathsf{IRT}^*_{\mathrm{XYZ}}$.

Proof. Conidis [4, Theorem 3.1] constructed a standard model which satisfies Δ^1_1-CA_0 but not ABW_0. By Theorem 7.7, this model does not satisfy $\mathsf{IRT}^*_{\mathrm{XYZ}}$. Since standard models satisfy full induction, this model does not satisfy $\mathsf{IRT}_{\mathrm{XYZ}}$ either (by Proposition 6.3). □

Theorem 7.10. $\mathsf{ABW}_0 \nvdash \mathsf{IRT}_{\mathrm{XYZ}}, \mathsf{IRT}^*_{\mathrm{XYZ}}$.

Proof. By Propositions 5.3 and 6.3, it suffices to show that $\mathsf{ABW}_0 \nvdash \mathsf{IRT}_{\mathrm{UYZ}}$. Van Wesep [31, I.1] constructed a standard model $\mathcal{N}$ which satisfies unique-Σ^1_1-AC_0 but not Δ^1_1-CA_0. Conidis [4, Theorem 4.1] showed that $\mathcal{N}$ satisfies ABW_0. We show below that $\mathcal{N}$ does not satisfy $\mathsf{IRT}_{\mathrm{UYZ}}$.

In order to define $\mathcal{N}$, van Wesep constructed a tree T^G and branches $(f^G_i)_{i\in\mathbb{N}}$ of T^G such that (1) $\mathcal{N}$ contains T^G and infinitely many (distinct) f^G_i (see [31, pg. 13 l. 1–11]); (2) $\mathcal{N}$ does not contain any infinite sequence of distinct branches of T^G (see [31, pg. 12 l. 7–9] and Steel [30, Lemma 7].) Then T^G is an instance of $\mathsf{IRT}_{\mathrm{UYS}}$ in $\mathcal{N}$ which has no solution in $\mathcal{N}$. This shows that $\mathcal{N}$ does not satisfy $\mathsf{IRT}_{\mathrm{UYS}}$ for trees.

Since $\mathcal{N}$ is a standard model, it satisfies full induction. By Theorem 6.16, it follows that $\mathcal{N}$ does not satisfy $\mathsf{IRT}_{\mathrm{UYD}}$ for forests. □

Figure 4 illustrates some of our results. In order to simplify the diagram, we have omitted all variants of IRT except $\mathsf{IRT}_{\mathrm{UVS}}$.

8. Isolating the Use of Σ^1_1-AC_0 in Proving IRT

We isolate the use of Σ^1_1-AC_0 in our proofs of $\mathsf{IRT}_{\mathrm{XYS}}$ and $\mathsf{IRT}_{\mathrm{UVD}}$ (Theorems 4.5, 5.10, 5.15) by identifying the following principles:

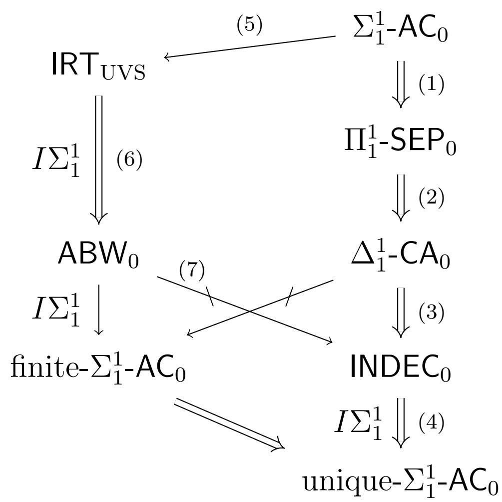

FIGURE 4. Partial zoo of theories of hyperarithmetic analysis. Single arrows indicate implication while double arrows indicate strict implication. The references for the above results are as follows: (1, 2) Montalbán [17, Theorems 2.1, 3.1]; (3, 4) Montalbán [16, Theorem 2.2], Neeman [20, Theorems 1.2, 1.3, 1.4], see also Neeman [21, Theorem 1.1]; (5) Theorem 4.5; (6) Theorems 7.7, 7.10; (7) Conidis [4, Theorem 4.1]. All results concerning finite-Σ^1_1-AC_0 are in Goh [10].

Definition 8.1. Let $\mathsf{SCR}_{\mathrm{XYZ}}$ be the assertion that if G is an X-graph with arbitrarily many Y-disjoint Z-rays, then there is a sequence of sets $(X_k)_k$ such that for each $k \in N$, X_k is a set of k Y-disjoint Z-rays in G.

Let $\mathsf{WIRT}_{\mathrm{XYZ}}$ be the assertion that if G is an X-graph and there is a sequence of sets $(X_k)_k$ such that for each $k \in N$, X_k is a set of k Y-disjoint Z-rays in G, then G has infinitely many Y-disjoint Z-rays.

SCR stands for Strongly Collecting Rays. WIRT stands for Weak Infinite Ray Theorem.

It is clear that Σ^1_1-AC_0 implies $\mathsf{SCR}_{\mathrm{XYZ}}$ and $\mathsf{SCR}_{\mathrm{XYZ}} + \mathsf{WIRT}_{\mathrm{XYZ}}$ implies $\mathsf{IRT}_{\mathrm{XYZ}}$. The only use of Σ^1_1-AC_0 in our proofs of $\mathsf{IRT}_{\mathrm{XYS}}$ and $\mathsf{IRT}_{\mathrm{UVD}}$ is to prove $\mathsf{SCR}_{\mathrm{XYS}}$ and $\mathsf{SCR}_{\mathrm{UVD}}$ respectively:

Theorem 8.2. ACA_0 *proves* $\mathsf{WIRT}_{\mathrm{XYS}}$ *and* $\mathsf{WIRT}_{\mathrm{UVD}}$.

Proof. For $\mathsf{WIRT}_{\mathrm{UVS}}$, see the proof of Theorem 4.5. For $\mathsf{WIRT}_{\mathrm{UVD}}$, see the proof of Theorem 5.10. For $\mathsf{WIRT}_{\mathrm{XES}}$, see the proof of Theorem 5.15. The desired result for $\mathsf{WIRT}_{\mathrm{DVS}}$ then follows from Lemma 5.4. □

Next we will use the above result to show that $\mathsf{SCR}_{\mathrm{XYS}}$ and $\mathsf{SCR}_{\mathrm{UVD}}$ are equivalent over RCA_0 to $\mathsf{IRT}_{\mathrm{XYS}}$ and $\mathsf{IRT}_{\mathrm{UVD}}$ respectively. First, observe that $\mathsf{IRT}_{\mathrm{XYZ}}$ implies $\mathsf{SCR}_{\mathrm{XYZ}}$ for each choice of XYZ. Second, Lemmas 5.2, 5.4 and 5.6 imply

Proposition 8.3. $\mathsf{SCR}_{\mathrm{DYZ}}$ *implies* $\mathsf{SCR}_{\mathrm{UYZ}}$*,* $\mathsf{SCR}_{\mathrm{DEZ}}$ *implies* $\mathsf{SCR}_{\mathrm{DVZ}}$ *and* $\mathsf{SCR}_{\mathrm{DYD}}$ *implies* $\mathsf{SCR}_{\mathrm{DYS}}$.

Proposition 8.4. $\mathsf{SCR}_{\mathrm{XYZ}}$ *implies* ACA_0.

Proof. By Proposition 8.3, it suffices to establish the desired result for the undirected variants of SCR. The proofs are almost identical to those of Proposition 4.6 and Theorem 5.9. There, we applied $\mathsf{IRT}_{\mathrm{UYZ}}$ to forests $G := \bigsqcup_n G_n$, where each G_n contains a Z-ray, and no two Z-rays in G_n can be Y-disjoint. Any infinite sequence of Y-disjoint Z-rays in G must contain a Z-ray in cofinally many graphs G_n. Therefore from such a sequence we can uniformly compute Z-rays in cofinally many graphs G_n, which establishes ACA_0 by the construction of $\bigsqcup_n G_n$. If we assume $\mathsf{SCR}_{\mathrm{UYZ}}$ instead of $\mathsf{IRT}_{\mathrm{UYZ}}$, we only have access to a sequence $(X_k)_k$ such that for each k, X_k is a set of k Y-disjoint Z-rays in G. From such a sequence we can still uniformly compute Z-rays in cofinally many graphs G_n, because for any k, X_{k+1} must contain a Z-ray in some G_n, $n \geq k$. □

By Theorem 8.2, Proposition 8.4, and the observation that $\mathsf{SCR}_{\mathrm{XYZ}} + \mathsf{WIRT}_{\mathrm{XYZ}} \vdash \mathsf{IRT}_{\mathrm{XYZ}}$, we obtain

Corollary 8.5. $\mathsf{SCR}_{\mathrm{XYZ}}$ *and* $\mathsf{IRT}_{\mathrm{XYZ}}$ *are equivalent over* RCA_0 *for the following choices of XYZ: XYS and UVD.*

We now turn our attention to $\mathsf{WIRT}_{\mathrm{XYZ}}$. As usual, Lemmas 5.2, 5.4 and 5.6 imply

Proposition 8.6. $\mathsf{WIRT}_{\mathrm{DYZ}}$ *implies* $\mathsf{WIRT}_{\mathrm{UYZ}}$*,* $\mathsf{WIRT}_{\mathrm{DEZ}}$ *implies* $\mathsf{WIRT}_{\mathrm{DVZ}}$ *and* $\mathsf{WIRT}_{\mathrm{DYD}}$ *implies* $\mathsf{WIRT}_{\mathrm{DYS}}$*.*

Recall that $\mathsf{WIRT}_{\mathrm{XYS}}$ and $\mathsf{WIRT}_{\mathrm{UVD}}$ are provable in ACA_0 (Theorem 8.2). $\mathsf{WIRT}_{\mathrm{DVD}}$ and $\mathsf{WIRT}_{\mathrm{DED}}$ are open, because Σ^1_1-$\mathsf{AC}_0 + \mathsf{WIRT}_{\mathrm{XYZ}}$ implies $\mathsf{IRT}_{\mathrm{XYZ}}$, and $\mathsf{IRT}_{\mathrm{DVD}}$ and $\mathsf{IRT}_{\mathrm{DED}}$ are open (see comments after Theorem 5.1). We do not have an upper bound on the proof-theoretic strength of $\mathsf{WIRT}_{\mathrm{UED}}$ (an upper bound on $\mathsf{WIRT}_{\mathrm{UED}}$ would yield an upper bound on $\mathsf{IRT}_{\mathrm{UED}}$, which we do not currently have).

We do not know if any $\mathsf{WIRT}_{\mathrm{XYZ}}$ is equivalent to ACA_0. In an effort to clarify the situation, we define an apparent strengthening of $\mathsf{WIRT}_{\mathrm{XYZ}}$ and show that it implies ACA_0:

Definition 8.7. Let nonuniform-$\mathsf{WIRT}_{\mathrm{XYZ}}$ be the assertion that if G is an X-graph and there is a sequence of Z-rays $R_0, R_1, \dots$ in G such that for each k, there are $i_0, \dots, i_k$ such that $R_{i_0}, \dots, R_{i_k}$ are pairwise Y-disjoint, then G has infinitely many Y-disjoint Z-rays.

Every instance of $\mathsf{WIRT}_{\mathrm{XYZ}}$ is also an instance of nonuniform-$\mathsf{WIRT}_{\mathrm{XYZ}}$, so nonuniform-$\mathsf{WIRT}_{\mathrm{XYZ}}$ implies $\mathsf{WIRT}_{\mathrm{XYZ}}$. Conversely, we have

Proposition 8.8. $\mathsf{ACA}_0 + \mathsf{WIRT}_{\mathrm{XYZ}}$ *implies nonuniform-*$\mathsf{WIRT}_{\mathrm{XYZ}}$*.*

Proof. Suppose G is an instance of nonuniform-$\mathsf{WIRT}_{\mathrm{XYZ}}$, i.e., G is an X-graph and $(R_n)_n$ is a sequence of Z-rays in G such that for each k, there are $i_0, \dots, i_k$ such that $R_{i_0}, \dots, R_{i_k}$ are pairwise Y-disjoint. Then ACA_0 can find such $i_0, \dots, i_k$ uniformly in k. Therefore by ACA_0, G is an instance of $\mathsf{WIRT}_{\mathrm{XYZ}}$. By $\mathsf{WIRT}_{\mathrm{XYZ}}$, G has infinitely many Y-disjoint Z-rays as desired. □

Theorem 8.9. *Nonuniform-*$\mathsf{WIRT}_{\mathrm{XYZ}}$ *implies* ACA_0 *over* RCA_0*. It follows that nonuniform-*$\mathsf{WIRT}_{\mathrm{XYS}}$ *and nonuniform-*$\mathsf{WIRT}_{\mathrm{UVD}}$ *are both equivalent to* ACA_0 *over* RCA_0*.*

Proof. By Proposition 8.8 and Theorem 8.2, ACA_0 implies nonuniform-$\mathsf{WIRT}_{\mathrm{XYS}}$ and nonuniform-$\mathsf{WIRT}_{\mathrm{UVD}}$.

Next, we show that nonuniform-$\mathsf{WIRT}_{\mathrm{XYZ}}$ implies ACA_0. By Lemma 5.2, it suffices to consider the undirected versions of nonuniform-WIRT. First, we shall prove that nonuniform-$\mathsf{WIRT}_{\mathrm{UYS}}$ implies ACA_0 by constructing a computable instance of nonuniform-$\mathsf{WIRT}_{\mathrm{UYS}}$ such that every nonuniform-$\mathsf{WIRT}_{\mathrm{UYS}}$ solution computes $\emptyset'$. (The desired result follows by relativization.) We will compute $\emptyset'$ by computing, for each i, some stage n by which the standard enumeration of $\emptyset'$ has enumerated all numbers in $\emptyset'$ which are less than i, i.e., $\emptyset'_n \upharpoonright i = \emptyset' \upharpoonright i$.

Construction of $G = (V, E)$: Our graph G will be on N^2. We construct E in stages as follows. First, initialize the set of dead columns to be the empty set. At stage s, suppose i enters $\emptyset'$. Let $t < s$ be the most recent stage at which any number less than i enters $\emptyset'$ (if no such stage exists, then $t = 0$). For all n such that $t < n < s$ and column n is not dead:

(1) add the edge between $\langle n, s\rangle$ and $\langle t, s+1\rangle$ to E;
(2) declare column n to be dead.

For all other n such that column n is not dead, we add the edge between $\langle n, s\rangle$ and $\langle n, s+1\rangle$ to E. This completes stage s of the construction of E.

In order to construct the single rays R_n, observe that if column m is not dead at the beginning of stage s, then, at the end of stage s, $\langle m, s\rangle$ is adjacent to $\langle t, s+1\rangle$ for some unique t. Furthermore, we can show by induction that column t is not dead at the beginning of stage $s+1$. So we can define R_n by starting at $\langle n, 0\rangle$ and recursively picking the unique adjacent vertex at the next level. (Note that this ensures that any two R_n and R_m are vertex-disjoint if and only if they are edge-disjoint.)

Verification: G and $(R_n)_n$ can be defined in RCA_0. Observe that RCA_0 proves that for each i, there is some least stage n_i such that $\emptyset'_{n_i} \upharpoonright i = \emptyset' \upharpoonright i$.

We claim that column n_i is never dead. At stages $s \leq n_i$, we cannot declare column n_i to be dead. As for stages $s > n_i$, suppose some j enters $\emptyset'$ at stage s. Since $\emptyset'_{n_i} \upharpoonright i = \emptyset' \upharpoonright i$, we have $j \geq i$. By minimality of n_i, some number less than i ($\leq j$) entered $\emptyset'$ at stage n_i. So we would not declare column n_i to be dead at stage s.

It follows that the vertices of R_{n_i} are $\langle n_i, s\rangle$, for $s \in N$. With this we can show that RCA_0 proves that G is an instance of nonuniform-$\mathsf{WIRT}_{\mathrm{XYS}}$: RCA_0 proves that for each k, the sequence $n_0, n_1, \ldots, n_{k-1}$ exists, and hence the sequence $(R_{n_i})_{i<k}$ exists (and is Y-disjoint).

Next, we shall prove that every column n which is not any n_i is eventually declared dead. RCA_0 proves that there is some maximum i such that $\emptyset'_n \upharpoonright i = \emptyset' \upharpoonright i$. By choice of i and our assumption on n, $n > n_i$. Consider the stage $s > n$ at which i enters $\emptyset'$. Then the most recent stage $t < s$ at which any number less than i enters $\emptyset'$ must be n_i. So we declared column n to be dead at stage s.

It is clear from the construction that if a ray R in G begins at some $\langle n, s\rangle$, then R can only intersect columns $m \in [n_i, n_{i+1})$, where i is defined as in the previous paragraph. Since column n_i is the only column among $[n_i, n_{i+1})$ which is never declared dead, R must share a tail with R_{n_i}.

With the above, we can conclude the proof that nonuniform-$\mathsf{WIRT}_{\mathrm{XYS}}$ implies that $\emptyset'$ exists. By nonuniform-$\mathsf{WIRT}_{\mathrm{XYS}}$, let $(S_i)_i$ be an infinite sequence of Y-disjoint single rays in G. For each i, define s_i to be the first coordinate of the first vertex of S_i. By thinning out $(s_i)_i$ if necessary, we may assume that $s_0 \leq s_1 \leq \ldots$.

We shall prove by induction that for each i, $s_{i+1} \geq n_i$ (hence $\emptyset' \upharpoonright i = \emptyset'_{s_{i+1}} \upharpoonright i$). The base case is trivial. Suppose $s_i \geq n_{i-1}$. If $s_i \geq n_i$, then we have $s_{i+1} \geq s_i \geq n_i$ as desired. Otherwise, $s_i \in [n_{i-1}, n_i)$. We cannot have $s_{i+1} \in [n_{i-1}, n_i)$ as well, because that would imply that R_{s_i} and $R_{s_{i+1}}$ both share a tail with $R_{n_{i-1}}$ and are hence not Y-disjoint.

Therefore we can use the sequence $(s_i)_i$ to define $\emptyset'$, as desired.

To show that nonuniform-$\mathsf{WIRT}_{\mathrm{UYD}}$ implies ACA_0, define G as above. Consider the graph G' on $N \times Z$ which contains G and its "reflection about the x-axis". For any single ray in G, we can consider its associated double ray in G', obtained by joining the ray with its reflection. This map from single rays in G to double rays in G' preserves Y-disjointness. Therefore, since G is an instance of nonuniform-$\mathsf{WIRT}_{\mathrm{UYS}}$, it follows that G' is an instance of nonuniform-$\mathsf{WIRT}_{\mathrm{UYD}}$. By nonuniform-$\mathsf{WIRT}_{\mathrm{UYD}}$, let $(S_i)_i$ be an infinite sequence of Y-disjoint double rays in G'. As before, we can define $(s_i)_i$ and use it to compute $\emptyset'$. □

We are unable to show that $\mathsf{WIRT}_{\mathrm{XYZ}}$ implies ACA_0, but we can prove

Theorem 8.10. $\mathsf{WIRT}_{\mathrm{XYZ}}$ *is not provable in* RCA_0.

Proof. By Proposition 8.6, it suffices to consider the undirected variants of WIRT. We shall present the proof for $\mathsf{WIRT}_{\mathrm{UYS}}$ before indicating how to modify it to obtain the proof for $\mathsf{WIRT}_{\mathrm{UYD}}$.

For $\mathsf{WIRT}_{\mathrm{UYS}}$, it suffices to construct a computable graph G on $\mathbb{N}$ and a computable sequence $((X_i^k)_{i<k})_k$ such that (1) for each $k \in \mathbb{N}$, the X_i^k for $i < k$ are pairwise vertex-disjoint single rays in G; (2) there is no computable sequence $(R_j)_j$ of edge-disjoint single rays in G. This will be a finite injury priority argument. At the end of stage s of our construction, for each $i < k \leq s$, we will have defined a path P_i^k containing at least $s - k + 1$ vertices which is intended to be an initial segment of the single ray X_i^k. For each $k \leq s$, we will obey the *disjointness rule*, namely, the paths $P_0^k, \ldots, P_{k-1}^k$ will be vertex-disjoint. In future stages, we will not add any edges between vertices which are currently in $P_0^k, \ldots, P_{k-1}^k$. Thus G will be a computable graph given by the union of the single rays X_i^k.

Apart from the above stipulations, we will satisfy the requirements

$$\begin{aligned} \mathcal{Q}_e : &\text{ if } R_0, R_1, \ldots \text{ is a sequence of single rays in } G \text{ defined by } \Phi_e, \\ &\text{then } R_0, R_1, \ldots \text{ are not edge-disjoint.} \end{aligned}$$

Arrange the requirements $\mathcal{Q}_0, \mathcal{Q}_1, \ldots$ in order of priority. During our construction, we will attempt to satisfy each $\mathcal{Q}_e$ by *merging* certain rays X_i^k and X_j^l, i.e., we take two new numbers and append them consecutively to both P_i^k and P_j^l. We also ensure that P_i^k and P_j^l henceforth agree.

Construction. At stage s of the construction, we are given a finite graph G_s consisting of, for each $k < s$, finite vertex-disjoint paths $P_0^k, \ldots, P_{k-1}^k$ which are intended to be initial segments of the rays $X_0^k, \ldots, X_{k-1}^k$.

First, for each such path in order we extend it by taking the least new number and appending it to the end of the path. (If some P_i^k and P_j^l have the same endpoint, we append the same new vertex to both of those paths.) Second, for each $i < s$, in order, we take the least new number and designate it as the first vertex of a new path P_i^s. To complete stage s of the construction, we will act for the requirement $\mathcal{Q}_e$ of highest priority which requires attention (if any), as defined below. We will attempt to satisfy each $\mathcal{Q}_e$ in two phases. Whenever $\mathcal{Q}_e$ is initialized, it lies in phase 1. We say that $\mathcal{Q}_e$ requires attention at stage s if it has not been (permanently) satisfied and either of the following hold:

(1) $\mathcal{Q}_e$ is in phase 1 and $\Phi_e(\langle a, 0\rangle)$ converges and equals $P_i^k(m)$ for some $a, i, k, m < s$ which are greater than the last stage $s_0(e)$ at which $\mathcal{Q}_e$ was initialized.
(2) $\mathcal{Q}_e$ is in phase 2 and $\Phi_e(\langle b, 0\rangle)$ converges and equals $P_j^l(n)$ for some $b, j, l, n < s$ which are greater than the last stage $s_1(e)$ at which $\mathcal{Q}_e$ was in phase 1.

Our action for $\mathcal{Q}_e$ in each case is as follows:

(1) We initialize all requirements of priority lower than $\mathcal{Q}_e$ which are not satisfied. We also declare that $\mathcal{Q}_e$ is in phase 2.
(2) We merge the rays X_i^k and X_j^l by appending the least two new numbers to the end of both P_i^k and P_j^l. (For any paths P_n^m which share an endpoint with P_i^k or P_j^l, we append the same new vertices to P_n^m.) We also initialize all requirements of priority lower than $\mathcal{Q}_e$ which are not satisfied, and declare that $\mathcal{Q}_e$ is (permanently) satisfied.

This completes stage s of the construction.

Verification. It is clear that, for each $i < k$, each X_i^k is a single ray and that G is a computable graph consisting of the union of these rays. Each requirement acts at most twice after each time it is initialized, so each requirement acts only finitely often, by induction.

To see that we never violate the disjointness rule, it suffices to show that we never merge any distinct rays X_n^m and $X_{n'}^m$, whether directly or indirectly. Suppose not. Consider the least stage s at which we have done so, say via action of some $\mathcal{Q}_e$. We have $s_0(e) < s_1(e) < s$.

We analyze which rays can be merged by requirements at stages $\leq s$. Requirements of priority higher than $\mathcal{Q}_e$ can only merge rays X_i^k and X_j^l if $k, l < s_0(e)$, otherwise we would initialize $\mathcal{Q}_e$ at some stage between $s_0(e)$ and s. Next, the rays X_i^k and X_j^l merged by $\mathcal{Q}_e$ satisfy $s_0(e) < k, l < s$. Finally, requirements of priority lower than $\mathcal{Q}_e$ can only merge rays X_i^k and X_j^l if $k, l < s_0(e)$, or $s_0(e) < k, l < s_1(e)$, or $s_1(e) < k, l < s$, because they are initialized at stages $s_0(e)$ and $s_1(e)$.

We now consider cases depending on the value of m.

<u>Case 1.</u> $m < s_0(e)$. Then X_n^m can only be (directly or indirectly) merged with rays X_j^l where $l \leq s_0(e)$. Likewise for $X_{n'}^m$. So the merger performed by $\mathcal{Q}_e$ at stage s (which merges some X_i^k and X_j^l, where $k, l > s_0(e)$) cannot cause X_n^m and $X_{n'}^m$ to be merged, contradicting minimality of s.

<u>Case 2.</u> $s_0(e) < m < s_1(e)$. Then X_n^m and $X_{n'}^m$ can only be (directly or indirectly) merged with rays X_j^l where $s_0(e) < l < s_1(e)$. The rest of the argument follows that in Case 1.

<u>Case 3.</u> $s_1(e) < m < s$. Similar to Case 2.

<u>Case 4.</u> $m = s_0(e)$ or $m = s_1(e)$. Then X_n^m and $X_{n'}^m$ are never merged with any ray (at stages $\leq s$), contradicting our choice of X_n^m and $X_{n'}^m$.

This proves that we never (directly or indirectly) merge distinct rays X_n^m and $X_{n'}^m$.

In order to prove that each requirement $\mathcal{Q}_e$ is satisfied, observe the following. Suppose a ray R in G begins at some vertex in X_i^k which lies in $G_{s+1} - G_s$.

(1) If X_i^k is never merged with any ray after stage s, then R shares a tail with X_i^k.
(2) Whenever X_i^k is merged with some ray after stage s, R has to pass through the edge between the two vertices we added in the merger.

Now, suppose Φ_e is total. Consider the possible actions for $\mathcal{Q}_e$ after (the final value of) stage $s_0(e)$. If we never define $s_1(e)$, then every value of $\Phi_e(\langle a, 0\rangle)$ for $a > s_0(e)$ lies in some X_i^k, for $i < k \leq s_0(e)$. In this case, there are $a \neq a'$ such that $\Phi_e(\langle a, 0\rangle)$ and $\Phi_e(\langle a', 0\rangle)$ lie in the same ray X_i^k. By observations (1) and (2) above, Φ_e cannot define a sequence of edge-disjoint rays. If we never define $s_2(e)$, the same argument shows that Φ_e cannot define a sequence of edge-disjoint rays. If we eventually define both $s_1(e)$ and $s_2(e)$, then our action at that stage guarantees (by observation (2) above) that Φ_e does not define a sequence of edge-disjoint rays.

To show that $\mathsf{WIRT}_{\mathrm{UYD}}$ is not provable in RCA_0, modify the above proof by growing and merging each path P_i^k at both of its endpoints. $\square$

9. Open questions

In addition to the variations of the Halin type theorems investigated here that remain open problems of graph theory ($\mathsf{IRT}_{\mathrm{DVD}}$ and $\mathsf{IRT}_{\mathrm{DED}}$) the most intriguing computational and reverse mathematical questions are about either separating the variants or providing additional reductions or equivalences among the $\mathsf{IRT}_{\mathrm{XYZ}}$.

Question 9.1. Can any additional arrows be added to Figure 1 over RCA_0 or $\mathsf{RCA}_0 + I\Sigma^1_1$? (This includes the question of whether $\mathsf{RCA}_0 \vdash \mathsf{IRT}_{\mathrm{UVD}} \rightarrow \mathsf{IRT}_{\mathrm{UVS}}$.)

As we noted in Remark 5.8 there is an apparent additional reduction in Bowler, Carmesin, Pott [3, pg. 2 l. 3–7]. They use an intermediate reduction to locally finite graphs in the sense of relying on the fact that if a graph has arbitrarily many disjoint rays it has a locally finite subgraph with arbitrarily many disjoint rays. This is the principle to which that Remark refers. It plus ACA_0 is a THA but over RCA_0 it does not imply ACA_0 and is provably very weak (in the sense of being highly conservative over RCA_0). Shore [26] analyzes this and many similar principles some related to the $\mathsf{IRT}_{\mathrm{XYZ}}$ and others to an array of classical logical principles.

Any reductions in RCA_0 as requested in the Question above would, of course, provide the analogous ones for the $\mathsf{IRT}^*_{\mathrm{XYZ}}$. However, it is possible that other implications can be proven for the $\mathsf{IRT}^*_{\mathrm{XYZ}}$:

Question 9.2. Can any implications of the form $\mathsf{IRT}^*_{\mathrm{XYZ}} \rightarrow \mathsf{IRT}^*_{\mathrm{X'Y'Z'}}$ be proven in RCA_0 other than the ones known to hold for the IRT versions?

Probably more challenging is the problem of separating the principles.

Question 9.3. Can one prove any nonimplication over RCA_0 or over $\mathsf{RCA}_0 + I\Sigma^1_1$ for any pair of the $\mathsf{IRT}_{\mathrm{XYZ}}$?

Of course a separation by standard models or even ones over $I\Sigma^1_1$ for the $\mathsf{IRT}_{\mathrm{XYZ}}$ would give nonimplication for the corresponding $\mathsf{IRT}^*_{\mathrm{XYZ}}$ but it might be that nonstandard models could be used to separate one pair of versions but not the other.

The most natural separation questions involve Σ^1_1-AC_0.

Question 9.4. Can one show that any of the $\mathsf{IRT}_{\mathrm{XYZ}}$ which are provable in Σ^1_1-AC_0 ($\mathsf{IRT}_{\mathrm{XYS}}$ and $\mathsf{IRT}_{\mathrm{UVD}}$) do not imply Σ^1_1-AC_0 over RCA_0 or even over $\mathsf{RCA}_0 + I\Sigma^1_1$? An intermediate result might be that $\mathsf{IRT}^*_{\mathrm{XYZ}}$ (for one of these versions) does not imply Σ^1_1-AC_0 over RCA_0.

The next natural question looks below ABW_0 in Figure 4.

Question 9.5. Can one prove that finite-Σ^1_1-AC_0 does not imply ABW_0 over RCA_0 or $\mathsf{RCA}_0 + I\Sigma^1_1$?

The weaker versions, $\mathsf{WIRT}_{\mathrm{XYZ}}$, of the $\mathsf{IRT}_{\mathrm{XYZ}}$, prompt a question about ACA_0.

Question 9.6. Do any of the $\mathsf{WIRT}_{\mathrm{XYZ}}$ (especially the ones provable from ACA_0) imply ACA_0? An easier question might be whether they imply WKL_0?

References

1. Ron Aharoni, *König's duality theorem for infinite bipartite graphs*, J. London Math. Soc. (2) **29** (1984), no. 1, 1–12. MR 734985

2. Ron Aharoni, Menachem Magidor, and Richard A. Shore, *On the strength of König's duality theorem for infinite bipartite graphs*, J. Combin. Theory Ser. B **54** (1992), no. 2, 257–290. MR 1152453
3. Nathan Bowler, Johannes Carmesin, and Julian Pott, *Edge-disjoint double rays in infinite graphs: a Halin type result*, J. Combin. Theory Ser. B **111** (2015), 1–16. MR 3315597
4. Chris J. Conidis, *Comparing theorems of hyperarithmetic analysis with the arithmetic Bolzano-Weierstrass theorem*, Trans. Amer. Math. Soc. **364** (2012), no. 9, 4465–4494. MR 2922598
5. Reinhard Diestel, *Graph theory*, fifth ed., Graduate Texts in Mathematics, vol. 173, Springer, Berlin, 2017. MR 3644391
6. Yu. L. Ershov, S. S. Goncharov, A. Nerode, J. B. Remmel, and V. W. Marek (eds.), *Handbook of recursive mathematics. Vol. 2*, Studies in Logic and the Foundations of Mathematics, vol. 139, North-Holland, Amsterdam, 1998, Recursive algebra, analysis and combinatorics. MR 1673582
7. Harvey Friedman, *Algorithmic procedures, generalized Turing algorithms, and elementary recursion theory*, Logic Colloquium '69 (Proc. Summer School and Colloq., Manchester, 1969), 1971, pp. 361–389. MR 0304140
8. ———, *Some systems of second order arithmetic and their use*, Proceedings of the International Congress of Mathematicians (Vancouver, B. C., 1974), Vol. 1 (1975), 235–242. MR 0429508
9. Harvey M. Friedman, *Subsystems of set theory and analysis*, 1967, Thesis (Ph. D.)–Massachusetts Institute of Technology, Dept. of Mathematics.
10. Jun Le Goh, *The strength of an axiom of finite choice for branches in trees*, in preparation.
11. Rudolf Halin, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, Math. Nachr. **30** (1965), 63–85. MR 0190031
12. ———, *Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen*, Mathematische Nachrichten **44** (1970), no. 1-6, 119–127.
13. Denis R. Hirschfeldt, *Slicing the truth*, Lecture Notes Series. Institute for Mathematical Sciences. National University of Singapore, vol. 28, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2015, On the computable and reverse mathematics of combinatorial principles, Edited and with a foreword by Chitat Chong, Qi Feng, Theodore A. Slaman, W. Hugh Woodin and Yue Yang. MR 3244278
14. Pierre Jullien, *Contribution à l'étude des types d'ordres dispersés*, 1968, Thesis (Ph.D.)–Marseille.
15. Georg Kreisel, *The axiom of choice and the class of hyperarithmetic functions*, Nederl. Akad. Wetensch. Proc. Ser. A 65 = Indag. Math. **24** (1962), 307–319. MR 0140418
16. Antonio Montalbán, *Indecomposable linear orderings and hyperarithmetic analysis*, J. Math. Log. **6** (2006), no. 1, 89–120. MR 2250955
17. ———, *On the Π^1_1-separation principle*, MLQ Math. Log. Q. **54** (2008), no. 6, 563–578. MR 2472464
18. ———, *Open questions in reverse mathematics*, Bull. Symbolic Logic **17** (2011), no. 3, 431–454. MR 2856080
19. ———, *Fraïssé's conjecture in Π^1_1-comprehension*, J. Math. Log. **17** (2017), no. 2, 1750006, 12. MR 3730562
20. Itay Neeman, *The strength of Jullien's indecomposability theorem*, J. Math. Log. **8** (2008), no. 1, 93–119. MR 2674002
21. ———, *Necessary use of Σ^1_1 induction in a reversal*, J. Symbolic Logic **76** (2011), no. 2, 561–574. MR 2830416
22. Klaus-Peter Podewski and Karsten Steffens, *Injective choice functions for countable families*, J. Combinatorial Theory Ser. B **21** (1976), no. 1, 40–46. MR 0409195
23. Hartley Rogers, Jr., *Theory of recursive functions and effective computability*, second ed., MIT Press, Cambridge, MA, 1987. MR 886890
24. Joseph G. Rosenstein, *Linear orderings*, Pure and Applied Mathematics, vol. 98, Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1982. MR 662564
25. Gerald E. Sacks, *Higher recursion theory*, Perspectives in Mathematical Logic, Springer-Verlag, Berlin, 1990. MR 1080970
26. Richard A. Shore, *Almost theorems of hyperarithmetic analysis*, to appear.
27. ———, *On the strength of Fraïssé's conjecture*, Logical methods (Ithaca, NY, 1992), Progr. Comput. Sci. Appl. Logic, vol. 12, Birkhäuser Boston, Boston, MA, 1993, pp. 782–813. MR 1281172
28. Stephen G. Simpson, *On the strength of König's duality theorem for countable bipartite graphs*, J. Symbolic Logic **59** (1994), no. 1, 113–123. MR 1264968

29. ———, *Subsystems of second order arithmetic*, second ed., Perspectives in Logic, Cambridge University Press, Cambridge; Association for Symbolic Logic, Poughkeepsie, NY, 2009. MR 2517689
30. John R. Steel, *Forcing with tagged trees*, Ann. Math. Logic **15** (1978), no. 1, 55–74. MR 511943
31. Robert Alan Van Wesep, *Subsystems of second-order arithmetic, and descriptive set theory under the axiom of determinateness*, ProQuest LLC, Ann Arbor, MI, 1977, Thesis (Ph.D.)–University of California, Berkeley. MR 2627617

Department of Mathematics, Yale University, New Haven CT 06520

Department of Mathematics, University of Wisconsin-Madison, Madison WI 53706
Email address: junle.goh@wisc.edu

Department of Mathematics, Cornell University, Ithaca NY 14853