

Provably Total Functions in the Polynomial Hierarchy

Noah Fleming
Memorial University

Deniz Imrek
UT Austin

Christophe Marciot
Memorial University

April 22, 2025

Abstract

TFNP studies the complexity of total, verifiable search problems, and represents the first layer of the total function polynomial hierarchy (TFPH). Recently, problems in higher levels of the TFPH have gained significant attention, partly due to their close connection to circuit lower bounds. However, very little is known about the relationships between problems in levels of the hierarchy beyond TFNP.

Connections to proof complexity have had an outsized impact on our understanding of the relationships between subclasses of TFNP in the black-box model. Subclasses are characterized by provability in certain proof systems, which has allowed for tools from proof complexity to be applied in order to separate TFNP problems. In this work we begin a systematic study of the relationship between subclasses of total search problems in the polynomial hierarchy and proof systems. We show that, akin to TFNP, reductions to a problem in $\text{TF}\Sigma_d$ are equivalent to proofs of the formulae expressing the totality of the problems in some Σ_d -proof system. Having established this general correspondence, we examine important subclasses of TFPH. We show that reductions to the **STRONGAVOID** problem are equivalent to proofs in a Σ_2 -variant of the (unary) Sherali-Adams proof system. As well, we explore the TFPH classes which result from well-studied proof systems, introducing a number of new $\text{TF}\Sigma_2$ classes which characterize variants of DNF resolution, as well as $\text{TF}\Sigma_d$ classes capturing levels of Σ_d -bounded-depth Frege.

Contents

1	Introduction	1
2	Preliminaries on the Total Function Polynomial Hierarchy	3
3	Proof Systems for the Total Function Polynomial Hierarchy	4
4	Sherali-Adams and Strong Range Avoidance	7
5	A Generic Correspondence	13
6	Characterizations in $\text{TF}\Sigma_2$	16
7	$\text{TF}\Sigma_{d+1}$ Consequences of Depth- $d.5$ Frege	25
8	Acknowledgments	33
	References	34

1 Introduction

The class TFNP consists of the *total* search problems whose solutions are verifiable in polynomial time. It has received considerable attention since it captures fundamental problems in a broad range of areas, whose lack of efficient algorithms is not readily explained by the theory of NP-completeness. Famous examples include NASH: output a Nash equilibria of a given bimatrix game; and FACTORING: output a prime divisor of given integer. TFNP itself is not believed to admit complete problems [Pud15], and as a consequence much of the work on TFNP has focused on studying subclasses which do. However, we are limited to proving conditional or oracle separations, as a separation between any TFNP subclasses would imply $P \neq NP$.

A flurry of recent results have established a *complete* picture of the relationships between the major TFNP subclasses in the *black-box* setting, where the input is presented as a black-box oracle which can be queried [BCE⁺98, GHJ⁺22b, GHJ⁺22a, FGPR24, GKRS19, FGPR24]. These results exploited a deep connection between black-box TFNP — denoted $TFNP^{dt}$ — and proof complexity, an area which studies efficient provability in certain propositional logics, known as proof systems. The connection of proof complexity to $TFNP^{dt}$ can be summarized as follows: A reduction between two total search problems is a proof that the first is total, assuming the totality of the second. By employing this lens, it has been shown that many important $TFNP^{dt}$ subclasses are *characterized* by provability in certain well-studied proof systems in the sense that there is a simple proof of the totality of a search problem if and only if there is an efficient reduction of that search problem to the complete problem for that subclass [GKRS19, GHJ⁺22b, BFI23, LPR24, DR23]. This connection has been highly impactful for the study of $TFNP^{dt}$, allowing for the rich set of tools in proof complexity to be leverage in order to provide separations between the major $TFNP^{dt}$ subclasses.

$TFNP = TF\Sigma_1$ is the first level of the *total function polynomial hierarchy* $TFPH = \bigcup_i TF\Sigma_i$ [KKMP21]. Recently, problems in higher levels of the polynomial hierarchy have received considerable attention, in part due to their close connection to circuit lower bounds. Indeed, consider the task of finding (the truth table of) a function which does not have circuits of size s . Using a standard encoding, any circuit of size s can be represented uniquely by $k = \text{poly}(s)$ -many bits. Consider the map $T : \{0, 1\}^k \rightarrow \{0, 1\}^n$ which maps circuits of size s to truth tables of the function that they compute. Finding a truth table of a function with high circuit complexity is equivalent to finding a string which is not in the range of T . This is an instance of the RANGEAVOIDANCE problem.

Definition 1.1. RANGEAVOIDANCE (or simply AVOID) is the following search problem: given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{n+1}$, find a $y \in \{0, 1\}^{n+1}$ such that for all x , $f(x) \neq y$.

Observe that any solution y to AVOID can be checked by a coNP verifier — check that for every $x \in \{0, 1\}^n$, $f(x) \neq y$. This means that AVOID belongs to the class $TF\Sigma_2$. If there is an algorithm for solving AVOID which belongs to a class \mathcal{C} then this implies the existence of a function in \mathcal{C} which does not have small circuits — a circuit lower bound against \mathcal{C} ! This approach led to the recent breakthrough circuit lower bounds against symmetric exponential time [Li24, CHR24, KP24]. Hence, understanding the complexity of $TF\Sigma_2$ is important for understanding circuit lower bounds. Indeed, the current best upper bound puts AVOID in the class of problems reducible to LOP — the $TF\Sigma_2$ problem of finding a minimum element in a total order.

$TF\Sigma_2$ contains numerous important problems beyond those connected to circuit lower bounds. For example, AVOID is the complete problem for the class APEPP which also captures the complexity of finding pseudo-random number generators, randomness extractors, and rigid matrices [Kor21]. We can restrict AVOID to only have one more element in its range than in its domain to obtain the problem STRONGAVOID.

Definition 1.2. STRONGRANGEAVOIDANCE (or simply STRONGAVOID) is the following search problem: given a function $f : \{0, 1\}^n \setminus \{0\} \rightarrow \{0, 1\}^n$, find an empty hole $y \in \{0, 1\}^n$, i.e., such that for all $x \in \{0, 1\}^n \setminus \{0\}$, $f(x) \neq y$.

STRONGAVOID is the complete problem for the class PEPP which captures the complexity of finding objects whose existence is guaranteed by the union bound, including all of APEPP [KKMP21]. Important problems have also been identified in higher levels of the polynomial hierarchy, such as those related to finding sets of large VC dimension [KKMP21].

Despite the importance of problems in levels of the polynomial hierarchy beyond TFNP, there has been little structural exploration into how they relate. Indeed, [KP24] provide the first black-box separation, showing that STRONGAVOID is not reducible to any problem in $TF\Sigma_2$ with a unique solution (in fact, they show that it cannot be solved with non-adaptive oracle calls to any language in Σ_2^P). Proof complexity has had an outsized impact on

proving black-box separations for TFNP. To facilitate further structural exploration of TFPH, we would like to explore to what degree proof complexity tools can be used to provide separations between classes within higher levels of the black-box total function polynomial hierarchy (denoted TFPH^{dt}).

Our Results

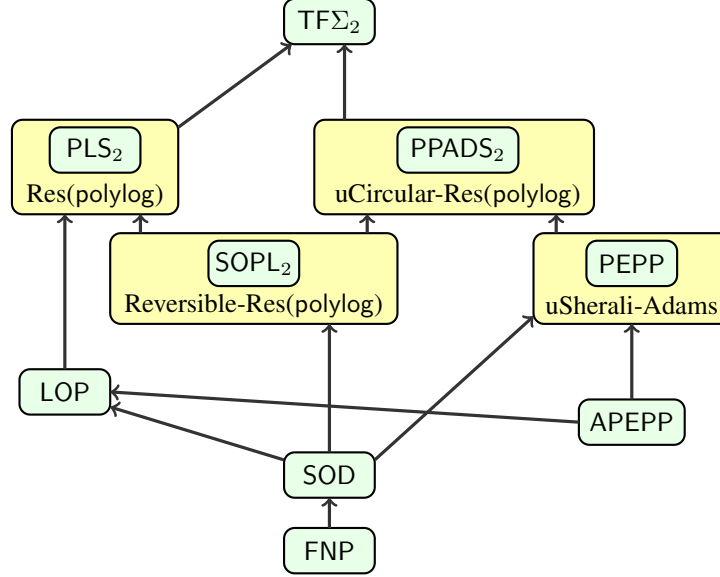


Figure 1: Relationships and characterizations of $\text{TF}\Sigma_2$ classes studied. An arrow indicates containment.

In this paper we begin a systematic study of the connections between the total function polynomial hierarchy in the black-box model and propositional proof complexity. First, we identify the form that proof systems which characterize $\text{TF}\Sigma_d^{dt}$ subclass take. In order to characterize $\text{TF}\Sigma_d$ subclasses these proof systems must be able to prove the validity of depth- $(d + 1)$ propositional formulas. However, they cannot be Cook-Reckhow proof systems (NP-verifiers) in general unless $\text{NP} = \text{coNP}$ as there are syntactic subclasses of $\text{TF}\Sigma_2^{dt}$ which contain all of TFNP^{dt} ; a characterization of which would imply a polynomially-bounded proof system. We show that in order to characterize $\text{TF}\Sigma_d^{dt}$ subclasses it suffices to augment Cook-Reckhow proof systems P with a Σ_d -weakening rule which generalizes the resolution weakening rule to Σ_d formulas; we call the resulting proof system Σ_d - P (see Definition 3.5).

To begin, we explore the limits of these characterizations, verifying that this is indeed the correct definition of a proof system for $\text{TF}\Sigma_d^{dt}$.

Theorem 1.3 (Informal). *The following hold:*

1. *For every syntactic $\mathcal{C} \subseteq \text{TF}\Sigma_d^{dt}$ there is a Σ_d -proof system P such that $R \in \mathcal{C}$ if and only if P efficiently proves that R is total.*
2. *For every well-behaved Σ_d -proof system P there is a syntactic $\text{TF}\Sigma_d^{dt}$ subclass \mathcal{C} such that $R \in \mathcal{C}$ if and only if P proves that R is total.*

Having established this scaffolding result, we begin to explore characterizations of specific $\text{TF}\Sigma_d^{dt}$ subclasses; our results for $\text{TF}\Sigma_2$ can be seen in Figure 1. First, we show that PEPP^{dt} is characterized by the Σ_2 -variant of the Sherali-Adams proof system.

Theorem 1.4 (Informal). *$R \in \text{PEPP}^{dt}$ iff there is an efficient Σ_2 -Sherali-Adams proof that R is total.*

This allows one to use an extension to the pseudo-expectation technique in order to exclude total search problems from PEPP, and hence also APEPP. Currently, no such exclusions are known.

We also consider several variants of the *DNF-resolution* proof system: DNF Resolution ($\text{Res}(\text{polylog})$), Circular DNF resolution ($\text{uCirRes}(\text{polylog})$), and Reversible DNF resolution ($\text{RevRes}(\text{polylog})$). We introduce new $\text{TF}\Sigma_2^{dt}$ classes which characterize them.

Theorem 1.5 (Informal). $\Sigma_2\text{-Res}(\text{polylog})$, $\Sigma_2\text{-uCirRes}(\text{polylog})$, $\Sigma_2\text{-RevRes}(\text{polylog})$ are characterized by the $\text{TF}\Sigma_2^{dt}$ subclasses PLS_2 , SOL_2 , SOPL_2 , respectively.

We explore how these new classes relate to natural $\text{TF}\Sigma_2$ classes, which can be seen in Figure 1. In doing so, we introduce a natural $\text{TF}\Sigma_2$ class SOD, of problems reducible to finding a source in a DAG given a sink, which we believe may be of independent interest.

Finally, we show that our characterization of DNF resolution can be extended to characterize bounded-depth Frege in higher levels of TFPH. The depth- d Frege system allows one to cut on depth- d propositional formulas; that is, with d -many quantifier alternations.

Theorem 1.6 (informal). $\Sigma_{d+2}\text{-Depth } d.5\text{-Frege}$ is characterized by the $\text{TF}\Sigma_{d+2}^{dt}$ class PLS_{d+2}^{dt} .

This result is inspired by the work of Beckmann and Buss who characterize PE_d and GI_d in bounded arithmetic [PT12]. It is also the $\text{TF}\Sigma_d$ analogue of Thapen’s recent TFNP characterization of depth- d Frege [Tha24].

Comparison with Bounded Arithmetic. Characterizations of TFPH classes have been studied in the *uniform* setting by theories of bounded arithmetic. Beckmann and Buss [BB09b] showed that Σ_k^b -definable functions of T_2^d correspond to the class $\text{PLS}_{d-1}^{\Sigma_k^b}$, which is defined by replacing the polynomial-time predicates and functions of the complete problem for the TFNP subclass PLS with predicates and functions from $\text{P}^{\Sigma_k^b}$. This results in the *generalized polynomial local search problem* GPLS_d of [PT12]. However, these correspondences do not stray outside of proof systems which correspond to bounded-depth Frege systems.

Open Problems. In this paper we provide the framework for characterizations between total search problems in the polynomial hierarchy, leaving open many natural questions.

1. We study decision-tree reductions, as these are the query analogue of polynomial-time reductions. However, it is natural also to consider more powerful reductions, such as P^{NP} -reductions. What characterizations does one obtain under such reductions?
2. There are several studied classes for which we do not yet have characterizations, such as APEPP and LOP. Due to the connection between STRONGAVOID and Sherali-Adams, it would appear that PEPP should correspond to a variant of Sherali-Adams which produces a large negative value, rather than -1 . However we cannot maintain this under decision-tree reductions.
3. $\text{TF}\Sigma_2$ problems with unique solutions play a critical role in the recent circuit lower bounds [Li24, CHR24, KP24]. What properties do proof systems which characterize such problems possess?

2 Preliminaries on the Total Function Polynomial Hierarchy

Subclasses of TFPH are typically defined by a simple existence principle to which everything in the class reduces. For example, any total order must have a minimal element. These existence principles naturally give rise to total search problems. Continuing the example:

Definition 2.1. The *Linear Ordering Principle* (LOP) asks, given $\prec: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, to find:

- A minimal element: $x \in \{0, 1\}^n$ such that $\forall y \neq x, x \prec y$.
- A violation to the total order: either (i) $x \in \{0, 1\}^n$ such that $x \prec x$, (ii) $x \neq y$ such that $x \not\prec y$ and $y \not\prec x$, or (iii) $x \prec y$ and $y \prec z$ and $x \not\prec z$.

To make these problems non-trivial the input is presented as a circuit C so that the search space is exponential in the number of input bits n . Formally, for any $x, y \in \{0, 1\}^n$, $C(x, y) = \prec(x, y)$. We call C a *white-box* encoding of the problem. Unfortunately, a separation between any pair of total search problems in the white-box model is hard to achieve, as it would imply $\text{P} \neq \text{NP}$.

Instead, we can gain intuition for the relationships between these classes by exploring their *black-box* variants. In this setting C is given as an oracle which can be queried, but we no longer have access to the description of C . A major benefit of considering the black-box model is that we can now prove *unconditional* separations between classes without having to resolve P versus NP. These separations imply oracle separation in the white-box setting.

A *query search problem* is a sequence of relations $R_n \subseteq \{0, 1\}^n \times \mathcal{O}_n$, one for each $n \in \mathbb{N}$. It is *total* if for every $x \in \{0, 1\}^n$ there is an $o \in \mathcal{O}_n$ such that $(x, o) \in R_n$. We think of $x \in \{0, 1\}^n$ as a bit string which can be accessed by querying individual bits, and we will measure the complexity of solving R_n as the number of bits that must be queried. Hence, an efficient algorithm for R_n will be one which finds a suitable o while making at most $\text{polylog}(n)$ -many queries to the input. We will not charge the algorithm for other computational steps, and therefore an efficient algorithm corresponds to a shallow decision tree. Total query search problems which can be computed by decision trees of depth $\text{polylog}(n)$ belong to the class FP^{dt} , where dt indicates that it is a black-box class. While search problems are formally defined as sequences $R = (R_n)$, we will often want to speak about individual elements in the sequence. For readability, we will abuse notation and refer to elements R_n in the sequence as total search problems; furthermore, we will often drop the subscript n , and rely on context to differentiate.

In this paper we will be considering total query search problems in the polynomial hierarchy TFPH^{dt} .

Definition 2.2. We say that a total search problem $R = (R_n)$, where $R_n \subseteq \{0, 1\}^n \times \mathcal{O}_n$, belongs to the d^{th} level of the query *total function polynomial hierarchy* $\text{TF}\Sigma_d^{dt}$ if for every $o \in \mathcal{O}_n$

$$(x, o) \in R \iff \forall z_1 \in \{0, 1\}^{\ell_1} \exists z_2 \in \{0, 1\}^{\ell_2} \dots Q z_{d-1} \in \{0, 1\}^{\ell_{d-1}} V_{o, (z_1, \dots, z_{d-1})}(x) = 1,$$

where $Q \in \{\exists, \forall\}$, $V_{o, z} = V_{o, (z_1, \dots, z_{d-1})}$ is a decision tree of $\text{polylog}(n)$ -depth, and each $\ell_i \in \text{polylog}(n)$.

Note that $\text{FP}^{dt} = \text{TF}\Sigma_0^{dt}$ and $\text{TFNP} = \text{TF}\Sigma_1^{dt}$. At this point one may ask about $\text{TF}\Pi_d^{dt}$. Kleinberg et al. [KKMP21] showed that $\text{TF}\Pi_d$ is efficiently reducible to $\text{TF}\Sigma_{d-1}$, and vice-versa. Hence, it does not offer a new perspective.

We can compare the complexity of total search problems by taking reductions between them. The following defines decision tree reductions between total search problems, the query analogue of polynomial-time reductions.

Definition 2.3. For total search problems $R \subseteq \{0, 1\}^n \times \mathcal{O}_n$, $S \subseteq \{0, 1\}^m \times \mathcal{O}'_m$, there is an *S-formulation* of R if, for every $i \in [m]$ and $o \in \mathcal{O}_S$, there are functions $f_i : \{0, 1\}^n \rightarrow \{0, 1\}$ and $g_o : \{0, 1\}^n \rightarrow \mathcal{O}_n$ such that

$$(f(x), o) \in S \implies (x, g_o(x)) \in R,$$

where $f(x) = (f_1(x) \dots f_m(x))$. The *depth* of the reduction is

$$d := \max(\{\text{depth}(f_i) : i \in [m]\} \cup \{\text{depth}(g_o) : o \in \mathcal{O}'_m\}),$$

where $\text{depth}(f)$ denotes the minimum depth of any decision tree which computes f . The *size* of the reduction is m , the number of input bits to S . The *complexity* of the reduction is $\log m + d$, and the complexity of reducing R to S is the minimum S -formulation of R .

We extend this definition to sequences in the natural way. If $S = (S_n)$ is a sequence and R_n is a single search problem then the complexity of reducing R_n to S is the minimum over m of the complexity of reducing R_n to S_m . For two sequences of search problems $s = (S_n)$ and $R = (R_n)$, the complexity of reducing S to R is the complexity of reducing R_n to S for each n . A reduction from R to S is efficient if its complexity is $\text{polylog}(n)$; we denote this by $R \leq_{dt} S$.

We say that a class of total search problems $\mathcal{C} \subseteq \text{TF}\Sigma_d^{dt}$ has $R \in \mathcal{C}$ as its *complete problem* if for every $S \in \mathcal{C}$, $S \leq_{dt} R$. We call subclasses with complete problems *syntactic*.

3 Proof Systems for the Total Function Polynomial Hierarchy

Search problems in the black-box model are intimately tied to the complexity of propositional theorem proving. A proof is a procedure for convincing a verifier that a statement is correct. In the propositional setting, a proof convinces the verifier that a propositional formula is unsatisfiable (equivalently, its negation is a tautology).

3.1 Recap: Proof Systems and TFNP

We begin by recalling how characterizations of proof systems by TFNP^{dt} subclasses occur. We will then generalize this to TFPH. Let UNSAT be the language of all unsatisfiable propositional formulas.

Definition 3.1. A *Cook-Reckhow proof system* is a polynomial-time function $P : \{0, 1\}^* \rightarrow \{0, 1\}$ such that for every propositional formula $F \in \{0, 1\}^*$,

$$F \in \text{UNSAT} \iff \exists \Pi \in \{0, 1\}^*, P(\Pi, F) = 1.$$

The size of proving an unsatisfiable formula F in P is $\min\{|\Pi| : P(\Pi, F) = 1\}$.

For many proof systems there is an associated width/degree measure. For example, in resolution it is the maximum number of literals in any clause appearing in a proof, and in algebraic systems such as Sherali-Adams and Sum-of-Squares it is the maximum degree of the polynomials appearing in the proof. Characterizations of TFNP^{dt} subclasses are in terms of a *complexity* parameter of the proof system, denoted

$$P(F) := \min \{ \text{width}(\Pi) + \log \text{size}(\Pi) : \Pi \text{ is a } P\text{-proof of } F \},$$

where width is the associated width measure for that system.

Typically one studies the complexity of proving the unsatisfiability of *CNF* formulas. As a CNF formula $F = C_1 \wedge \dots \wedge C_m$ is falsified only when one of its clauses is falsified, a proof convinces the verifier that for every assignment $x \in \{0, 1\}^n$ there is some clause C_i of F such that $C_i(x) = 0$. Hence, the complexity of proving that F is unsatisfiable is intimately related to the complexity of exhibiting a falsified clause, given an assignment. This is known as the *false clause search problem* $\text{SEARCH}_F \subseteq \{0, 1\}^n \times [m]$, defined as

$$(x, i) \in \text{SEARCH}_F \iff C_i(x) = 0.$$

As F is unsatisfiable this search problem is total, and if each clause of F contains at most $\text{polylog}(n)$ -many variables it belongs to TFNP^{dt} .

The above intuition suggests that understanding TFNP^{dt} (or at least the false clause search problem) is important for understanding proof complexity. Remarkably, proof complexity is also crucial for understanding TFNP^{dt} . It turns out that TFNP^{dt} is *equivalent* to a large sub-area of proof complexity! The intuition is the following: A reduction between two total search problems is a *proof* that the first is total, assuming the totality of the second. By employing this lens, works have shown that many common proof systems are *characterized* by certain well-studied tautologies in the sense that they can prove a tautology iff there is a short reduction of that tautology to the characterizing one.

The heart of this connection is the following claim which shows that TFNP^{dt} is *exactly* the study of the false clause search problem. The proof proceeds by expressing the totality of any problem R in TFNP^{dt} as a tautology and then taking its negation.

Claim 3.2. If $R \in \text{TFNP}^{dt}$ then there is an unsatisfiable $\text{polylog}(n)$ -width CNF formula F_R such that $\text{SEARCH}_{F_R} \in \text{TFNP}^{dt}$ and $R =_{dt} \text{SEARCH}_{F_R}$.

From this, characterizations of TFNP^{dt} subclasses by proof systems have been derived. We say that a syntactic subclass $\mathcal{C} \subseteq \text{TFNP}^{dt}$ is *characterized* by a proof system P if for every $\text{SEARCH}_F \in \text{TFNP}^{dt}$, $\text{SEARCH}_F \in \mathcal{C}$ iff $P(F) = \text{polylog}(n)$.

3.2 Proof Systems and TFPH

The aim of this paper is to explore characterizations of classes of problems belonging to higher levels of TFPH^{dt}. These will correspond to the provability of quantified formulas.

Definition 3.3. A $\Sigma_{d.5}$ formula F is the propositional translation of any quantified formula of the form

$$\exists z_1 \in \{0, 1\}^{\ell_1} \forall z_2 \in \{0, 1\}^{\ell_2} \dots Q z_d \in \{0, 1\}^{\ell_d} L(x, z_1, \dots, z_d),$$

where $\ell_i \in \text{polylog}(n)$, $Q \in \{\exists, \forall\}$, and L is a formula which depends on at most $\text{polylog}(n)$ -many free variables (x) . That is, a $\Sigma_{d.5}$ formula is of the form

$$F = \bigvee_{z_1 \in \{0, 1\}^{\ell_1}} \bigwedge_{z_2 \in \{0, 1\}^{\ell_2}} \dots \bigcirc_{z_d \in \{0, 1\}^{\ell_d}} L_{z_1, \dots, z_d}(x),$$

where $\bigcirc \in \{\wedge, \vee\}$, and $L_{z_1, \dots, z_d}(x) := L(x, z_1, \dots, z_d)$. Similarly, $\Pi_{d.5}$ formulas are negations of $\Sigma_{d.5}$ formulas.

Note that because L_z depends on $\text{polylog}(n)$ -many variables, we may assume without loss of generality (with a quasi-polynomial blow-up in size) that L_z is a CNF/DNF formula with clauses/terms of width $\text{polylog}(n)$. Hence, a $\Sigma_{d,5}$ -formula is a layered circuit of depth d where the gates at each layer are the same, and the gates at the first d layers are allowed $2^{\text{polylog}(n)}$ fanin, while the final layer is restricted to have $\text{polylog}(n)$ fanin. Observe that a $\Pi_{1,5}$ -formula is a low-width CNF formula.

Our aim is to characterize the higher levels of the total function polynomial hierarchy. Towards this, we generalize the false clause search problem to $\Sigma_{d,5}$ formulas.

False Formula Search. For a formula $F := \bigwedge_{o \in [m]} H_o$ where each H_o is a $\Sigma_{d,5}$ -formula, the *False Formula* search problem $\text{FF}_F \subseteq \{0, 1\}^n \times [m]$ is defined as

$$(x, o) \in \text{FF}_F \iff H_o(x) = 0.$$

Observe that if F is unsatisfiable then FF_F is total and $\text{FF}_F \in \text{TF}\Sigma_{d+1}^{dt}$. The following lemma generalizes Claim 3.2 to say that $\text{TF}\Sigma_d^{dt}$ is *exactly* the study of the false formula search problem.

Lemma 3.4. *For every $R \in \text{TF}\Sigma_d$ there is an unsatisfiable $\Pi_{d,5}$ -formula F_R such that $(x, o) \in R$ iff $(x, o) \in \text{FF}_{F_R}$.*

Proof. Let $R \subseteq \{0, 1\}^n \times [m] \in \text{TF}\Sigma_d$. Then there are $\text{polylog}(n)$ -depth decision trees $V_{o,(z_1, \dots, z_{d-1})}$ such that

$$(x, o) \in R \iff \forall z_1 \in \{0, 1\}^{\ell_1} \exists z_2 \in \{0, 1\}^{\ell_2} \dots Q z_{d-1} \in \{0, 1\}^{\ell_{d-1}} V_{o,(z_1, \dots, z_{d-1})}(x) = 1,$$

where $Q \in \{\exists, \forall\}$, $V_{o,\vec{z}} = V_{o,(z_1, \dots, z_{d-1})}$ is a decision tree of $\text{polylog}(n)$ -depth, and each $\ell_j \in \text{polylog}(n)$. Slightly abusing notation, let V_o be a propositional translation of the verifier as $\Sigma_{(d-1),5}$ -formula:

$$V_o(x) := \bigwedge_{z_1 \in \{0, 1\}^{\ell_1}} \bigvee_{z_2 \in \{0, 1\}^{\ell_2}} \dots \bigcirc_{z_{d-1} \in \{0, 1\}^{\ell_{d-1}}} V_{o,\vec{z}}(x),$$

where $\bigcirc \in \{\wedge, \vee\}$, and $V_{o,\vec{z}}(x)$ is computable by a $\text{polylog}(n)$ -depth decision tree, and hence propositionalized as a $\text{polylog}(n)$ -width CNF formula if $\bigcirc = \wedge$ or a $\text{polylog}(n)$ -width DNF if $\bigcirc = \vee$, collapsing the top gate into \bigcirc . This is done as follows: say that a root-to-leaf path in $V_{o,\vec{z}}$ is a b -path if it ends at a leaf labeled $b \in \{0, 1\}$. Then, $V_{o,\vec{z}}$ is propositionalized as

- If $d - 1$ is even: $\bigvee_{1\text{-path } p \in V_{o,\vec{z}}} p$,
- If $d - 1$ is odd: $\bigwedge_{0\text{-path } p \in V_{o,\vec{z}}} \neg p$,

where p is the conjunction of literals queried along p (if a variable x is queried and we take branch-0 then we consider this as literal $\neg x$ and otherwise as x). Note that in this case the outer gate of $V_{o,\vec{z}}$ matches \bigcirc , and the depth collapses by 1. Consider the following $\Pi_{d,5}$ -formula which states that R is not total:

$$F := \bigwedge_{o \in \mathcal{O}} \neg V_o(x).$$

Observe that if $(x, o) \in R$ then there is some z_1, \dots, z_d such that $V_{o,z}(x) = 1$, and hence $(x, o) \in \text{FF}_F$. Conversely, if $(x, o) \in \text{FF}_F$ then $(x, o) \in R$. \square

We will call the formula F_R the *propositionalization* of R . This lemma allows us to relate the complexity of total search problems to the provability of propositional formulas. In the remainder we will develop what provability means in this context. In particular, what are the properties of a proof system which proves the formulas that result from TFPH^{dt} search problems.

A characterization of a TFPH^{dt} class by a proof system proceeds by showing that the proof system can prove the correctness of reductions to the class. To discuss this we will need to propositionalize reductions.

Reduced Formula. Let $R \subseteq \{0, 1\}^n \times \mathcal{O}$ be a problem in $\text{TF}\Sigma_d^{dt}$ and let $V_{\vec{z},o}$, $o \in \mathcal{O}$ be its verifiers. Let (f, g) be an R -formulation where $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$, $g : \{0, 1\}^m \rightarrow \mathcal{O}$, then the *reduced formula* $F_R(f, g)$ is the $\Pi_{d,5}$ -formula defined as

$$F_R(f, g) := \bigwedge_{o \in \mathcal{O}} \bigwedge_{\text{path } p \in g_o} \neg V_{o,p}(f(x)),$$

where $V_{o,p}(f(x)) = \bigwedge_{z_1 \in \{0,1\}^{\ell_1}} \bigvee_{z_2 \in \{0,1\}^{\ell_2}} \cdots \bigcirc_{z_{d-1} \in \{0,1\}^{\ell_{d-1}}} (V_{o,\bar{z}}(f(x)) \wedge p)$ and $V_{o,\bar{z}}(f(x))$ can be represented as a polylog(n)-width CNF/DNF as in [Lemma 3.4](#), using that both $V_{o,\bar{z}}$ and f are computable by polylog(n)-depth decision trees.

Reduced formulas capture formulations in the following sense. Let $H := \bigwedge_{o \in \mathcal{O}_H} H_o$ and (f, g) be an FF_F -formulation of FF_H , where $F = \bigwedge_{o \in \mathcal{O}_F} F_o$. Then for any $o \in \mathcal{O}_F$ and any path p in g_o labelled with some $o^* \in \mathcal{O}_H$ we have that

$$V_{o,p}(f(x)) = 0 \implies H_{o^*}(x) = 0. \quad (1)$$

That is, $H_{o^*} \implies V_{o,p}(f)$, and we say that $V_{o,p}(f)$ is a *weakening* of H_{o^*} .

A proof system P is characterized by a TFPH^{dt} class \mathcal{C} with complete problem FF_F if efficient provability of F in that proof system implies low-complexity reductions to the complete problem FF_F for that class, and membership in the class \mathcal{C} implies that P can prove the correctness of the reduction to P . The latter takes the following form: if (f, g) is a FF_F -formulation of a $\text{FF}_H \in \mathcal{C}$ then

- i) From H , P can efficiently derive the reduced formula $F(f, g)$.
- ii) P has an efficient proof of $F(f, g)$.

What properties must a proof system possess in order to perform (i) and (ii) for a subclass $\mathcal{C} \subseteq \text{TFPH}^{dt}$? If $\text{TFNP}^{dt} \subseteq \mathcal{C}$ then a Cook-Reckhow proof system (an NP-verifier) does not suffice unless $\text{NP} = \text{coNP}$ ¹. Interestingly, what fails is step (i) — [Theorem 5.1](#) shows that step (ii) can always be carried out by a Cook-Reckhow system. We will need to augment Cook-Reckhow proof systems in order to perform step (i). The issue is that Cook-Reckhow systems cannot always perform the weakening from (1). That is, if $F(f, g) = \bigwedge_{o \in \mathcal{O}_{F(f,g)}} F_o$ and $H = \bigwedge_{o' \in \mathcal{O}_H} H_{o'}$ then by correctness of the reduction we know that for every $o \in \mathcal{O}_{F(f,g)}$, F_o is a weakening of some $H_{o'}$. However, Cook-Reckhow proof systems cannot necessarily derive F_o efficiently given H . For example, if $F_o = \top$, the trivial tautology, then this is tantamount to proving that F_o is a tautology, which is a coNP-complete task. It will suffice to augment our proof systems to be able to do so.

Definition 3.5. Let P be a Cook-Reckhow proof system. A proof of a Π_d formula $F = \bigwedge_{i \in [m]} F_i$ in the proof system Σ_d - P is a pair (H, Π) such that

- 1. Π is a P -proof that the Π_{d+1} -formula $H = \bigwedge_{j \in [k]} H_j$ is unsatisfiable.
- 2. Each H_j is a Σ_d -formula such that there is some $i \in [m]$ for which $F_i \implies H_j$. That is, H_j is a Σ_d -weakening of F_i .

The *complexity* of the proof (H, Π) is $\log |H| + \log s + d$ where $\log s + d$ is the complexity of the proof Π .

Clearly such proofs are verifiable in Σ_d . As we will see, they suffice to characterize subclasses of $\text{TF}\Sigma_d^{dt}$.

4 Sherali-Adams and Strong Range Avoidance

We begin with an example of a characterization by showing that $\text{STRONGRANGEAVOIDANCE}$ is characterized by Σ_2 -Sherali-Adams. A full treatment of this proof system is given in the monograph [\[FKP19\]](#).

For any boolean formula F we will assume without loss of generality that all negations occur at the leaves and let $\text{Vars}^+(F)$ be the positive literals in F and $\text{Vars}^-(F)$ be the negative literals. For any conjunct $t = \bigwedge_{x \in \text{Vars}^+(t)} x \wedge \bigwedge_{x \in \text{Vars}^-(t)} \neg x$ we associate the polynomial $\prod_{x \in \text{Vars}^+(t)} x \prod_{x \in \text{Vars}^-(t)} (1 - x)$, and refer to them also as conjuncts. A *conical junta* is a sum of conjuncts $\mathcal{J} := \sum t$.

Let $D = \bigvee_t t$ be any DNF. We can express D as a degree $\deg(D) := \max_{t \in D} \deg(t)$ polynomial

$$\sum_{t \in D} t - 1.$$

Observe that for any $x \in \{0, 1\}^n$, $D(x) = 1$ iff $\sum_{t \in D} t(x) - 1 \geq 0$. Henceforth we will abuse notation and refer to D as both the DNF and the associated polynomial.

Throughout this section we will work with *multi-linear arithmetic* associating $x_i^2 = x_i$ for every variable x . This has the effect of restricting the underlying linear program to $\{0, 1\}$ -points.

¹ Indeed, for any unsatisfiable 3-CNF formula F , $\text{FF}_F \in \text{TFNP}^{dt}$.

Definition 4.1. Let $F = \{D_i\}_{i \in [m]}$ be an unsatisfiable collection of DNFs. A Σ_2 -Unary DNF Sherali-Adams (which we denote by uSA) proof Π of F is a weakening $F' = \{D'_i\}_{i \in [m']}$ of F together with a list of canonical juntas $\mathcal{J}_i, \mathcal{J}$, such that

$$\sum_{i \in [m']} D'_i \mathcal{J}_i + \mathcal{J} = -1.$$

The *degree* $\deg(\Pi)$ is the maximum degree among $D_i, D'_i \mathcal{J}_i$, and \mathcal{J} , and the *size* $\text{size}(\Pi)$ is the number of monomials (counted with multiplicity) in $D_i, D'_i \mathcal{J}_i, \mathcal{J}$. The *complexity* of the proof is $\text{uSA}(\Pi) := \deg(\Pi) + \log \text{size}(\Pi)$, and the complexity of proving F is $\text{uSA}(F) := \min_{\Pi} \text{uSA}(\Pi)$, where the minimum is taken over all uSA proofs of F .

Note also that weakening subsumes the need to explicitly allow the additional conical junta in a uSA proof; we could instead defined uSA as a Nullstellensatz proof $\sum D'_i \mathcal{J}_i = -1$. This is because the additional junta \mathcal{J} may be introduced using weakening: for each conjunct t of \mathcal{J} , weaken some D_i in F to *true* or t . For example, D_i can be weakened to $x_i \vee \neg x_i \vee t$, the polynomial encoding of which is $x_i + (1 - x_i) + t - 1 = t$.

Claim 4.2. uSA is sound and complete.

Proof. Suppose that uSA is not sound, then there exists a uSA refutation of a satisfiable DNF $F = \{D_i\}_{i \in [m]}$,

$$\sum_{i \in [m']} D'_i \mathcal{J}_i + \mathcal{J} = -1.$$

Let $x \in \{0, 1\}^n$ be a satisfying assignment to F , meaning that for every i , $D'_i(x) = 1$ for any weakening D'_i of D_i , and in particular the polynomial representation of $D'_i(x) \geq 0$. As juntas are non-negative over $\{0, 1\}^n$ we have that

$$\sum_{i \in [m']} D'_i(x) \mathcal{J}_i(x) + \mathcal{J}(x) \geq 0,$$

which is a contradiction.

For completeness, let $F = \{D_i\}_{i \in [m]}$ be an unsatisfiable formula. Each assignment $x \in \{0, 1\}^n$ must falsify some DNF of F , which we will denote by D_x . Let I_x be the indicator polynomial $I_x := \prod_{i: x_i=1} x_i \prod_{i: x_i=0} (1 - x_i)$ of the assignment x . We claim that the polynomial

$$\sum_{x \in \{0, 1\}^n} I_x D_x = -1,$$

and is therefore a uSA proof. To see this, since we are working over the ideal $\langle x_i - x_i^2 \rangle$, it suffices to show that the polynomial evaluates to -1 on every $x \in \{0, 1\}^n$. Observe that if $y \in \{0, 1\}^n$ falsifies D_x then $D_x(y) = -1$, additionally, if $x \neq y$, then $I_x(y) = 0$. Hence, for every $y \in \{0, 1\}^n$,

$$\sum_{x \in \{0, 1\}^n} I_x(y) D_x(y) = I_y(y) D_y(y) = D_y(y) = -1.$$

□

In the rest of this section, we show that uSA is closely related to STRONGRANGEAVOIDANCE. We restate an equivalent definition next.

Definition 4.3. An instance of STRONGRANGEAVOIDANCE (STRONGAVOID) is given by a map $f : [n] \rightarrow [n + 1]$. A solution is any $h \in [n + 1]$ such that for every $p \in [n]$, $f(p) \neq h$.

STRONGAVOID can be encoded as a CNF formula by introducing, for every $p \in [n]$, $\log n + 1$ -many binary variables $p_1, \dots, p_{\log n + 1}$ naming in binary the hole $h \in [n + 1]$ to which pigeon p flies. For exposition, it will be convenient to think of p as an $(n + 1)$ -ary variable and we will denote by $\llbracket p = h \rrbracket$ the indicator conjunct that is satisfied iff p maps to $h \in [n + 1]$ under the given assignment

$$\llbracket p = h \rrbracket := p_1^{h_1} \wedge \dots \wedge p_{\log n + 1}^{h_{\log n + 1}},$$

where $p_i^{h_i} = p_i$ if the i^{th} bit of h is 1 and $\neg p_i$ otherwise. Note that $\sum_{h \in [n + 1]} \llbracket p = h \rrbracket = 1$ as polynomials.

We can express STRONGAVOID as the unsatisfiable family of DNFs,

$$\bigvee_{p \in [n]} \llbracket p = h \rrbracket \quad \forall h \in [n+1].$$

The main theorem of this section is the following.

Theorem 4.4. *For any $\text{FF}_F \in \text{TF}\Sigma_2^{dt}$, there is a complexity c STRONGAVOID-formulation of FF_F iff there is a Σ_2 -Sherali-Adams proof of complexity $\Theta(c)$.*

We break the proof of this theorem into [Lemma 4.5](#) and [Lemma 4.8](#) which are proven over the following two sub-sections. This theorem gives necessary and sufficient conditions for separating other $\text{TF}\Sigma_2$ classes \mathcal{C} from STRONGAVOID: exhibit a pseudo-expectation (see e.g., [\[FKP19\]](#)) against any $\text{polylog}(n)$ -width Σ_2 -weakening of the propositionalization of STRONGAVOID.

4.1 SA Proofs Imply sRA Reductions

Lemma 4.5. *If there is a size s and degree d uSA proof of F then there is a depth- d reduction from FF_F to an instance of STRONGAVOID of size $O(s)$.*

To prove this lemma, it will be convenient to work with the following problem which is equivalent to STRONGAVOID.

Definition 4.6. The *Unmetered Source of Dag* (USOD) problem is defined as follows. The input is a “successor” function $S : [n] \rightarrow [n]$ which defines a graph in which each vertex has fan-out ≤ 1 but arbitrary fan-in. There is an edge from i to j if $S(i) = j$. To make the problem total, we enforce that the vertex 1 is a sink, it will have fan-out 0 but fan-in at least 1. The goal is to find a source; the solutions are:

- 1 is a solution if either $S(1) \neq 1$ or $\forall v \neq 1 \in [n], S(v) \neq 1$ (1 is not a sink).
- $v \in [n]$ is a solution if $S(v) \neq v$ but $\forall u \in [n], S(u) \neq v$ (v is a source).

Lemma 4.7. $\text{USOD} =_{dt} \text{STRONGAVOID}$. *Furthermore, this reduction is by depth-1 decision trees.*

Proof. From an instance $S : [n] \rightarrow [n]$ of USOD, we construct an instance $f : [n] \rightarrow [n+1]$ of STRONGAVOID as follows. For $v \neq 1 \in [n]$, let $f(v) := S(v)$ and let $f(1) := n+1$. We claim that any solution u to this STRONGAVOID instance is a source in S . First observe that $u \neq n+1$ as $f(1) = n+1$. Hence, by construction, we have that $\forall v \in [n], S(v) \neq u$, and in particular $S(u) \neq u$, so u is a source.

For the converse direction, from an instance $f : [n] \rightarrow [n+1]$ of STRONGAVOID we construct an instance $S : [n+1] \rightarrow [n+1]$ of USOD by defining $S(v+1) := f(v)$ for all $v \in [n]$ and let $S(1) = 1$. Let v be a solution to this instance of USOD, if $v = 1$, then, since $S(1) = 1$, for all $u \in [n]$, $f(u) \neq 1$. Otherwise, $v \neq S(u)$ for all $u \in [n+1]$, and so $v \neq f(u)$ for all $u \in [n]$. \square

Proof of Lemma 4.5. Let $F = \bigwedge_{o \in \mathcal{O}} D_o$ and let Π be a size s and degree d uSA proof of F over n variables, where

$$\Pi := \sum_{i \in [m]} \sum_{j \in I_i} D'_i J_j + \sum_{k \in K} j_k + 1 = 0,$$

for sets of indices I_i, K , each D'_i is a weakening of some $D_o \in F$ and each J_j, j_k is a conjunct. We construct an instance of USOD with one node per occurrence of a (signed) monomial in Π . Therefore, for simplicity, we will refer to monomials as nodes and vice-versa. The constant 1 will be our distinguished sink, and we will set $S(1) = 1$. We will define the remaining successor pointers as follows:

Negative Monomials. Since $\Pi = 0$, there is a positive and negative copy of every monomial occurring in the proof; construct a pairing of the monomials in this way. Furthermore, under any assignment $x \in \{0, 1\}^n$ the number of monomials which evaluate to 1 and to -1 is equal. For each negative monomial $-m$ in Π , the decision tree $S(-m)$ queries the variables of m and outputs as follows:

- i) If $m(x) = 0$ then $S(-m) = -m$.
- ii) Otherwise, let m be the positive copy of $-m$ that $-m$ is paired with and set $S(-m) = m$.

This completes the description of the successor pointer for negative monomials.

Positive Monomials. For any positive monomial m , the decision tree for $S(m)$ first queries the (at most d -many) variables of m to determine the value of $m(x)$. If $m(x) = 0$, then $S(m) = m$. Otherwise, we will define S as follows.

We define the successor pointer for the positive monomials which belong to each $D'_i J_j$ first, and handle the monomials from the conjuncts j_d later. Fix some $D'_i J_j$ in Π , where $D'_i = \sum_{k \in [\ell]} t_d - 1$ and consider the monomials within it. We would like to satisfy the following property: there is a source within the monomials $D'_i J_j$ iff $D'_i(x) = -1$ (i.e., the DNF $D'_i(x) = 0$). To get some intuition, first suppose that $J_j = 1$ and that all monomials m in D'_i are positive — that is, $D'_i J_j = \sum_{k \in [\ell]} m_k - 1$. Then, the current assignment to S affects $D'_i J_j$ as follows:

- Each monomial m_k such that $m_k(x) = 0$ is an isolated vertex for which $S(m_k) = m_k$.
- Each monomial m_k for which $m_k(x) \neq 0$ has a single incoming edge (from $-m_k$).
- The monomial -1 has an outgoing edge.

If at least one of the monomials m_k is non-zero we can send it to -1 , and otherwise -1 becomes a source (see Figure 2). Therefore, the only sources will come from the “ -1 nodes” of falsified DNFs. To handle the general case, we use the fact that in every conjunct, under any assignment, there are at least as many non-zero positive monomials as non-zero negative monomials.

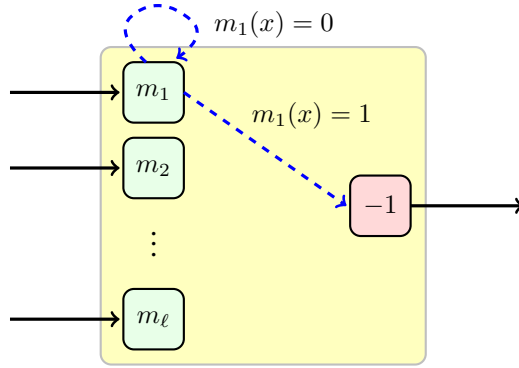


Figure 2: The “gadget” for a $D'_i J_j$ where $J_j = 1$ and D'_i contains only positive literals (each conjunct is a monomial).

We now describe the construction in general. Consider a $D'_i J_h$ in Π . For each positive monomial m in $D'_i J_j = (\sum_{k \in [\ell]} t_d - 1) J_j$, belonging to some conjunct $t_d J_j$, the pointer $S(m)$ will query the (at most d -many) variables in $t_d J_j$. Let $\alpha \in \{0, 1\}^{\text{Vars}(J_j)}$ be the assignment to the variables of J_j that was discovered.

If $J_j \upharpoonright \alpha = 0$: Then $D'_i J_j \upharpoonright \alpha = 0$. Hence, for every positive monomial m in $D'_i J_j$, either $m \upharpoonright \alpha = 0$, in which case we have already set $S(m) = m$, or m must cancel with another monomial $-m'$ in $D'_i J_j$ under α . That is, $m \upharpoonright \alpha = -m' \upharpoonright \alpha$, and so we define $S(m) = -m'$. Note that in this case there are no sources within $D'_i J_j$: every monomial m in $D'_i J_j$ either evaluates to 0 and nothing points to it, or has exactly one incoming and one outgoing edge.

If $J_j \upharpoonright \alpha \neq 0$: We define the successor pointer for the monomials in $D'_i J_j$ so that there is a source iff every $t_i(x) = 0$. Let $\text{Mons}(J_j)^+$, $\text{Mons}(J_j)^-$ be the (non-zero) positive and negative monomials in J_j respectively. Let

$$\delta := |\text{Mons}^+(J_j \upharpoonright \alpha)| - |\text{Mons}^-(J_j \upharpoonright \alpha)|$$

be the difference between the number of positive and negative monomials, and note that $\delta > 0$ as J_j is a conjunct and $J_j \upharpoonright \alpha \neq 0$. Recall that $D'_i J_j = \sum_{k \in [\ell]} t_d J_j - J_j$. For each term, we will define a matching so that $-J_j$ has only δ -many negative monomials without incoming edges, and every negative monomial in $t_d J_j$ has an incoming edge.

- For $-J_j$: Define an arbitrary pairing $P := \{(m, -m')\} \subseteq \text{Mons}^+(J_j \upharpoonright \alpha) \times \text{Mons}^-(J_j \upharpoonright \alpha)$ such that each positive monomial occurs in exactly one pair and each negative monomial occurs in at most one pair. Hence we have δ -many negative monomials that are not paired. For each pair $(m, -m') \in P$ define $S(m) = -m'$.

Note that we have now defined the successor of every positive monomial in J_j .

- For each $t_d J_j$: Observe that as t_d is a conjunct, under any assignment it contains at least as many positive monomials as negative monomials. Define an arbitrary pairing $P := \{(m, -m')\} \subseteq \text{Mons}^+(t_d J_j \upharpoonright \alpha) \times \text{Mons}^-(t_d J_j \upharpoonright \alpha)$ such that each negative monomial occurs in exactly one pair and each positive monomial occurs in at most one pair. For each pair $(m, -m') \in P$ define $S(m) = -m'$. Let $\beta \in \{0, 1\}^{\text{Vars}(t_d J_j)}$ be the assignment to the variables of $t_d J_j$ that was discovered by the trees made by the decision tree of any of the monomials m in $t_d J_j$. Let

$$c := |\text{Mons}^+(t_d \upharpoonright \beta)| - |\text{Mons}^-(t_d \upharpoonright \beta)|$$

be the difference between the number of non-zero positive and negative monomials in t_d under β .

If $c = 0$, then $t_d J_j \upharpoonright \beta = 0$ and so the the number of non-zero positive and negative monomials is equal. In this case, each negative monomial has an incoming edge which is provided by this pairing.

Otherwise, if $t_d J_j \upharpoonright \beta \neq 0$ then there are $c\delta$ -many non-zero positive monomials whose successor is still undefined, and partition them into c -many groups of C_1, \dots, C_δ of δ -many monomials each. Recall that $-J_j$ has exactly δ -many negative monomials with no incoming edge, $-m_1, \dots, -m_\delta$. For each $m \in C_i$ define $S(m) = -m_i$. In this case, each monomial in $t_d J_j$ and $-J_j$ has an incoming edge.

Finally, we define the successor for each positive monomial in each conjunct j_d , for some $k \in K$, in the conical junta. To do so, we use the fact that j_d contains at least as many positive monomials as negative monomials in order to ensure that there is never any source among the monomials of j_d . The successor for each positive monomial m of j_d queries the (at most d -many) variables in j_d for an assignment $\alpha \in \{0, 1\}^{\text{Vars}(j_d)}$. Define an arbitrary pairing $P := \{(m, -m')\} \subseteq \text{Mons}^+(j_d \upharpoonright \alpha) \times \text{Mons}^-(j_d \upharpoonright \alpha)$ such that each negative monomial occurs in exactly one pair and each positive monomial occurs in at most one pair. For each pair $(m, -m') \in P$ define $S(m) = -m'$. For the remaining positive monomials m in j_d whose successor is not defined, set $S(m) = 1$ (this choice is somewhat arbitrary).

This completes the description of the successor function S (the f -part of the formulation). It remains to define the output function g of the formulation. For each potential solution m ,

- If m is a monomial from some $D'_i J_j$, then D'_i is the weakening of some D_o of F , and we output o .
- Otherwise, we output an arbitrary index $o \in [m]$.

Finally, we prove that this formulation is correct. To do so, we show that the only monomials which do not have incoming edges belong to some $D'_i J_j$ for which $D'_i(x)$ is falsified. This suffices, as if m belongs to $D'_i J_j$ where $D'_i(x) = 0$ then $g_m(x) = o$ for some D_o of F of which D'_i is a weakening of. Hence, $D_o(x)$ is falsified and we have found a solution to FF_F . By the negative monomial case in the formulation, every positive monomial has an incoming edge. By the pairings constructed in the formulation, every negative monomial in each j_d in the conical junta also has an incoming edge. As well, for each $D'_i J_j = \sum_k t_d J_j - J_j$, each negative monomial in each $t_d J_j$ has an incoming edge. Hence, the only potential sources belong to the $-J_j$ terms of each $D'_i J_j$. As we argued before, if $J_j(x) = 0$ then there is no source in the monomials of $D'_i J_j$, so suppose that this is not the case. As we have paired off positive and negative monomials in $-J_j$, the only incoming edge to each of the δ -many remaining negative monomials of J_j must come from some $t_d J_j$. If there is a t_d such that $t_d(x) \neq 0$ (and hence $D'_i(x)$ is satisfied) then $t_d J_j$ has $c\delta$ -many monomials which map to the δ -many remaining negative monomials of J_j , meaning that there is no source in $D'_i J_j$. Thus, $D'_i J_j$ becomes a source only if $J_j(x) \neq 0$ and $D'_i(x)$ is falsified. \square

4.2 sRA Reductions Imply SA Proofs

We begin by observing that there is a trivial Sherali-Adams refutation of Range Avoidance:

$$\sum_{h \in [n+1]} \left(\sum_{p \in [n]} \llbracket p = h \rrbracket - 1 \right) = \sum_{p \in [n]} \sum_{h \in [n+1]} \llbracket p = h \rrbracket - (n+1) = n - (n+1) = -1,$$

where the third equality follows as we $\sum_{h \in [n+1]} \llbracket p = h \rrbracket = 1$.

In the remainder of this section we will show that Sherali-Adams can prove reductions to **STRONGAVOID**.

Lemma 4.8. *If f, g is a **STRONGAVOID**-formulation of FF_F of depth d and size s then there is a degree- $O(d \log n)$ and size $\text{poly}(s \cdot n^d)$ uSA proof of F .*

If (f, g) is a STRONGAVOID-formulation of FF_F for some formula $F = \bigwedge_{i \in [m]} D_i$, let $P(g_h), P(f_p)$ be the set of all root-to-leaf paths in the decision trees g_h and f_p respectively. As well, for any hole $h \in [n+1]$, let $P_h(f_p)$ be the set of paths in f_p whose leaf is labelled by hole h .

We can express the reduction from FF_F to STRONGAVOID as the unsatisfiable formula $\text{STRONGAVOID}(f, g)$ defined as

$$\begin{aligned} & \forall h \in [n+1], \forall \sigma^* \in P(g_h), \bigvee_{\sigma \neq \sigma^* \in P(g_h)} \bigvee_{p \in [n]} \llbracket p = h \rrbracket \wedge \sigma \\ &= \forall h \in [n+1], \forall \sigma^* \in P(g_h), \bigvee_{\sigma \neq \sigma^* \in P(g_h)} \bigvee_{p \in [n]} \bigvee_{\rho \in P_h(f_p)} \rho \wedge \sigma \end{aligned}$$

Letting $D_{h, \sigma^*} := \bigvee_{\sigma \neq \sigma^* \in P(g_h)} \bigvee_{p \in [n]} \bigvee_{\rho \in P_h(f_p)} \rho \wedge \sigma$, this becomes the unsatisfiable family of DNFs

$$\text{STRONGAVOID}(f, g) := \{D_{h, \sigma^*}\}_{h \in [n+1], \sigma^* \in P(g_h)}.$$

The following lemma shows that uSA can deduce $\text{STRONGAVOID}(f, g)$ from F .

Proof of Lemma 4.8. We will abuse notation and let $\llbracket p = h \rrbracket := \sum_{\delta \in P_h(f_p)} \delta$ denote the decision-tree substitution of the indicator $\llbracket p = h \rrbracket$. To begin, we will weaken F to $\text{STRONGAVOID}(f, g)$, the polynomials of which are

$$D_{h, \sigma^*} := \sum_{\sigma \neq \sigma^* \in P(g_h)} \sum_{p \in [n]} \llbracket p = h \rrbracket \cdot \sigma - 1$$

for $h \in [n+1]$ and $\sigma^* \in P(g_h)$. As each $\llbracket p = h \rrbracket$ contains $O(\log n)$ -many Boolean variables, and we are replacing each one by a depth- d decision tree, the degree of $\text{STRONGAVOID}(f, g)$ is $O(d \log n)$. Similarly, the size blows up by a factor of n^d .

For any $h \in [n+1]$,

$$\begin{aligned} \sum_{\sigma^* \in P(g_h)} D_{h, \sigma^*} &= \sum_{\sigma^* \in P(g_h)} \left(\sum_{\sigma \neq \sigma^* \in P(g_h)} \sum_{p \in [n]} \llbracket p = h \rrbracket \sigma - 1 \right) \\ &= \sum_{p \in [n]} \sum_{\sigma^* \in P(g_h)} \sum_{\sigma \neq \sigma^* \in P(g_h)} \llbracket p = h \rrbracket \sigma - |P(g_h)| \\ &= \sum_{p \in [n]} \llbracket p = h \rrbracket \sum_{\sigma^* \in P(g_h)} \sum_{\sigma \neq \sigma^* \in P(g_h)} \sigma - |P(g_h)| \\ &= \sum_{p \in [n]} \llbracket p = h \rrbracket \left(|P(g_h)| - 1 \right) \sum_{\sigma \in P(g_h)} \sigma - |P(g_h)| \\ &= \sum_{p \in [n]} \llbracket p = h \rrbracket \left(|P(g_h)| - 1 \right) - |P(g_h)| \quad (\text{Summing all paths in the DT } g(h)) \\ &= (|P(g_h)| - 1) \left(\sum_{p \in [n]} \llbracket p = h \rrbracket - 1 \right) \end{aligned}$$

By padding, we can assume without loss of generality that all decision trees g_h have the same number of paths; that is

$|P(g_h)| = |P(g_{h'})| = \alpha$ for all $h, h' \in [n+1]$, and some $\alpha \in \mathbb{N}$ with $\alpha > 1$. Then,

$$\begin{aligned}
\sum_{h \in [n+1]} \sum_{\sigma^* \in P(g_h)} D_{h, \sigma^*} &= \sum_{h \in [n+1]} (\alpha - 1) \left(\sum_{p \in [n]} \llbracket p = h \rrbracket - 1 \right) \\
&= (\alpha - 1) \left(\sum_{p \in [n]} \sum_{h \in [n+1]} \llbracket p = h \rrbracket - (n+1) \right) \\
&= (\alpha - 1) \left(\sum_{p \in [n]} \sum_{h \in [n+1]} \sum_{\delta \in P_h(p)} \delta - (n+1) \right) \\
&= (\alpha - 1) \left(\sum_{p \in [n]} \sum_{\delta \in P(p)} \delta - (n+1) \right) \\
&= (\alpha - 1) \left(\sum_{p \in [n]} 1 - (n+1) \right) \quad (\text{Summing all paths in the DT } f(p)) \\
&= (\alpha - 1)(n - (n+1)) = -(\alpha - 1) \leq -1.
\end{aligned}$$

□

5 A Generic Correspondence

In this section we establish a general correspondence between syntactic subclasses of total search problems in the polynomial hierarchy and proof systems. Our characterizations will rely on the following two properties of a Σ_d -proof system:

- *Reduction-Closed.* For unsatisfiable $\Pi_{d,5}$ formulas F, H , if P has a complexity- s proof of F and there is a complexity- c FF_F -formulation of FF_H then $P(H) = \text{poly}(cs)$.
- *Reflective.* P has $\text{polylog}(n)$ -complexity proofs of a *reflection principle* about itself—a formula encoding the soundness of this proof system; we expand on the meaning of this below.

We show the following, generalizing [BFI23].

Theorem 5.1. *The following hold:*

- i) *Every syntactic subclass of $\text{TF}\Sigma_d$ is characterized by a Σ_d -proof system.*
- ii) *Every Σ_d -proof system which is reduction-closed and reflective is characterized by a subclass of $\text{TF}\Sigma_d$.*

We prove (i) in [subsection 5.1](#) and (ii) in [subsection 5.2](#).

5.1 A Proof System for any $\text{TF}\Sigma_d$ Problem

In this section we show how to construct a proof system from any total search problem $R \subseteq \{0, 1\}^n \times \mathcal{O}$, which we think of as the complete problem for some syntactic subclass. The key insight is that one can view a decision tree reduction from a total search problem $Q \subseteq \{0, 1\}^m \times \mathcal{O}_Q$ to R as a *proof* that Q is total, if we take the totality of R as an axiom. In what follows we formalize this intuition. We define proofs in the *canonical proof system* for a $\text{TF}\Sigma_d^{dt}$ subclass as reductions to one of its complete problems.

Definition 5.2. Let $\text{FF}_F \in \text{TF}\Sigma_d^{dt}$ where $F = \bigwedge_{o \in [m]} F_o$. The *canonical proof system* for FF_F , denoted P_F , is defined as follows. A proof Π in P_F consists of a triple $(f, g, F(f, g))$, where

- (f, g) is a FF_F -formulation (i.e., a set of decision trees), and
- $F(f, g) = \bigwedge_{o \in [m^*]} L_o$ is the [reduced formula](#) associated with this formulation.

Π is a P_F proof of an unsatisfiable formula $H = \bigwedge_{t \in [m']} H_t$, where each H_t is a Σ_d -formula, if for every L_o in $F(f, g)$ there exists some $t \in [m']$ such that L_o is a Σ_d -weakening of H_t ; that is,

$$H_t \implies L_o.$$

The *size* of the proof Π is the number of bits needed to write down Π , and the width of Π is the maximum depth among the decision trees in the formulation,

$$\text{depth}(\Pi) := \max_{i \in [n], o \in [m]} \{\text{depth}(f_i), \text{depth}(g_o)\}.$$

The *complexity* of proving H in P_F is the minimum over all P_F -proofs of H ,

$$P_F(H) := \min \{\text{width}(\Pi) + \log \text{size}(\Pi) : \Pi \text{ is a } P_F\text{-proof of } H\}.$$

This proof system is sound, since any substitution of an unsatisfiable formula remains unsatisfiable. As well, it is complete for unsatisfiable Π_{d+1} formulas as depth- n decision trees suffice to solve any total search problem. Note that this proof system agrees with the definition of [BFI23] when $d = 1$.

We will show that P_F characterizes the subclass with complete problem FF_F , proving the first direction of [Theorem 5.1](#).

Lemma 5.3. *If $\text{FF}_F, \text{FF}_H \in \text{TF}\Sigma_d^{dt}$ then there is a complexity- c FF_F -formulation of FF_H iff $P_F(H) \leq c \cdot \text{polylog}(n)$.*

Proof. Let (f, g) be a complexity- c FF_F -formulation of FF_H . We claim that $(f, g, \text{FF}_F(f, g))$ is a P_F proof of H . As $\text{FF}_F \in \text{TF}\Sigma_d^{dt}$, F is a $\Pi_{d.5}$ formula, and so the [reduced formula](#) $\text{FF}_F(f, g)$ is a Π_{d+1} -formula ($\Pi_{d.5}$ if $c = \text{polylog}(n)$). As well, the size of $\text{FF}_F(f, g)$ is at most $\text{size}(\text{FF}_F) \cdot \exp(O(c))$, as each clause/term on the bottom layer of F has width at most $\text{polylog}(n)$ and we replace it by the CNF/DNF representation of a depth- $O(c)$ decision tree, which has width $O(c)$ and size at most $\exp(O(c))$. Finally, for $F(f, g) := \bigwedge_{o \in [m^*]} L_o$ and $H := \bigwedge_{t \in [m]} H_t$, by the correctness of the formulation, we can conclude that for every $o \in [m^*]$ there exists some $t \in [m']$ such that $H_t \implies L_o$, and so L_o is a Σ_d -weakening of H_t .

For the converse direction, suppose that $(f, g, F(f, g))$ is a P_F proof of an unsatisfiable formula $H := \bigwedge_{i \in [m]} H_i$, where each H_i is a Σ_d -formula. By definition, (f, g) constitutes a complexity- c FF_F -formulation of FF_H . Indeed, each decision tree of (f, g) has depth at most c and there are at most 2^c -many of them, and so this is a complexity- c formulation. \square

5.2 A $\text{TF}\Sigma_d$ Problem for any Proof System which Reflects

In this section we show that a Σ_d -proof system P corresponds to a $\text{TF}\Sigma_d$ -problem if that proof system is *reduction closed* and *reflective*.

A reflection principle states that P -proofs are sound; we will restrict ourselves to proofs of $\Sigma_{d.5}$ formulas. Typically, the provability of a proof system's reflection principle is sufficient in order to simulate that system. In our setting, a reflection principle will falsely assert that there is a complexity- c P -proof Π of a $\Sigma_{d.5}$ -formula H and that H is satisfied by a truth assignment α :

$$\text{REF}_P := \text{PROOF}(H, \Pi) \wedge \text{SAT}(H, \alpha).$$

This formula will be parameterized by n_H , the number of variables of H , as well as c the complexity of the proof Π .

SAT. The formula $\text{SAT}(H, \alpha)$ states that $\alpha \in \{0, 1\}^{n_H}$ is a satisfying assignment to H , where $\alpha \in \{0, 1\}^n$ and H are given as input. A generic $\Pi_{d.5}$ -formula has the following structure:

$$H = \bigwedge_{o \in \mathcal{O}} \bigvee_{z_1 \in \{0, 1\}^{\ell_1}} \bigwedge_{z_2 \in \{0, 1\}^{\ell_2}} \dots \bigcirc_{z_{d-1} \in \{0, 1\}^{\ell_{d-1}}} H_{o, z_1, \dots, z_{d-1}}$$

where $\bigcirc \in \{\wedge, \vee\}$ and $H_{o, z_1, \dots, z_{d-1}}$ is a width $w \in \text{polylog}(n)$ clause if $\bigcirc = \wedge$ or conjunct if $\bigcirc = \vee$. Each $H_{o, \bar{z}} := H_{o, z_1, \dots, z_{d-1}}$ is specified by w -many $(2n+1)$ -ary variables $v_{o, z, 1}, \dots, v_{o, z, w} \in [2(n+1)]$, where $v_{o, z, i} = j$ denotes the variable

- x_j if $i \in [n]$,
- $\neg x_{j-n}$ if $j \in \{n+1, \dots, 2n\}$,
- constant 1 if $j = 2n+1$,
- constant 0 if $j = 2n+2$.

We could allow the formula REF_P to be parameterized by $|\mathcal{O}|, \ell_1, \dots, \ell_{d-1}$. However, for simplicity, since we are considering complexity- c proofs, it suffices to simply set all of these parameters to 2^c and $w = c$. In this case, the number of $H_{o,\vec{z}}$ is 2^{cd} , and hence the number of Boolean variables of H is $c \log(2n_H + 2) \cdot 2^{cd}$. Then the $\Pi_{d,5}$ formula SAT can be written as

$$\text{SAT}(H, \alpha) := \bigwedge_{o \in \mathcal{O}} H_o(\alpha) := \bigwedge_{o \in \mathcal{O}} \bigvee_{z_1 \in \{0,1\}^{\ell_1}} \bigwedge_{z_2 \in \{0,1\}^{\ell_2}} \dots \bigcirc_{z_{d-1} \in \{0,1\}^{\ell_{d-1}}} \llbracket H_{o,\vec{z}}(\alpha) = 1 \rrbracket,$$

where $\llbracket H_{o,\vec{z}}(\alpha) = 1 \rrbracket$ is the width- $O(w \log n_H)$ DNF (if $\bigcirc = \vee$) or CNF (if $\bigcirc = \wedge$) defined by the following decision tree $T_{o,\vec{z}}$: First query the $w \log(2n_H + 2)$ -many Boolean variables $H_{o,\vec{z},1}, \dots, H_{o,\vec{z},w}$ to determine the literals ℓ_1, \dots, ℓ_w of $H_{o,\vec{z}}$. Then, query the corresponding bits of α to determine if $H_{o,\vec{z}}$ is satisfied. If it is, then $T_{o,\vec{z}}$ outputs 1 and otherwise it outputs 0. This can be converted into a DNF or CNF in the usual way.

Proof. The formula $\text{PROOF}(H, \Pi)$ states that Π is a P -proof of H . A complication is that there are many different ways by which one could encode a P -proof as a formula, some of which may change the difficulty of proving the reflection principle drastically. Following [BF123] we define one reflection principle for each encoding of a P -proof; we call such an encoding a *verification procedure*.

Definition 5.4. A *verification procedure* V for a Σ_d -proof system P , parameterized by n_H, c , is $\Pi_{d,5}$ -formula which generically encodes a complexity- c P -proof Π of an n_H -variate formula H . Specifically, the formula $V_{n_H,c}(\Pi, H)$ has two sets of variables H, Π , where:

- An assignment to the variables $H = \{H_{o,\vec{z},i} | i \in [n_H]\}$ specifies a $\Pi_{d,5}$ formula as before.
- An assignment to the variables Π specifies a purported P -proof of H of complexity c , such that any error in Π can be verified by an efficient Σ_{d-1} -algorithm (placing $\text{REF} \in \text{TF}\Sigma_d$).
- V has $2^{\Theta(c)}$ -many variables.

As c bounds the logarithm of the size of the proof, and the number of variables is exponential in $\Theta(c)$, the second condition ensures that a violated sub-formula of V can be verified by a Σ_{d-1} -algorithm making $\text{polylog}(c)$ -many queries.

A *reflection principle* for a proof system P and verification procedure V is

$$\text{REF}_{P,V} := \text{PROOF}_{n_H,c}(H, \Pi) \wedge \text{SAT}_{n_H,c}(H, \alpha),$$

where $\text{PROOF}_{n_H,c}(H, \Pi) := V_{n_H,c}(H, \Pi)$. Often, we will suppress the subscripts P, V .

We now prove point (ii) of Theorem 5.1.

Lemma 5.5. Let P be a Σ_d -proof system that is reduction closed and reflective for some $\text{REF} := \text{REF}_{P,V}$. Then for any $\text{FF}_H \in \text{TF}\Sigma_d$,

- If there is a complexity- c FF_{REF} -formulation of FF_H then $P(H) = \text{poly}(c \cdot P(\text{REF}))$.
- There is a complexity $O(P(H))$ FF_{REF} -formulation of FF_H .

Proof. To prove (i), suppose that there is a complexity- c FF_{REF} -formulation of H . By the definition of being reduction closed, there is a P proof of H of complexity $\text{poly}(c \cdot P(\text{REF}))$.

For (ii), let Π be a complexity- c proof of H in P . We construct a FF_{REF} -formulation (f, g) of FF_H as follows. f will hard-wire (Π, H) as the input to REF , and map the input variables of FF_H to the variables $\alpha_1, \dots, \alpha_{n_H}$ of REF . Since Π is a valid proof of H , $\text{PROOF}(\Pi, H)$ is always satisfied and we can set g_o arbitrarily for any solution o corresponding to a subformula of $\text{PROOF}(\Pi, H)$. As $\text{PROOF}(\Pi, H)$ is always satisfied under this reduction the only solutions which may occur belong to $\text{SAT}(H, \alpha)$. In particular, as we have mapped the input variables of H to the bits $\alpha_1, \dots, \alpha_{n_H}$, for any assignment $x \in \{0, 1\}^n$, $H_o(x) = 0 \iff H_o(\alpha) = 0$. Hence, we define $g_o = o$. \square

6 Characterizations in $\text{TF}\Sigma_2$

In this section we uncover $\text{TF}\Sigma_2$ characterizations of several well-studied proof systems — DNF Resolution, DNF Circular Resolution [AL23, DR23], and DNF Reversible Resolution [GHJ⁺22b, DR23]. Along the way we introduce several new $\text{TF}\Sigma_2$ classes which are inspired by TFNP classes. These are analogs to the *coloured* TFNP classes introduced in [KST07a, DR23]. In subsection 6.3 we explore the relationships between these and prominent $\text{TF}\Sigma_2$ subclasses.

The *DNF resolution* proof systems are extensions of the resolution proof system (and restrictions of) to allow them to operate with DNF formulas, rather than only clauses. Davis and Robere [DR23] gave characterizations of these systems by coloured TFNP classes. We introduce several classes which characterize the Σ_2 -variants of these proof systems; we believe these $\text{TF}\Sigma_2$ classes *herbrandize* to the coloured classes.

Definition 6.1. A $\text{Res}(\text{polylog})$ *refutation* of a Π_2 -unsatisfiable formula $F = \bigwedge_{i=1}^m A_i$ is a sequence of $\text{polylog}(n)$ -width DNF formulas $\Pi = (D_1, \dots, D_s = \perp)$ where each D_i is deduced from previous DNFs by one of the following rules:

- *Axiom Introduction.* Introduce A_i for some $i \in [m]$.
- *Symmetric Cut.* From $D \vee t$ and $D \vee \bar{t}$ derive D , where t is any term.
- *Reverse Cut.* From D derive $D_i = D \vee t$ and $D_{i+1} = D \vee \bar{t}$, for some term t .

The *size* s of Π is the sum of the sizes of DNFs involved in Π , and the *width* w is the maximum width of any DNF in Π . The *complexity* of Π is $\log s + w$.

A $\text{RevRes}(\text{polylog})$ proof is a $\text{Res}(\text{polylog})$ proof in which every DNF in the sequence is used as the premise to a derivation rule at most once.

A $\text{uCircRes}(\text{polylog})$ proof has access to the additional rule

- *DNF Creation.* $S_i = S_{i-1} \cup \{D\}$, where D is any DNF formula.

provided that each copy of D that is created in this way is derived at least as many times as it is used as the premise to a derivation rule.

The following technical lemma will be key to our characterizations.

Lemma 6.2. [Theorem 3.6 in [DR23]] $\text{Res}(\text{polylog})$, $\text{RevRes}(\text{polylog})$, and $\text{uCircRes}(\text{polylog})$ are reduction closed.

Davis and Robere proved Lemma 6.2 for DNF resolution proofs of $\Pi_{1.5}$ -formulas (that is, when the axioms are clauses). It is straightforward to see that it holds by exactly the same argument (Claim 1) when the axioms are DNF formulas. In section 7 we prove this theorem for depth- $d.5$ Frege, for every d , of which $\text{Res}(k)$ is $d = 1$.

In the following subsections we will prove Theorem 1.5, characterizing each of these proof systems by new $\text{TF}\Sigma_2^{dt}$ subclasses. To define each of these classes it will be convenient to use the following notion of a *meta-pointer*.

Definition 6.3. Given a function $S : [m] \times [t] \rightarrow [m]$, the *meta-pointer* $\tilde{S} : [m] \rightarrow [m] \cup \{\text{undefined}\}$ is defined as

$$\tilde{S}(u) = \begin{cases} v & \text{if for every } i \in [t], S(u, i) = v, \\ u & \text{if there is } i \in [t] \text{ such that } S(u, i) = u \\ \text{undefined} & \text{if there is } i, j \in [t] \text{ such that } u \neq S(u, i) \neq S(u, j). \end{cases}$$

Note that, if $u \neq v$, $\tilde{S}(u) = v$ is Π_1 -verifiable: For all $i \in [t]$, we need to verify that $S(u, i) = v$, which takes $\log(m)$ queries. Moreover, $\tilde{S}(u) = u$ and $\tilde{S}(u) = \text{undefined}$ are Σ_1 verifiable: We can non-deterministically guess $i \in [t]$ such that $S(u, i) = i$, or $i \neq j \in [t]$ such that $u \neq S(u, i) \neq S(u, j)$, in other words, they are efficiently computable if we are given i (and j) as witnesses. The inclusion in $\text{TF}\Sigma_2$ of the problems presented in this chapter follows directly from this fact.

6.1 DNF Resolution

$\text{polylog}(n)$ -width resolution was characterized by the TFNP^{dt} subclass PLS [BKT14]. In this section we introduce a $\text{TF}\Sigma_2$ -variant of the PLS-complete problem *iteration* and show that it characterizes $\Sigma_2\text{-Res}(\text{polylog})$. The *iteration* problem encodes the principle that *every DAG has a sink*. The input is given by a pointer function $S : [n] \rightarrow [n]$ giving the *successor* of a node $u \in [n]$, thought of as the next node on a root-to-leaf walk in the dag. A solution is (i) an invalid source $S(1) = 1$, (ii) a u which points backwards $S(u) < u$, (iii) a sink: $u \in [n]$ such that $S(u) \neq u$ but $S(S(u)) = S(u)$, or (iv) a node u with an undefined pointer $S(u) = \text{undefined}$. Our $\text{TF}\Sigma_2$ variant obfuscates the successor function. Similar ideas were used to define the RWPHP_2 problem in [KT21].

Definition 6.4. An instance of ITER_2 is given by a function $S : [m] \times [t] \rightarrow [m]$. A solution is a *witness* of a solution to the iteration instance defined by the meta-pointer \tilde{S} :

- (u, i, i') such that $S(u, i), S(u, i') \neq u$ and $S(u, i) \neq S(u, i')$, ($\tilde{S}(u)$ is undefined.)
- (u, i) such that $S(u, i) < u$. (A pointer which points backwards)
- $(1, i)$ if $S(1, i) = 1$. (1 is not a source)
- (u, v, i) such that $\tilde{S}(u) = v$ and $S(v, i) = v$. (v is a proper sink)

The class PLS_2^{dt} is the set of $R \in \text{TF}\Sigma_2^{dt}$ such that $R \leq_{dt} \text{ITER}_2$.

Theorem 6.5. For any $\text{FF}_F \in \text{TF}\Sigma_2^{dt}$, there is a complexity- c ITER_2 -formulation of FF_F iff there is a complexity $O(c)$ $\Sigma_2\text{-Res}(\text{polylog})$ proof of F .

We prove this theorem in the following two lemmas, each giving one direction.

Lemma 6.6. For $\text{FF}_F \in \text{TF}\Sigma_2$, if $\Sigma_2\text{-Res}(\text{polylog})(F) = c$ then there is a complexity- $O(c)$ ITER_2 -formulation of FF_F .

Proof. Let (Π, H) be a $\Sigma_2\text{-Res}(\text{polylog})(F)$ proof of $F = \bigwedge_{i \in [\ell]} F_i$, where $H = \bigwedge_{i \in [k]} A_i$ and each A_i is a Σ_2 -weakening of a DNF of F . Up to padding, we may assume that each DNF in the proof has the same number of terms t . Consider the proof $\Pi = D_1, \dots, D_m$ in reverse order so that $D_1 = \perp$; this will be our designated source.

Let $t_{u,i}$ be the i^{th} term of D_u . Given an assignment $\alpha \in \{0, 1\}^n$ to the variables of F , we construct a function $S_\alpha : [m] \times [t] \rightarrow [m]$ by setting $S_\alpha(u, i)$ to be:

- u if D_u is an axiom, or if $t_{u,i}(\alpha) = 1$;
- v if $t_{u,i}(\alpha) = 0$ and D_u was derived from D_v by the reverse cut rule or semantic weakening of an axiom;
- v if $t_{u,i}(\alpha) = 0$ and D_u was derived from $D_v = D_u \vee t$ and $D_w = D_u \vee \bar{t}$ via symmetric cut and $t(\alpha) = 0$ and w if $\bar{t}(\alpha) = 0$;

Finally, for each solution o to the instance S_α we define the output of the reduction $g_o(\alpha)$ to be arbitrary if o does not correspond to an axiom A_i of H , and otherwise this axiom A_i is a weakening of a DNF F_j of F , and we set $g_o(\alpha) = j$. Note that in this case $A_i(\alpha) = 0 \implies F_j(\alpha) = 0$. Observe that computing $S_\alpha(u, i)$ involves evaluating at most two terms, and hence the depth of the reduction is at most twice the width of the proof. It remains to argue that the reduction is correct.

Claim. The function \tilde{S}_α satisfies the following properties:

- i) \tilde{S}_α is defined everywhere.
- ii) If D_u is not an axiom of H then $D_u(\alpha) = 0$ iff $\tilde{S}_\alpha(u) \neq u$.
- iii) If $\tilde{S}_\alpha(u) = v \neq u$, then $D_v(\alpha) = 0$.

Assuming the claim, we see that the only type of solution to this ITER_2 instance S_α are proper sinks corresponding to falsified axioms of H , which are weakenings of (falsified) axioms of F . Hence, g returns a correct solution to $\text{FF}_F(\alpha)$.

Proof of Claim. We prove each item, beginning with (i). Clearly \tilde{S}_α is well defined for any u that was not derived using the cut rule since $S_\alpha(u, i)$ only has one choice of value other than u . So now consider u such that D_u was

derived from $D_v = D_u \vee t$ and $D_w = D_u \vee \bar{t}$. For $i \in [t]$, we see that $S_\alpha(u, i)$ depends on two values: $t_{u,i}(\alpha)$, and $t(\alpha)$ in the case where $t_{u,i}(\alpha) = 0$. Thus, $t(\alpha)$ being independent of i , $S_\alpha(u, i)$ is always identical when not equal to u .

(ii) follows from the fact that $D_u(\alpha) = 0$ iff $t_{u,i}(\alpha) = 0$ for all i , and $\tilde{S}_\alpha(u) = u$ iff $t_{u,i}(\alpha) = 1$ for at least one i . Finally, (iii) follows by definition. \square

We will now prove the converse. First, we describe the encoding of ITER_2 as a unsatisfiable formula. For each $(u, i) \in [m] \times [t]$, the m -ary value of $S_{u,i}$ will be described by $\log m$ -many boolean variables $S_{u,i,b}$, where the indicator function

$$\llbracket S_{u,i} = v \rrbracket := \bigwedge_{b \in [\log m]} S_{u,i,b}^{v_b},$$

where we think of v as being written in its binary encoding, v_b is its b^{th} bit, and $S_{u,i,b}^1 = S_{u,i,b}$ and $S_{u,i,b}^0 = \neg S_{u,i,b}$. As well, $\llbracket S_{u,i} \neq v \rrbracket = \neg \llbracket S_{u,i} = v \rrbracket$, and

$$\llbracket \tilde{S}_u \neq v \rrbracket := \bigvee_{i \in [t]} \llbracket S_{u,i} \neq v \rrbracket.$$

Then ITER_2 is the conjunction of the following subformulae:

- $\llbracket S_{1,i} \neq 1 \rrbracket$ for each $i \in [n]$. (1 is not a source)
- $\llbracket S_{u,i} \neq v \rrbracket \vee \llbracket S_{u,i'} \neq v' \rrbracket$ for all $v \neq v'$ and $i \neq i'$ such that $u \neq v, v'$ (\tilde{S} is defined everywhere)
- $\llbracket S_{u,i} \neq v \rrbracket$ for all $v < u$ and $i \in [n]$. (Nothing points backwards)
- $\llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket S_{v,j} \neq v \rrbracket$ for all $u < v$ and $j \in [n]$. (v is not a proper sink)

Note that the subformulae of the ITER_2 formula are clauses making the formula a CNF. We may then question what makes ITER_2 a $\text{TF}\Sigma_2^{dt}$ problem and not a TFNP^{dt} one. The key to understanding this resides in the size of said clauses. Indeed, for a the false formula problem corresponding to a CNF to be in TFNP^{dt} , we need to be able to verify if a given clause is falsified by an assignment by only querying a $\text{polylog}(n)$ amount of bits. This in turns directly implies that we would need each clause to be of $\text{polylog}(n)$ -width. This is not the case here because of the fourth type of axioms which are of $\text{poly}(n)$ -width. On the other hand, considering clauses a 1-width DNFs, we see that this false formula problem corresponding to this formula lands indeed in $\text{TF}\Sigma_2^{dt}$. We now state the converse.

Lemma 6.7. *For $\text{FF}_F \in \text{TF}\Sigma_2$, if there is a complexity- c ITER_2 -formulation of FF_F then there is a complexity- $O(c)$ $\text{Res}(\text{polylog})$ proof of F .*

Observe that the set of formulas $\{\llbracket S_{u,i} \neq v \rrbracket\}_{v \in [m]}$ contains all clauses containing all of the variables $S_{u,i,b}$. Hence they can be cut in $O(m \log m)$ -many steps to obtain \perp . Throughout the proof we will write the

$$\frac{D_v \vee \llbracket S_{u,i} \neq v \rrbracket, \forall v}{D}$$

as a shorthand for this derivation with $D = \bigvee_{v \in [m]} D_v$.

Proof of Lemma 6.7. By Lemma 6.2 it suffices to show that $\text{Res}(\text{polylog})$ can prove ITER_2 . By induction from $u = m$ to $u = 1$ we will derive a set of formulae that state that does not point forward in \tilde{S} . Combining this with the fact that the image of u by \tilde{S} cannot be undefined and u may not point backwards, this is semantically equivalent to stating that u points to itself. We then reach a contradiction when reaching $u = 1$ since 1 must be a proper source of our graph. This will be achieved by deducing

$$L_u := \{\llbracket \tilde{S}_u \neq v \rrbracket : u < v\},$$

which can be combined with axioms stating that no node points backwards for the desired statement.

The base case is trivial, as $L_m = \emptyset$. Consider some $u \in [m]$ and suppose that we have derived L_v for all $v > u$. We derive the formula $\llbracket \tilde{S}_u \neq v \rrbracket \in L_u$ as follows: consider some $w > v > u$ and apply the reverse cut rule to $\llbracket \tilde{S}_v \neq w \rrbracket$ in order to obtain $\llbracket \tilde{S}_v \neq w \rrbracket \vee \llbracket \tilde{S}_u \neq v \rrbracket$. Now consider the cuts from $a = t$ to $a = 2$

$$\frac{\llbracket \tilde{S}_u \neq w \rrbracket \vee \bigvee_{i < a+1} \llbracket S_{v,i} \neq v \rrbracket \quad \llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket S_{v,a} \neq v \rrbracket \quad \llbracket S_{v,1} \neq w \rrbracket \vee \llbracket S_{v,a} \neq w' \rrbracket, \forall w' \neq v, w}{\llbracket \tilde{S}_u \neq v \rrbracket \vee \bigvee_{i < a} \llbracket S_{v,i} \neq v \rrbracket}$$

to the set of formulae $\llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket S_{v,1} \neq w \rrbracket$. Finally, we do one last cut

$$\frac{\llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket S(v,1) \neq w \rrbracket, \forall w > v \quad \llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket S_{v,1} \neq w \rrbracket \quad \llbracket S_{v,1} \neq w \rrbracket, \forall w < v}{\llbracket \tilde{S}_u \neq v \rrbracket},$$

which derives the formula $\llbracket \tilde{S}_u \neq v \rrbracket \in L_u$.

Finally, once we have derived L_1 we can derive \perp as follows. For a fixed $v > 1$, starting from $a = t$ down to $a = 2$ we operate the cuts

$$\frac{\bigvee_{i < a+1} \llbracket S_{1,i} \neq v \rrbracket \quad \llbracket S_{1,a} \neq 1 \rrbracket \quad \llbracket S_{1,1} \neq v \rrbracket \vee \llbracket S_{1,n} \neq v' \rrbracket, \forall v' \neq v, 1}{\bigvee_{i < a} \llbracket S_{1,i} \neq v \rrbracket}.$$

Once we have derived $\llbracket S_{1,1} \neq v \rrbracket$, we do one final cut

$$\frac{\llbracket S_{1,1} \neq v \rrbracket, \forall v \neq 1 \quad \llbracket S_{1,1} \neq 1 \rrbracket}{\perp} \quad \square$$

6.2 Circular and Reversible DNF Resolution

In this section we characterize the Σ_2 -uCircRes(polylog) proof system by a $\text{TF}\Sigma_2$ -variant of the *Sink-of-Line* problem. An instance of Sink-of-Line is given by functions $S, P : [m] \times [t] \rightarrow [m] \cup \{\text{undefined}\}$ which define a graph G as follows: there is a directed edge (u, v) if $\tilde{S}(u) = v$ and $\tilde{P}(v) = u$. A solution to this instance is either i) 1 if 1 is not a source in G , ii) a sink u in G , iii) a vertex u for which $\tilde{P}(u)$ or $\tilde{S}(u)$ is undefined. We now describe the $\text{TF}\Sigma_2$ variant.

Definition 6.8. An instance of SoL_2 is given by functions $S, P : [m] \times [t] \rightarrow [m]$. A solution is a witness to a solution to the SOL instance defined by the meta-pointers (\tilde{S}, \tilde{P}) :

- (u, i, i') if $S(u, i') \neq u$ and $S(u, i) \neq S(u, i')$ or $P(u, i), P(u, i') \neq u$ and $P(u, i) \neq P(u, i')$.
(Predecessor or Successor of u is undefined)
- $(1, i)$ if $S(1, i) = 1$ or $\tilde{S}(1) = v \neq 1$ and $P(v, i) \neq 1$.
(1 is not a source)
- (u, i) if $u \neq 1$ and $S(u, i) = 1$.
(u has a pointer to 1)
- (u, v, i) for $u \neq v$ if $\tilde{S}(u) = v, \tilde{P}(v) = u$ and $S(v, i) = v$; or $\tilde{S}(u) = v, \tilde{P}(v) = u, \tilde{S}(v) = w$ and $P(w, i) \neq v$.
(v is a proper sink) [Noah: make sure this reads okay]

Theorem 6.9. For any $\text{FF}_F \in \text{TF}\Sigma_2$, there is a complexity- c SoL_2 -formulation of FF_F iff there is a complexity $O(c)$ Σ_2 -uCircRes(polylog) proof of F .

This theorem follows by combining Lemma 6.10 and Lemma 6.14. We begin with the backwards direction, showing that uCircRes(polylog) can prove SoL_2 formulations. SoL_2 is encoded as an unsatisfiable formula which is the conjunction of the following

- $\llbracket S_{u,i} \neq 1 \rrbracket$ for $u \in [m], i \in [t]$, and $\llbracket \tilde{S}_1 \neq v \rrbracket \vee \llbracket P_{u,i} \neq v \rrbracket$ for all $u, v \neq 1, i \in [t]$.
(1 is a source)
- $\llbracket S_{u,i} \neq v \rrbracket \vee \llbracket S_{u,i'} \neq v' \rrbracket$ for all $i \neq i', v \neq v'$.
(\tilde{S} is not undefined)
- $\llbracket P_{u,i} \neq v \rrbracket \vee \llbracket P_{u,i'} \neq v' \rrbracket$ for all $i \neq i', v \neq v'$.
(\tilde{P} is not undefined)
- $\llbracket S_{u,i} \neq 1 \rrbracket$ for all $i \in [t]$ and $u \neq 1$.
(Nothing points to 1)
- Let $\bar{E}_{u,v} := \llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket \tilde{P}_v \neq u \rrbracket$, we include
(No proper sinks)
 - $\bar{E}_{u,v} \vee \llbracket S_{v,i} \neq v \rrbracket$ for each $u \neq v$ and $i \in [m]$, and
 - $\bar{E}_{u,v} \vee \llbracket \tilde{S}_v \neq w \rrbracket \vee \llbracket P_{w,k} \neq w' \rrbracket$ for $u \neq v \neq w \neq w'$ and $k \in [t]$.

Lemma 6.10. For $\text{FF}_F \in \text{TF}\Sigma_2$, there is a complexity-uCircRes(polylog)(F) SoL_2 -formulation of FF_F .

Proof. By Lemma 6.2 it suffices to show that uCircRes(polylog) can prove SoPL_2 . For each $u \in [m]$ we would like to derive the set of formulas

$$L_u = \{\bar{E}_{u,v} : v \neq u, 1\},$$

stating that u has no outgoing edges. Our proof will proceed by the following three steps:

1. Assume L_u for each $u \neq 1$;

2. From L_v for $v \neq u$, deduce L_u . Since L_v is semantically equivalent to saying that node v points to itself, if u were to point to any other node, then said node would be a proper sink. Hence L_u follows.
3. L_1 is in direct contradiction with axioms stating that 1 is a source.

For step 1, we use the *DNF creation* rule

$$\overline{\bar{E}_{u,v}}$$

For step 2 and $u \in [m]$, we perform the following. For $w \neq v \neq u$ with $w, v \neq 1$, consider $\bar{E}_{v,w} \in L_v$ and weaken it successively to get

$$\overline{\bar{E}_{u,v} \vee \bar{E}_{v,w}},$$

then we cut as follows: starting with $c = n$ down to $c = 1$,

$$\frac{\bar{E}_{u,v} \vee [\tilde{S}_v \neq w] \vee \bigvee_{k < c+1} [P_{w,k} \neq v] \quad \bar{E}_{u,v} \vee [\tilde{S}_v \neq w] \vee [P_{w,c} \neq w'], \forall w' \neq w}{\bar{E}_{u,v} \vee [\tilde{S}_v \neq w] \vee \bigvee_{k < c} [P_{w,k} \neq v]}$$

to get $\bar{E}_{u,v} \vee [\tilde{S}_v \neq w]$. Next, starting from $b = n$ down to $b = 2$,

$$\frac{\bar{E}_{u,v} \vee \bigvee_{j < b+1} [S_{v,j} \neq w] \quad \bar{E}_{u,v} \vee [S_{v,b} \neq v] \quad [S_{v,1} \neq w] \vee [S_{v,b} \neq w'], \forall w' \neq v, w}{\bar{E}_{u,v} \vee \bigvee_{j < b} [S_{v,j} \neq w]}$$

and end up with the formulae $\bar{E}_{u,v} \vee [S_{v,1} \neq w]$. Finally,

$$\frac{\bar{E}_{u,v} \vee [S_{v,1} \neq w], \forall w \neq v, 1 \quad \bar{E}_{u,v} \vee [S_{v,1} \neq v] \quad [S_{v,1} \neq 1]}{\bar{E}_{u,v}}$$

derives $\tilde{E}_{u,v} \in L_u$. Having derived L_1 allows us to take $\bar{E}_{1,v} \in L_1$ and, starting with $b = n$ down to $b = 1$, we may cut

$$\frac{[\tilde{S}_1 \neq v] \vee \bigvee_{j < b+1} [P_{v,j} \neq 1] \quad [\tilde{S}_1 \neq v] \vee [P_{v,b} \neq w], \forall w \neq 1}{[\tilde{S}_1 \neq v] \vee \bigvee_{j < b} [P_{v,j} \neq 1]}$$

to get $[\tilde{S}_1 \neq v]$ for each $v \neq 1$. Next, starting from $a = n$ down to $a = 2$, we cut

$$\frac{\bigvee_{i < a+1} [S_{1,i} \neq v] \quad [S_{1,a} \neq 1] \quad [S_{1,1} \neq v] \vee [S_{1,a} \neq v'], \forall v' \neq v}{\bigvee_{i < a} [S_{1,i} \neq v]}$$

to get $[S_{1,1} \neq v]$ for $v \neq 1$. We may then cut one final time

$$\frac{[S_{1,1} \neq v], \forall v \neq 1 \quad [S_{1,1} \neq 1]}{\perp}.$$

□

We delay the proof of the other direction until the end of this section, and complete it together with the proof of the same direction $\text{RevRes}(\text{polylog})$ as they are similar.

We characterize the $\text{RevRes}(\text{polylog})$ by a $\text{TF}\Sigma_2$ variant of the *Sink-of-Potential-Line* (SOPL) problem. This is a *metered* variant of SOL, meaning that edges must always point towards larger numbers. An instance of SOPL is given by functions $S, P : [m] \rightarrow [m] \cup \{\text{undefined}\}$ which defines a graph G with edges (u, v) iff $S(u) = v$ and $P(v) = u$. A solution is either i) 1 if 1 is not a source in G , ii) a sink u in G , iii) a vertex which points backwards $S(u) < u$, or iv) a vertex u if $S(u)$ or $P(u)$ is undefined.

Definition 6.11. An instance of SOPL_2 is given by functions $S, P : [m] \times [t] \rightarrow [m]$. A solution is a witness to a solution to the SOPL instance defined by the meta-pointers (\tilde{S}, \tilde{P}) :

- (u, i, i') if $S(u, i') \neq u$ and $S(u, i) \neq S(u, i')$ or $P(u, i), P(u, i') \neq u$ and $P(u, i) \neq P(u, i')$.
(Predecessor or Successor of u is undefined)
- $(1, i)$ if $S(1, i) = 1$ or $\tilde{S}(1) = v \neq 1$ and $P(v, i) \neq 1$.
(1 is not a source)
- (u, i) if $S(u, i) < u$.
(u points backwards)

- (u, v, i) for $u < v$ if $\tilde{S}(u) = v, \tilde{P}(v) = u$ and $S(v, i) = v$; or $\tilde{S}(u) = v, \tilde{P}(v) = u, \tilde{S}(v) = w$ and $P(w, i) \neq v$.
(v is a proper sink)

Theorem 6.12. *For any $\text{FF}_F \in \text{TF}\Sigma_2$, there is a complexity- c SOPL_2 -formulation of FF_F iff there is a complexity $O(c)$ Σ_2 -RevRes(polylog) proof of F .*

This theorem follows by combining [Lemma 6.13](#) and [Lemma 6.14](#). We begin with the backwards direction, showing that RevRes(polylog) can prove SOPL_2 formulations. SOPL_2 is encoded as an unsatisfiable formula which is the conjunction of the following

- $\llbracket S_{u,i} \neq 1 \rrbracket$ for $u \in [m], i \in [t]$, and $\llbracket \tilde{S}_1 \neq v \rrbracket \vee \llbracket P_{u,i} \neq v \rrbracket$ for all $u, v \neq 1, i \in [t]$. (1 is a source)
- $\llbracket S_{u,i} \neq v \rrbracket \vee \llbracket S_{u,i'} \neq v' \rrbracket$ for all $i \neq i', v \neq v'$. (\tilde{S} is not undefined)
- $\llbracket P_{u,i} \neq v \rrbracket \vee \llbracket P_{u,i'} \neq v' \rrbracket$ for all $i \neq i', v \neq v'$. (\tilde{P} is not undefined)
- $\llbracket S_{u,i} \neq v \rrbracket$ for all $i \in [t]$ and $v < u$. (No backwards edges)
- Let $\bar{E}_{u,v} := \llbracket \tilde{S}_u \neq v \rrbracket \vee \llbracket \tilde{P}_v \neq u \rrbracket$, we include (No proper sinks)
 - $\bar{E}_{u,v} \vee \llbracket S_{v,i} \neq v \rrbracket$ for each $u < v$ and $j \in [m]$, and
 - $\bar{E}_{u,v} \vee \llbracket \tilde{S}_v \neq w \rrbracket \vee \llbracket P_{w,k} \neq w' \rrbracket$ for $u < v < w$ and $w \neq w'$ and $k \in [t]$.

Lemma 6.13. *For $\text{FF}_F \in \text{TF}\Sigma_2$, there is a complexity-RevRes(polylog)(F) SOPL_2 -formulation of FF_F .*

Proof. By [Lemma 6.2](#) it suffices to show that RevRes(polylog) can prove SOPL_2 . We will prove by induction on $u = m \dots 1$ that u does not have any outgoing edges. That is, we will derive the set of formulas

$$L_u := \{ \bar{E}_{u,v} : u > v \}.$$

First observe that the base case is given by the *no backwards edges* axioms. Assuming that we can derive L_1 , we show how to complete the proof. For $v > 1$, starting with $b = n$ down to $b = 1$, we cut

$$\frac{\llbracket \tilde{S}_1 \neq v \rrbracket \vee \bigvee_{j < b+1} \llbracket P_{v,j} \neq 1 \rrbracket \quad \llbracket \tilde{S}_1 \neq v \rrbracket \vee \llbracket P_{v,b} \neq w \rrbracket, \forall w \neq 1}{\llbracket \tilde{S}_1 \neq v \rrbracket \vee \bigvee_{j < b} \llbracket P_{v,j} \neq 1 \rrbracket}.$$

Next, starting from $a = n$ down to $a = 2$, we successively cut

$$\frac{\bigvee_{i < a+1} \llbracket S_{1,i} \neq v \rrbracket \quad \llbracket S_{1,a} \neq 1 \rrbracket \quad \llbracket S_{1,1} \neq v \rrbracket \vee \llbracket S_{1,a} \neq w \rrbracket, \forall w \neq 1, v}{\bigvee_{i < a} \llbracket S_{1,i} \neq v \rrbracket}.$$

Once all those formulae are derived, we cut one final time to finish the proof

$$\frac{\llbracket S_{1,1} \neq 1 \rrbracket \quad \llbracket S_{1,1} \neq v \rrbracket, \forall v > 1}{\perp}.$$

We now describe how to derive L_u from all L_v with $v > u$. For a given v and $\bar{E}_{v,w} \in L_v$, we start by weakening it to get $\llbracket \tilde{P}_v \neq u \rrbracket \vee \bar{E}_{v,w}$ and again to get $\bar{E}_{u,v} \vee \bar{E}_{v,w}$. Once this is done, starting at $c = n$ down to $k = 1$, we cut

$$\frac{\bar{E}_{u,v} \vee \llbracket \tilde{S}_v \neq w \rrbracket \vee \bigvee_{k < c+1} \llbracket P_{w,k} \neq v \rrbracket \quad \bar{E}_{u,v} \vee \llbracket \tilde{S}_v \neq w \rrbracket \vee \llbracket P_{v,c} \neq w' \rrbracket, \forall w' \neq w}{\bar{E}_{u,v} \vee \llbracket \tilde{S}_v \neq w \rrbracket \vee \bigvee_{k < c} \llbracket P_{w,k} \neq v \rrbracket}$$

to get $\bar{E}_{u,v} \vee \llbracket \tilde{S}_v \neq w \rrbracket$. Finally, from $b = n$ down to $b = 2$, we cut

$$\frac{\bar{E}_{u,v} \vee \bigvee_{j < b+1} \llbracket S_{v,j} \neq w \rrbracket \quad \bar{E}_{u,v} \vee \llbracket S_{v,b} \neq v \rrbracket \quad \llbracket S_{v,1} \neq w \rrbracket \vee \llbracket S_{v,c} \neq w \rrbracket, \forall w' \neq v, w}{\bar{E}_{u,v} \vee \bigvee_{j < b} \llbracket S_{v,j} \neq w \rrbracket}$$

and once we derived $\bar{E}_{u,v} \vee \llbracket S_{v,1} \neq w \rrbracket$ for each $w > v$, we have one final cut

$$\frac{\bar{E}_{u,v} \vee \llbracket S_{v,1} \neq w \rrbracket, \forall w > v \quad \bar{E}_{u,v} \vee \llbracket S_{v,1} \neq v \rrbracket \quad \llbracket S_{v,1} \neq w \rrbracket, \forall w < v}{\bar{E}_{u,v}}$$

to get $\bar{E}_{u,v} \in L_u$. □

Finally, we prove the other direction of [Theorem 6.12](#) and [Theorem 6.9](#).

Lemma 6.14. *Let $\text{FF}_F \in \text{TF}\Sigma_2$. Suppose that F admits a complexity- c Σ_2 -uCircRes(polylog) ($-\Sigma_2$ -RevRes(polylog)) proof, then there is a complexity- $O(c)$ SOL_2 -(SOPL_2 -)formulation of FF_F .*

Proof. We first handle Circular DNF resolution, and discuss what needs to be changed in order to handle Reversible DNF resolution at the end of the proof. The idea for the transformation of a uCircRes(polylog) proof into a SOL_2 formulation is the same as the transformation of a Res(polylog) into an ITER_2 formulation ([Lemma 6.6](#)) with the addition of defining a predecessor function. Let $\Pi = (D_1, \dots, D_m)$ be such a proof. By padding, we may assume that each DNF in the proof has the same number of terms. Let us consider the proof in reverse order such that $D_1 = \bigvee_{i \in [t]} \perp$.

Let $t_{u,i}$ be the i^{th} term of D_u . Given an assignment $\alpha \in \{0, 1\}^n$ to the variables of F , we construct a function $S_\alpha : [m] \times [t] \rightarrow [m]$ by setting $S_\alpha(u, i)$ to be:

- u if D_u is an axiom, or if $t_{u,i}(\alpha) = 1$;
- v if $t_{u,i}(\alpha) = 0$ and D_u was derived from D_v by the reverse cut rule or semantic weakening of an axiom;
- v if $t_{u,i}(\alpha) = 0$ and D_u was derived from $D_v = D_u \vee t$ and $D_w = D_u \vee \bar{t}$ via symmetric cut and $t(\alpha) = 0$ and w if $\bar{t}(\alpha) = 0$;

As well, define the predecessor function $P_\alpha : [m] \times [t] \rightarrow [m]$, as $P_\alpha(u, i)$:

- u if either $u = 1$, or the formula D_u was deduced but never used as the premise of a rule, or if $t_{u,i}(\alpha) = 1$;
- v if $t_{u,i}(\alpha) = 0$ and u is used as a premise to derive D_v via any of the rules but the reverse cut;
- v or w if $t_{u,i}(\alpha) = 0$ and D_u was used as the premise of the reverse cut rule to derive $D_v = D_u \vee t$ and $D_w = D_u \vee \bar{t}$. If $t(\alpha) = 0$, then $P_\alpha(u, i) = v$ and $P_\alpha(u, i) = w$ otherwise.

Finally, for each solution o to the instance S_α we define the output of the reduction $g_o(\alpha)$ to be arbitrary if o does not correspond to an axiom A_i of H , and otherwise this axiom A_i is a weakening of a DNF F_j of F , and we set $g_o(\alpha) = j$. Note that in this case $A_i(\alpha) = 0 \implies F_j(\alpha) = 0$. Observe that computing $S_\alpha(u, i)$ and $P_\alpha(u, i)$ involve evaluating at most two terms, and hence the reduction is efficient.

It remains to argue that the reduction is correct.

Claim. The following hold:

1. \tilde{P}_α and \tilde{S}_α are defined everywhere;
2. If D_u was used as the premise of a rule, $D_u(\alpha) = 0$ if and only if $\tilde{P}_\alpha(u) \neq u$ and $\tilde{S}_\alpha(u) \neq u$;
3. If $\tilde{P}_\alpha(u) = v \neq u$, then $D_v(\alpha) = 0$;
4. For a pair $u \neq v$, $\tilde{S}_\alpha(u) = v$ if and only if $\tilde{P}_\alpha(v) = u$.

Assuming the claim, the only solutions are proper sinks corresponding to falsified axioms of H , which are weakenings of (falsified) axioms of F . Hence, g returns a correct solution to $\text{FF}_F(\alpha)$.

Proof of Claim. The proof of this claim is, at heart, the same as the proof of the claim in [Lemma 6.6](#). The behavior of both functions implies that the only solutions one might get in the instance are proper sinks and that these proper sinks can only be falsified axioms.

Finally, when Π is a RevRes(polylog) proof, $S_\alpha(u, i) \geq u$ and $P_\alpha(v, j) \leq v$ for any u and v since the graph representation of Π does not include cycles, and thus we would not have *fake solutions* corresponding to edges pointing backwards making our formulation a valid SOPL_2 -formulation. \square

6.3 Relationships in $\text{TF}\Sigma_2$

In this subsection, we prove all the new inclusions in [1](#), relating the classes whose combinatorial principle is *artificially* brought up from TFNP to some naturally $\text{TF}\Sigma_2$ classes.

Here the characterization helps us build the reduction.

Proposition 6.15. $(\text{USOD})_n$ admits an efficient $\text{uCircRes}(\text{polylog}(n))$ -proof and $\text{STRONGAVOID} \in \text{SOL}_2$.

Proof. Let us first give the explicit encoding of the $(\text{USOD})_n$ by giving the axioms:

1. **1 is a sink:** This is encoded by two axioms.

(a) $[S_1 = 1];$

(b) $\bigvee_{t \neq 1} [S_t = 1];$

2. **No sources:** $\bigvee_t [S_t = u].$

The strategy for the proof is:

1. assume that $S(u) = u$ for any $u \neq 1$;
2. from the fact that $S(v) = v$ for all $v \neq u$, deduce that $S(u) = u$. Indeed if all other nodes point to themselves, u can not point to anything but itself since otherwise it would qualify as a source. We also derive $S(u) \neq 1$ during this process;
3. once this is done, we will be left with the fact that $S(u) \neq 1$ for each $u \neq 1$ which is in direct contradiction with the second axiom.

We start by introducing $[S_u = u]$ for each $u \neq 1$ via the DNF creation rule for step 1.

Now, fixing u , for $t \neq u$, let us weaken $[S_t = t]$ to obtain $[S_t = t] \vee [S_u \neq w]$ for all $w \in [n]$ and consider the case $w = t$. Since $u \neq t$, the formula $[S_t \neq t] \vee [S_u \neq t]$ is a tautology and thus we can introduce it. Then we may operate the cut

$$\frac{[S_t = t] \vee [S_t \neq u] \quad [S_t \neq t] \vee [S_t \neq u]}{[S_t \neq u]}$$

and get $[S_t \neq u]$ for each $t \neq u$. Once this is done, we may cut

$$\frac{\bigvee_t [S_t = u] \quad [S_t \neq u], \forall t \neq u}{[S_u = u]}$$

Observe that, as announced, we have derived $[S_t \neq 1]$ for all $t \neq 1$. That is the end for step 2.

We can now simply cut one last time to finish the proof

$$\frac{\bigvee_{t \neq 1} [S_t = 1] \quad [S_t \neq 1], \forall t \neq 1}{\perp}$$

The size of the proof and the characterization theorem shows that $\text{USOD} \in \text{SOL}_2$. Also, the equivalence $\text{USOD} =_{dt} \text{STRONGAVOID}$ gives us $\text{STRONGAVOID} \in \text{SOL}_2$ \square

The *sink-of-DAG* problem is the canonical PLS-complete problem in which one is given a source of a DAG and one wants to find a sink. Our characterization of Sherali-Adams by STRONGAVOID proceeded via a (equivalent) *unmetered source-of-DAG* problem. Hence, it is natural to also consider a metered version of these problem, where one is given a sink of a DAG and one wants to find a source.

Definition 6.16. The *Source of DAG* (SOD) problem is defined as follows. The input is a “successor” function $S : [n] \rightarrow [n]$ which defines a graph in which each vertex has fan-out ≤ 1 but arbitrary fan-in. There is an edge from i to j if $S(i) = j$. A solution to the instance S is:

1. i if $S(i) < i$; (i has a backward edge)
2. n if for all $i < n$, $S(i) \neq n$; (n is not a sink)
3. i if for all $j \in [n]$, $S(j) \neq i$. (A source)

The following result then arise naturally.

Proposition 6.17. SOD_n admits an efficient $\text{RevRes}(\text{polylog}(n))$ -proof, and $\text{SOD} \in \text{SoPL}_2$.

Proof. The axioms of SOD_n are:

1. **n is a proper sink:** $\bigvee_{t \neq n} [S_t = n];$

2. **No sources:** $\bigvee_t \llbracket S_t = u \rrbracket$ for each $u \neq 1$;
3. **No edges pointing backwards:** $\llbracket S_u \neq v \rrbracket$ for any pair of nodes $v < u$.

The strategy of the proof is as follows:

1. given that $S(t) = t$ for each $t < u$, deduce that $S(u) = u$. This must be true since otherwise u is a source.
2. use the fact that the derived formulae directly contradicts the first axiom.

Let us start by describing step 1. Assume we have derived $\llbracket S_t = t \rrbracket$ for each $t < u$. Then we may weaken these formulae to get $\llbracket S_t = t \rrbracket \vee \llbracket S_t \neq v \rrbracket$ and let us consider the case $v = u$. Since $t \neq u$, the formula $\llbracket S_t \neq t \rrbracket \vee \llbracket S_t \neq u \rrbracket$ is a tautology that we introduce, and we cut

$$\frac{\llbracket S_t \neq t \rrbracket \vee \llbracket S_t \neq u \rrbracket \quad \llbracket S_t = t \rrbracket \vee \llbracket S_t \neq u \rrbracket}{\llbracket S_t \neq u \rrbracket}$$

to obtain $\llbracket S_t \neq u \rrbracket$. Next, we cut

$$\frac{\bigvee_t \llbracket S_t = u \rrbracket \quad \llbracket S_t \neq u \rrbracket, \forall t < u \quad \llbracket S_t \neq u \rrbracket, \forall t > u}{\llbracket S_u = u \rrbracket}$$

to derive $\llbracket S_u = u \rrbracket$. Once this is done, we may perform one final cut

$$\frac{\bigvee_{t \neq n} \llbracket S_t = n \rrbracket \quad \llbracket S_t \neq n \rrbracket, \forall t \neq n}{\perp}$$

hence $\text{SoD} \in \text{SOPL}_2$. □

These proofs are interesting since they do indicate that up to complexifying a function, it is possible to build an inverse that is also hard-to-compute with an efficient reduction. Also, since we know how to transform $\text{uCirRes}(\text{polylog})$ -refutations (resp. $\text{RevRes}(\text{polylog})$ -refutations) into SOL_2 -instances (resp. SOPL_2 -instances), following the instructions lets us concretely build those inverses.

The other inclusions are proved more directly, and do not rely on the characterization.

Proposition 6.18. $\text{LOP} \leq_{dt} \text{ITER}_2$.

Proof. Let \prec be an LOP instance on $[n]$. By encoding it with $\binom{n}{2}$ variables such that, for $i < j \in [n]$, $x_{i,j} = 1$ means $i \prec j$, and $x_{i,j} = 0$ means $j \prec i$, we can force the purported order to always be total. An output to the LOP instance would thus either be a \prec -minimal element, or a proof that \prec is not an order, i.e., that the transitivity does not hold. Consider the ITER_2 instance on $\binom{n}{2} + n$ meta-nodes with a meta-node for each $(i, j) \in [n]^2$ with $i \geq j$. Let $(1, 1)$ be the source. It helps to think of the meta-nodes as arranged in n levels, with the first element in the label being the level a meta-node is at.

The idea is that (i, j) is valid (i.e., has an outgoing edge) if and only if \prec is transitive and j is the \prec -minimal value in $[i]$. If $i < n$, it will point to $(i + 1, j')$, where $j' = j$ if j is still \prec -minimal in $[i + 1]$, and $j' = i + 1$ otherwise.

We now formally define the nodes with index (i, j) . If $i = n$, then it contains a single node that points to itself. Otherwise, there is two kind of nodes:

- $\binom{n}{3}$ nodes verifying the transitivity of \prec . Each of those nodes are associated with 3 distinct elements $(a, b, c) \in [n]^3$. We define $S((i, j), (a, b, c))$ as follows:
 - Query $a \prec b, b \prec c$ and $a \prec c$. If the answers show that \prec is not transitive on (a, b, c) , point to (i, j) .
 - Query $j \prec i + 1$. If it holds, point to $(i + 1, j)$. Otherwise, point to $(i + 1, i + 1)$.
- $i - 1$ nodes verifying the validity of (i, j) . Each of those nodes are associated with a value $k \in [i] \setminus \{j\}$. We define $S((i, j), k)$ as follows:
 - Query $j \prec k$. If it does not hold, point to (i, j) .
 - Query $j \prec i + 1$. If it holds, point to $(i + 1, j)$, otherwise, point to $(i + 1, i + 1)$.

Since every node that points out of its index does the same query to decide where to point, the meta successor is well-defined. If \prec is not transitive, every meta-node will point to itself. The solution can thus only be of type $((1, 1), i)$, with this node being of the transitive type. This immediately gives us a triple in $[n]$ proving \prec is not transitive. If \prec is indeed a total order, then it is clear that every level has a single active node, the only proper sink on level n indicates the \prec -minimal value in $[n]$. \square

Proposition 6.19. $\text{SOD} \leq_{dt} \text{LOP}$.

Proof. Let S be an SOD instance on n vertices. Consider an LOP instance \prec on $2n$ values split into two groups $C = [n]$ and $L = [n]$, we denote elements of C by i_C and elements of L by i_L , for $i \in [n]$. The group C 's goal is to "check for backward pointers", if the \prec -minimal element is i_C , then i points backwards. The group L checks for loops: if the \prec -minimal element is i_L , then there are no backward edges. Moreover, if $i = n$, then n is not a proper sink. Otherwise, i is the first node (in regular order) to not point to itself in S , i.e., i is a source. Formally, for $i, j \in [n]$, we define \prec as follows:

- $i_C \prec j_C$ if and only if $i < j$;
- $i_C \prec j_L$ if and only if $S(i) < i$;
- $i_L \prec j_L$ if and only if either one of the following holds:
 - $S(i_L) = i, S(j_L) = j$ and $i > j$;
 - $S(i_L) \neq i$ and $S(j_L) = j$;
 - $S(i_L) \neq i, S(j_L) \neq j$ and $i < j$.

It is clear that \prec is total. If it is transitive, then the minimal element is either the first source in S , or n if it is not a proper sink. If it is not transitive, the minimal element allows us to find a backward pointer. \square

Theorem 1 in [KKMP21] proves that $\text{FNP} \subseteq \text{PEPP}$, we prove that actually $\text{FNP} \subseteq \text{SOD}$. As it is straightforward that SOURCEOFDAG reduces to $\text{UNMETEREDSOURCEOFDAG}$, which is equivalent to empty, this implies that every $\text{TF}\Sigma_2$ class studied in this paper, apart from APEPP , contains FNP .

Proposition 6.20. $\text{FNP} \subseteq \text{SOD}$.

Proof. Let x be an instance of P_n , a FNP problem, and let \mathcal{O} be its set of solutions. By definition of FNP , this set is at most quasipolynomial in n . Consider the PLS instance with $|\mathcal{O}| + 1$ nodes. Consider the extra node as n . To define $S(o)$, run the verification decision tree of P_n for (x, o) . If it accepts, point to n , otherwise, point $S(o)$ points to itself. If n is not a proper sink, it means P_x has no solution on x , if o is a proper source, then it is a solution for x . \square

7 $\text{TF}\Sigma_{d+1}$ Consequences of Depth- $d.5$ Frege

The PK proof system and its fragments PK_k whose derived formulae are limited to depth- d formulae is one of the most well-known and studied proof system in proof complexity. One particular use for these proof systems is that they correspond to the theories T_2^d [Kra94, Kra01, ST11] of bounded arithmetic introduced by Buss in his seminal thesis [Bus86]. These theories, together with the theories S_2^d form a hierarchy akin to the polynomial hierarchy found in complexity theory - with the caveat that $S_2^d \subseteq T_2^d$ instead of the double helix usually drawn to describe PH. As for the latter questions of whether this hierarchy collapses, and if so at what level, is an important, if not the most important, open problem in the field. Multiple pathways have been followed in tackling this problem, but one that stuck out is looking at the $\forall\Sigma_1$ consequences of such said theories. The characterization results for T_2^1 by the PLS class of problems in [BB09a] and T_2^2 by the class C-PLS - a version of PLS whose nodes are each assigned a set of colors - in [KST07b] translate to the characterization of the TFNP classes PLS and C-PLS by the propositional proof systems Res and $\text{Res}(\text{polylog})$ [DR23] respectively in the field of propositional proof complexity.

In the following section, we prove a general characterization theorem for each of the proof systems $\text{Res}(C_d)$ - where C_d is the class of formulae $\Sigma_{d.5} \cup \Pi_{d.5}$ - with Σ_{d+1} -weakening. The main idea for the complete problem is more or less the same as in the case of ITER_2 . We aim at obfuscating a function corresponding to an ITER instance that needs the power of cutting over $\Sigma_{d.5} \cup \Pi_{d.5}$ formulae. As usual, the Σ_{d+1} -semantic weakening of axioms is mainly part of the definition to ensure that the class defined by the proof system is closed under efficient decision tree

reduction. What differs really from the level 2 case is how we obfuscate. Here the trick of asking for all indices to agree does not work anymore - or at least not to our best efforts. That is where the problems GPLS_d and PE_d from [PT12] inspired us with the idea of alternating min's and max's. The fact that these problems also already characterize the $\forall\Sigma_1$ consequences of T_2^d leads us to believe that this the right way to proceed.

We start by defining our proof system $\text{Res}(C_d)$ subsection 7.1 and prove that it is closed under decision tree substitution. We then go on with defining our problem ITER_d and the corresponding class PLS_d^{dt} in subsection 7.2. We also prove that our problem is indeed in $\text{TF}\Sigma_d^{dt}$. Finally, in subsection 7.3 we attack the characterization result. We apply the template we have already profusely used to prove characterizations:

1. prove that we can transform efficient proof into efficient decision tree reductions to the complete problem in Proposition 7.11;
2. find an efficient refutation of the formula corresponding to the complete problem in Proposition 7.14.

For the second part, we will use a popular method used to find refutations of some proof systems like Res: The Prover-Delayer game.

7.1 The $\text{Res}(C_d)$ proof system

Writing $C_d = \Sigma_{d,5} \cup \Pi_{d,5}$, the $\text{Res}(C_d)$ proof system follows a simple intuition: in proof complexity, it is usual to consider CNF formulae and resolve over opposite literals, and CNFs are basically conjunctions of disjunctions of literals which we denote by $\Pi\Sigma\text{L}$ with L being the class of literals. When dealing with higher order search problems, as seen previously, we see that their translation does not yield CNFs anymore, but unsatisfiable $\Pi_{(d+1),5} = \Pi\Sigma C_d$ -formulae and this motivates the considering the following proof system.

Definition 7.1. For $d \geq 1$, a $\text{Res}(C_d)$ refutation of a Π_{d+2} -unsatisfiable formula $F = \bigwedge_{i=1}^m A_i$ is a sequence of $\text{polylog}(n)$ -width Σ_d -formulae $\Pi = (\pi_1, \dots, \pi_l = \perp)$ where each π_i is deduced from the previous Σ_d -formulae by one of the following rules:

- *Axiom Introduction.* Introduce A_i for some $i \in [m]$.
- *C_d -Cut* From $\pi \vee C$ and $\sigma \vee \bar{C}$ derive $\pi \vee \sigma$ where C is any C_d -formula.
- *C_d -weakening.* From π derive $\pi \vee C$ where C is any C_d -formula.

The size s of Π is $\sum_{i=1}^s |\pi_i|$ and its width w is the maximum width of any formula in Π . The complexity of Π is $\log(s) + w$.

We usually see a $\text{Res}(C_d)$ proof in the form of a DAG where each node corresponds to a formula of the proof and has zero, one, or two parents depending on what rule was used to derive it. Looking at proof this way leads to a following characterization very close to what we already have for the resolution proof system: *The Prover-Delayer game*.

Pascal the Prover and Danielle the Delayer have a disagreement over a particular $\Pi\Sigma C$ -formula $F = \bigwedge_{i=1}^m A_i$. Pascal says they are convinced that the formula F is unsatisfiable, while Danielle insists on the contrary. Danielle goes even so far as asserting that they have in their possession an assignment α satisfying the formula F , but they refuse to give it to Pascal. Pascal, not being born yesterday, trusts their intuition, and says they can not simply believe Danielle at their word and need more information to verify such a claim. An argument ensues, but both of our protagonists ultimately find an agreement. Danielle will not give the value of α to Pascal, however they agree on answering the following type of queries: Pascal may choose a formula $C \in C_d$ and Danielle answers with the alleged evaluation $C(\alpha)$. The game then proceeds as follows: Pascal chooses a formula $C \in C_d$, Danielle answers with $b \in \{0, 1\}$, and Pascal remembers the equality $C(\alpha) = b$. The game ends when the set of equalities remembered by Pascal $\{C_i(\alpha) = b_i\}_{i \in I}$ is incompatible with all axioms A_i being satisfied.

Definition 7.2. Let $F = \bigwedge_{i=1}^m A_i$ be an unsatisfiable Π_{d+2} -formula. A strategy for the Prover is the data of a DAG G of maximum fan-out two and one root. Each node n is labeled with a set of boolean equalities of C_d -formulae M_n which we call the *memory of the Prover at node n* . Given $M_n = \{C_i = 0\}_{i \in I}$, we write $V(M_n) := \{\alpha \mid C_i(\alpha) = 0, \forall i \in I\}$. The memories relate to each other in the following way:

- if $n = r$ is the root, then $M_r = \emptyset$;

- if n has one child c , then $M_c = M_n \setminus \{C = 0\}$ for some $C \in \mathcal{C}$;
- if n has two children nodes c_0 and c_1 , then $M_{c_0} = M_n \cup \{C = 0\}$ and $M_{c_1} = M_n \cup \{\bar{C} = 0\}$ for some $C \in \mathcal{C}$.
- if l is a leaf, then there exists $i \in [m]$ such that $A_i(\alpha) = 0$ for each $\alpha \in V(M_l)$.

The size of the strategy is the $\sum_{n \in V(G)} \sum_{C \in M_n} |C|$ and its width is the maximal width of any \mathcal{C}_d -formula queried.

As for the original Prover-Delayer game for the resolution proof system and the Buss-Pudlák game for bounded-depth Frege proofs, finding a strategy for a formula closely relates to finding a refutation.

Lemma 7.3. *Let F be a Π_{d+2} unsatisfiable formula. If F admits a $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ refutation of size s and width w , then there is a strategy for the Prover for F of size s and width w and conversely.*

Proof. Given a $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ proof of F (H, Π), the strategy is given by the following.

- The Prover starts at the end of the proof Π , that is on the node corresponding to \perp and its memory is empty.
- At the node corresponding to the formula π , if it was derived via:
 - the weakening rule, i.e. $\pi = \pi' \vee C$ for some $C \in \mathcal{C}_d$: the Prover forgets the equality $\{C = 0\}$;
 - the cut rule over $C \in \mathcal{C}_d$ from $\pi_0 = A \vee C$ and $\pi_1 = B \vee \bar{C}$: Then the Prover queries the formula C . When the Delayer answers with response $b \in \{0, 1\}$, then the Prover adds the equality $C = 0$ to its memory if $b = 0$ and moves to the node labeled with π_0 , and else they add the equality $\bar{C} = 0$ and moves to the node labeled with π_1 .
 - the axiom introduction rule: the Prover stops

Moving this way through the graph of the proof ensures that the set of satisfying assignments for the current memory is incompatible with the formula labeling the node. Thus, when the Prover stops, the current memory is incompatible with either an axiom of H .

Now suppose we have a strategy for the Prover and consider its graph. Then at each node n replace the memory $M_n = \{C_1 = 0, \dots, C_{m_n} = 0\}$ with the formula $\bigvee_{i \in [m_n]} C_i$. This is indeed a derivation as querying a formula $C \in \mathcal{C}$ a formula $C \in \mathcal{C}$ ends being an instance of the cut rule and the weakening of the memory corresponds to an instance of the weakening rule. Moreover, the leafs are labeled with weakenings of axioms of F . We may then take our weakening H to be the conjunction of the formulae labeling the leaves. \square

To make sure that $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ defines indeed a TFP class, we need the proof system to be closed under decision tree reductions, since otherwise the mere notion of the class does not apply to the proof system. We check here that $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ verifies this property soon, but first let us consider the following lemma about \mathcal{C}_d -formulae.

Lemma 7.4. *For $d \geq 1$, $f = (f_1, \dots, f_n)$ a vector of $\text{polylog}(n)$ -depth decision trees and $C \in \mathcal{C}_d$ a formula on n variables. Then $C(f) \in \mathcal{C}_d$.*

We can now prove that $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ verifies the wanted property.

Lemma 7.5. *Let F be an unsatisfiable Π_{d+2} -formula on n variables and let $f = (f_1, \dots, f_n)$ be a vector of $\text{polylog}(n)$ -depth decision trees. Then if F admits a $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ refutation of size s , $F(f)$ admits a refutation of size $s \text{polylog}(n)$.*

Proof. We prove this result using the characterization of proof by the Prover-Delayer game. Say we have a strategy of size s . The strategy is then modified in the following way: replace the memory of a node v $M_v = \{C_1 = 0, \dots, C_{m_v} = 0\}$ by $M_{v,f} = \{C_1(f) = 0, \dots, C_{m_v}(f) = 0\}$ which is a valid strategy by Lemma 7.4. What this transformation amounts to is that instead of querying C , now the Prover will query the formula $C(f)$ which represents a correct strategy for the Prover for formula $F(f)$. \square

Using this result, we are able to prove that the $\Sigma\mathcal{C}\text{-Res}(\mathcal{C})$ proof system is closed under decision tree reductions. This part is important if we want to be able to make sense of the class.

Lemma 7.6. *Let $F = \bigwedge_{i=1}^m F_i$ and $G = \bigwedge_{j=1}^l G_j$ be unsatisfiable Π_{d+2} -formulae on n variables and suppose that we have a proof of F of size s . If there is a decision tree reduction (f, g) from FF_G to FF_F of $\text{polylog}(n)$ -depth, then G admits a $\Sigma_{d+2}\text{-Res}(\mathcal{C}_d)$ proof of size $s 2^{\text{polylog}(n)}$.*

Proof. Consider the proof of F as a strategy for the Prover. Our strategy for G amounts to the following: Consider the strategy for $F(f)$ described in the previous result. Then, at any leaf corresponding to the weakening of an axiom $F_i(f)$ for $i \in [m]$, the Prover queries variables down the decision tree g_i . Once this is over, they consider the label of the path j associated to the path in g_i that the Prover went down to. By the correctness of the reduction, the current memory of the Prover is such that the set of satisfying assignment for said memory contains the set of satisfying assignments for the axiom G_j of G . Taking the formula being the co \square

7.2 The ITER_d problem and the PLS_d subclass

Definition 7.7. For an integer $k \geq 1$, and a product set $\mathbf{r} = [r_1] \times \cdots \times [r_k]$ we set the following notation when k is odd:

$$\begin{aligned} \text{MAX}(\mathbf{r}) &:= \max_{i_1 \in [r_1]} \min_{i_2 \in [r_2]} \cdots \max_{i_d \in [r_k]}, \\ \text{MIN}(\mathbf{r}) &:= \min_{i_1 \in [r_1]} \max_{i_2 \in [r_2]} \cdots \min_{i_d \in [r_k]}, \end{aligned}$$

and if k is even we only change the last min or max to its opposite.

As for the ITER_2 problem, the an instance of ITER_d problem for $k \geq 3$ is an obfuscated instance of ITER . The main idea of ITER_d is to render the successor function S computable only trough the help of a $\Sigma_{d,5}$ oracle by an alternation of max and min. Indeed, what ITER_d amounts to is to give a function $S : [n]^k \rightarrow [n]$ and construct an ITER instance by setting $\tilde{S}(u) = \text{MIN}([n]^{k-1})\{S(u, i)\}$ for a node u . This closely resembles the approach used in [KT21] for the definition of the RWPHP_2 problem in $\text{TF}\Sigma_2$ from the classical retraction pigeonhole principle in TFNP . What we prove in the rest of this section is a generalization of the classical result $\text{Res}^{dt} = \text{PLS}^{dt}$ by showing that if we define PLS_d to be the class of problems efficiently reducible to ITER_d then $\text{PLS}_{d+2}^{dt} = \Sigma_{d+1} - \text{Res}(\text{C}_d)^{dt}$ for all $k \geq 1$. We start by giving a formal definition of ITER_d .

Definition 7.8. An instance of $\text{ITER}_{d,m,t}$ with $t = (t_1, \dots, t_{d-1})$ is given by a successor function $S : [m] \times [t_1] \times \cdots \times [t_{d-1}] \rightarrow [m]$ that describes a graph on m vertices via the \tilde{S} function the same way we have done so far. An output is given by a quadruple (u, i, v, j) with $u, v \in [m]$ and $i_1^*, j_1^* \in [t_1]$ such that

$$\begin{aligned} i_1^* &= \arg \min \{ \max_{i_1 \in [t_1]} \cdots \min_{i_{d-1} \in [t_{d-1}]} S(u, i) \} \\ j_1^* &= \arg \min \{ \max_{j_2 \in [t_2]} \cdots \min_{j_{d-1} \in [t_{d-1}]} S(v, j) \} \end{aligned}$$

when k is even and the last min is replaced by a max when k is odd, and $\tilde{S}(u) = v$, and

- $u = v = 1$ (1 is not a source);
- $v < u$ (u admits a backward pointer);
- $u < v$ and $\tilde{S}(u) = \tilde{S}(v) = v$ (v is a proper sink).

The problem ITER_d corresponds to the case where all functions are equal to the identity function. The class PLS_d is defined as the syntactic class of problems that admit an efficient decision tree reduction to ITER_d .

One way to see the indices i_1^* and j_1^* in the solution is as *certificates of computation for u and v* . Once again, this follows closely the approach in [KT21] where a solution of to a RWPHP_2 instance is not only a solution to the usual underlying RWPHP instance but also a certificate of computation for the functions involved. One reason this problem is hard is that for the solutions where $u < v$, the verifier must be able to assert whether $\tilde{S}(u) = v$ and $\tilde{S}(v) = v$ or, in other words, it must be able to assert that the certificates of computation i and j indeed witness a correct computation for their respective input nodes.

Proposition 7.9. $\text{ITER}_d \in \text{TF}\Sigma_d$ for $d \geq 1$.

Proof. Let us assume k is even, the odd case being handled in the same way up to changing a min into a max. Consider an output $o = (u, i_1^*, v, j_1^*)$ and an instance S . First, let us understand what the assertion $\tilde{S}(u) = v$ logically means.

In fact, this statement can be written as the conjunction of the assertions $\tilde{S}(u) \geq v$ and $\tilde{S}(u) \leq v$. Both statement may be rewritten as

$$\begin{aligned}\tilde{S}(u) \geq v &\equiv \forall i_1 \exists i_2 \cdots \forall i_{d-1} S(u, i) \geq v \\ \tilde{S}(u) \leq v &\equiv \exists i_1 \forall i_2 \cdots \exists i_{d-1} S(u, i) \leq v\end{aligned}$$

In particular the latter formula is true by taking i_1 to be the argument of the minimum. Thus, for the output o , the verifier $V_o(S, i', i, j', j)$ for $i' = (i_1, \dots, i'_{d-1})$, $j' = (j'_1, \dots, j'_{d-1})$ and $i = (i_2, \dots, i_{d-1})$, $j = (j_2, \dots, j_{d-1})$ follows this procedure:

1. checks that $S(u, i') \geq v$ and $S(u, i_1^*, i) \leq v$, and outputs 0 otherwise;
2. outputs 1 if $u < v$ or $u = v = 1$;
3. checks that $S(v, j') \geq v$ and $S(v, j_1^*, j) \leq v$, and outputs 1 if it case and 0 otherwise.

From the paragraph above, we then see that the sentence

$$\forall(i'_1, i_2, j'_1, j_2) \exists(i'_2, i_3, j'_2, j_3) \cdots \exists(i'_{k-2}, i_{d-1}, j'_{k-2}, j_{d-1}) \forall(i'_{d-1}, j'_{d-1}) V_o(S, i', i, j', j)$$

is true if and only if S admits o as an output. □

7.3 $\text{RES}(C_d)$ Characterizes PLS_{d+2}

Now we are ready to state a general characterization theorem.

Theorem 7.10. *For any $\text{FF}_F \in \text{TF}\Sigma_{d+2}$ and d , there is a complexity- c PLS_{d+2}^{dt} -formulation of FF_F iff there is a complexity $\Theta(c)$ $\Sigma_{d+2}\text{-Res}(C_d)$ proof of F .*

Proof. Proposition 7.14 and Lemma 7.6 gives the inclusion $\text{PLS}_{d+2}^{dt} \subseteq \Sigma_{d+1}\text{-Res}(C_d)^{dt}$. Proposition 7.11 gives us the converse $\Sigma_{d+2}\text{-Res}(C_d)^{dt} \subseteq \text{PLS}_{d+2}^{dt}$. □

We start by proving PLS_{d+2}^{dt}

Proposition 7.11. *Let F be an unsatisfiable $\Pi_{(k+1),5}$ formula on n variables. Suppose F admits a $\Sigma_{d+2}\text{-Res}(C_d)$ -refutation of size s and width w , then there are functions $m, t_1, \dots, t_d, t_{d+1} : \mathbb{N} \rightarrow \mathbb{N}$ such that $m \cdot t_1 \cdots t_{d+1} = s$ such that FF_F admits a $(\text{ITER}_{d+2,m,t})$ -formulation of size s and depth w .*

Before going about the proof of this result, let us talk a little bit about $\Pi_{d,5}$ and $\Sigma_{d,5}$ formulae via the following lemma.

Lemma 7.12. *Let $F \in \Sigma_{d,5}$ on n variables and let us write*

$$F = \bigvee_{i_1 \in [r_1]} \bigwedge_{i_2 \in [r_2]} \cdots \bigcirc_{i_d \in [r_k]} c_i$$

with the notation $i = (i_1, \dots, i_d)$. Let us write $\mathbf{r} = [r_1] \times \cdots \times [r_k]$ and let c_i be the corresponding C_0 -subformula. Then for $\alpha \in \{0, 1\}^n$, we have that the evaluation of F at α is equal to

$$F(\alpha) = \text{MAX}(\mathbf{r})\{c_i(\alpha)\}.$$

With the same notation, if $F \in \Pi_{d,5}$, then we have

$$F(\alpha) = \text{MIN}(\mathbf{r})\{c_i(\alpha)\}.$$

Proof. For $k = 0, 1$, this is clear. This is realized by induction for $k \geq 2$, since, writing $F = \bigvee_{i_1 \in [r_1]} \bigwedge_{i_2 \in [r_2]} F_{i_1, i_2}$ with $F_{i_1, i_2} \in \Sigma_{d-2}$ for each $(i_1, i_2) \in [r_1] \times [r_2]$. Then $F(\alpha) = \max_{i_1 \in [r_1]} \min_{i_2 \in [r_2]} \{F_{i_1, i_2}(\alpha)\}$. Since $F_{i_1, i_2} \in \Sigma_{d-2}$, we get that $F_{i_1, i_2}(\alpha) = M([r_3] \times \cdots \times [r_k])\{c_{i_1, \dots, i_d}(\alpha)\}$. Thus

$$F(\alpha) = \max_{i_1 \in [r_1]} \min_{i_2 \in [r_2]} \{\text{MAX}([r_3] \times \cdots \times [r_k])\{c_{i_1, \dots, i_d}(\alpha)\}\} = \text{MAX}([r_1] \times \cdots \times [r_k])\{c_i(\alpha)\}.$$

The proof of the second assertion is practically the same. □

We are now ready to prove the proposition.

Proof. Let $\Pi = (\pi_1, \dots, \pi_m)$ be a proof of size s . We start by reverting the order of the proof such that $\pi_1 = \perp$. Also, if $\pi_u = A \vee B$ was derived from $\pi_v = A \vee \bar{C}$ and $\pi_w = B \vee C$ by an application of the cut rule with $C \in \Pi_{d,5}$, let us assume that we have that $u < v < w$ up to reordering the proof. Also, let us replace $\pi_u = A \vee B$ with $\pi_u = A \vee B \vee C$. Up to padding the subformulae with \top or \perp depending on the main connectives, we may assume that each π_u appearing in the proof is of the form

$$\pi_u = \bigvee_{i_1 \in [t_1]} \bigwedge_{i_2 \in [t_2]} \cdots \bigcirc_{i_{d+1} \in [t_{d+1}]} c_{u,i}$$

with the notation $i = (i_1, \dots, i_{d+1})$ and $c_{u,i} \in \mathbf{C}_0$. For any assignment $\alpha \in \{0, 1\}^n$, let us define the following successor function S_α .

1. $S_\alpha(u, i) = u$ if π_u is an axiom;
2. $S_\alpha(u, i) = c_{u,i}(\alpha) \cdot u + (1 - c_{u,i}(\alpha)) \cdot v$ if π_u was derived from π_v by the weakening rule;
3. $S_\alpha(u, i) = c_{u,i}(\alpha) \cdot u + (1 - c_{u,i}(\alpha)) \cdot w$ if $\pi_u = A \vee B \vee C$ was derived from $\pi_v = A \vee \bar{C}$ and $\pi_w = B \vee C$ with $C \in \Pi_{d,5}$ and $c_{u,i}$ is a \mathbf{C}_0 -subformula of A or B ;
4. $S_\alpha(u, i) = c_{u,i}(\alpha) \cdot v + (1 - c_{u,i}(\alpha)) \cdot w$ if $\pi_u = A \vee B \vee C$ was derived from $\pi_v = A \vee \bar{C}$ and $\pi_w = B \vee C$ with $C \in \Pi_{d,5}$ and $c_{u,i}$ is a \mathbf{C}_0 -subformula of C ;

where the multiplications are symbolic. The output function g_o for any $o = (u, i, v, j)$ is given by the constant function returning v . The size of the reduction is indeed the size of the proof and since each computation amounts to evaluating a term or clause of width w at most the depth is also the width of the proof.

We now need to prove that the reduction is correct. Consider the following claim.

Claim 7.13. The function \tilde{S}_α has the following properties:

1. $\tilde{S}_\alpha(u) \geq u$;
2. if $u \in [m]$ is such that π_u is not an axiom, then $\tilde{S}_\alpha(u) = u$ if and only if $\pi_u(\alpha) = 1$ (where here we go back to $\pi_u = A \vee B$ rather than $\pi_u = A \vee B \vee C$);
3. for $u \in [m]$, if $\tilde{S}_\alpha(u) = v$ with $v \neq u$, then $\pi_v(\alpha) = 0$

Assuming the claim to be right, we see that it concludes the proof. Indeed, the claim implies that the only proper sinks in the graph described by \tilde{S} are falsified axioms. Then the output function g_o simply outputs the sink, i.e. a falsified axiom. \square

We now only need to prove the claim.

Proof of Claim 7.13. Observe this about $\tilde{S}_\alpha(u)$.

1. if π_u is an axiom, then $\tilde{S}_\alpha(u) = u$;
2. if π_u was derived from π_v using the weakening rule, then

$$\begin{aligned} \tilde{S}_\alpha(u) &= \text{MIN}(k+1, i) \{c_{u,i}(\alpha) \cdot u + (1 - t_{u,i}(\alpha)) \cdot v\} \\ &= \text{MIN}(k+1, i) \{c_{u,i}(\alpha)(u - v) + v\} \\ &= \text{MAX}(k+1, i) \{c_{u,i}(\alpha)\}(u - v) + v \\ &= \pi_u(\alpha) \cdot u + (1 - \pi_u(\alpha)) \cdot v \end{aligned}$$

with MIN changing to MAX because the quantity $u - v$ is negative and the last equality is given by [Lemma 7.12](#).

3. if $\pi_u = A \vee B \vee C$ was derived from $\pi_v = A \vee \bar{C}$ and $\pi_w = B \vee C$, writing $\mathbf{t}_{A \vee B} = [t_1 - 1] \times [t_2] \times \cdots [t_{d+1}]$ and $\mathbf{t}_C = [t_2] \times \cdots \times [t_{d+1}]$

$$\begin{aligned} \tilde{S}_\alpha(u) &= \min[\text{MIN}(\mathbf{t}_{A \vee B}) \{c_{u,i}(\alpha) \cdot u + (1 - c_{u,i}(\alpha)) \cdot w\}, \text{MAX}(\mathbf{t}_C) \{c_{u,t_1,i}(\alpha) \cdot v + (1 - c_{u,t_1,i}(\alpha)) \cdot w\}] \\ &= \min[\text{MAX}(\mathbf{t}_{A \vee B}) \{c_{u,i}(\alpha)\} \cdot (u - w) + w, \text{MIN}(\mathbf{t}_C) \{c_{u,t_1,i}(\alpha)\} \cdot (v - w) + w] \\ &= \min[(A \vee B)(\alpha) \cdot u + (1 - (A \vee B)(\alpha)) \cdot w, C(\alpha) \cdot v + (1 - C(\alpha)) \cdot w] \\ &= (A \vee B)(\alpha) \cdot u + (1 - (A \vee B)(\alpha))(C(\alpha) \cdot v + (1 - C(\alpha)) \cdot w) \end{aligned}$$

with MIN switching to MAX (resp. MAX switching to MIN) because $(u - w)$ (resp. $(v - w)$) are negative quantities and the third equality is given by [Lemma 7.12](#).

The three properties follow from these equalities. \square

As stated above $\Sigma_{d+2}\text{-Res}(C_d)^{dt}$ is closed under efficient decision tree reductions. In our situation to prove our theorem it is sufficient to prove the following.

Proposition 7.14. *ITER_{d+2} admits a $\Sigma_{d+2}\text{-Res}(C_d)$ proof of size $O(n^{k+6})$.*

Before moving to the proof, let us have a word about the encoding of the formula. For each node $u \in [n]$ and vector of indices $i \in [n]^{k+1}$, as in the case of ITER₂, we have $\log(n)$ many variable $S_{u,i,\alpha}$ for each $\alpha = 0, \dots, \log(n) - 1$. The idea is to encode the value of $S(u, i)$ as the conjunction of these variables or their negation. For example, if v has binary expansion $v = \sum_{\alpha=0}^{\log(n)-1} b_\alpha 2^\alpha$, then the formula $\llbracket S_{u,i} = v \rrbracket := \bigwedge_{\alpha=0}^{\log(n)-1} S_{u,i,\alpha}^{b_\alpha}$ with the notation that $x^1 = x$ and $x^0 = \bar{x}$ for any variable x . Also, as in the case of ITER₂, we write $\llbracket S_{u,i} \neq j \rrbracket$ for the negation of $\llbracket S_{u,i} = j \rrbracket$ where the negation has been propagated, i.e. the conjunction becomes a disjunction and all literals are negated.

We now describe some formulae that will come in handy in the proof below. Some of them depends on the parity of k so we may need to give two different definitions. To ease the notation, we simply state that all variables take value in $[n]$ if not stated explicitly otherwise. We also adopt the symbolic notation $i' = (i'_1, \dots, i'_{k+1})$, and $i'' = (i''_1, \dots, i''_{k+1})$ and analogously by replacing i with j . The formulae are:

1. For two node u and v , and index i_1^* , $\llbracket i_1^* \neq \arg \min \vee \tilde{S}_u \neq v \rrbracket$ is the formula given by

$$\begin{aligned} & \left(\bigvee_{i'_1} \bigwedge_{i'_2} \cdots \bigwedge_{i'_{k+1}} \bigwedge_{v \leq v'} \llbracket S_{u,i'} \neq v' \rrbracket \right) \vee \left(\bigvee_{i''_2} \bigwedge_{i''_3} \cdots \bigvee_{i''_{d+1}} \bigvee_{v < v''} \llbracket S_{u,i_1^*,i''} = v'' \rrbracket \right) & \text{if } k \text{ is odd;} \\ & \left(\bigvee_{i'_1} \bigwedge_{i'_2} \cdots \bigvee_{i'_{d+1}} \bigvee_{v' < v} \llbracket S_{u,i'} = v' \rrbracket \right) \vee \left(\bigvee_{i''_2} \bigwedge_{i''_3} \cdots \bigwedge_{i''_{d+1}} \bigwedge_{v'' \geq v} \llbracket S_{u,i_1^*,i''} \neq v'' \rrbracket \right) & \text{if } k \text{ is even.} \end{aligned}$$

It encodes the fact that $\tilde{S}(u)$ is not v by arguing that either i_1^* is not a certificate of computation for u (left hand side of the disjunction) or that the minimum is greater than v (right hand side of the disjunction). These are $\Sigma_{(k+1).5}$ -formulae that will be used for the axioms of ITER_{d+2}.

2. For a two nodes u and v and an index i_1 , $\llbracket S_{u,i_1} \leq v \rrbracket$, with the convention that $i = (i_1, \dots, i_{d+1})$ is the formula given by

$$\begin{aligned} & \bigwedge_{i_2} \bigvee_{i_3} \cdots \bigwedge_{i_{d+1}} \bigwedge_{v < v'} \llbracket S_{u,i} \neq v' \rrbracket & \text{if } k \text{ is odd;} \\ & \bigwedge_{i_2} \bigvee_{i_3} \cdots \bigvee_{i_{d+1}} \bigvee_{v' \leq v} \llbracket S_{u,i} = v' \rrbracket & \text{if } k \text{ is even.} \end{aligned}$$

It encodes the fact that $\max_{i_2} \min_{i_3} \cdots \circ_{i_{d+1}} S(u, i)$ is less or equal to v . These are $\Pi_{d.5}$ -formulae that will be used as cuts in the proof (or queries in strategy for the Prover).

We are now finally ready to describe the $\Sigma_{d+2}\text{-Res}(C_d)$ -refutation of the ITER_{d+2}-formula.

Proof. First, let us describe the axioms of the formula:

1. **1 is a source:** $\llbracket i^* \neq \arg \min \vee \tilde{S}_1 \neq 1 \rrbracket$ for each i^* ;
2. **No backwards pointer:** $\llbracket i^* \neq \arg \min \vee \tilde{S}_u \neq v \rrbracket$ for each pair of nodes $v < u$ and index i^* ;
3. **No proper sinks:** For each pair of nodes $u, v \in [n]$ with $u < v$ and indices i_1^*, j_1^*

$$\llbracket i_1^* \neq \arg \min \vee \tilde{S}_u \neq v \rrbracket \vee \llbracket j_1^* \neq \arg \min \vee \tilde{S}_v \neq v \rrbracket$$

Instead of giving an explicit refutation, we give a strategy for the Prover for the formula formed from the axioms above. The Delayer claims they have (the binary encoding of) a function $S : [n]^{k+1} \rightarrow [n]$ violating none of the axioms above. This is indeed impossible by the totality of the problem and the goal of the Prover is then to query formulae about the behavior of the function until they notice that its behavior being incompatible with one of the axioms. The idea of the strategy is as follows The Prover navigates on the graph described by the function \tilde{S} . Starting at the node 1, they reach a goal if and only if they have confirmation from the Delayer that they witness a solution. In other words, when they have confirmation that:

- 1 points to itself;
- u points to a lesser node;
- $\tilde{S}(u) = v$ and $\tilde{S}(v) = v$ with $u < v$.

To achieve this, the Prover, starting at node $u = 1$, tries to guess $\tilde{S}(u)$. They are not able to ask this directly to the Delayer since this would require the Prover to query a formula of higher alternating depth. Finding the value of $\tilde{S}(u)$ is done through the *auction procedure* described below. Once they get the value $\tilde{S}(u) = v$, if $v < u$ or if $u = v = 1$, the Prover stops since it witnesses a violation of the formula. Otherwise, if $u < v$, the Prover keeps in memory the fact - here we mean the set of equalities implying said fact - that $\tilde{S}(u) = v$, forgets about anything else and moves on with searching for the value of $\tilde{S}(v)$ via the auction procedure. Then if they find that $\tilde{S}(v) \leq v$, they as they witness a solution in this case as well. Otherwise, if $\tilde{S}(v) = w$ with $v < w$ they forget that $\tilde{S}(u) = v$ and goes on with parcouring the rest of the graph until they witnesses a solution.

In order to set out the strategy, we first need to describe the so called *auction procedure*. Given the node $u \in [n]$, the goal of the procedure is to determine the value v such that $\tilde{S}(u) = v$ is the only value compatible with the answers given by the Delayer. The procedure happens in rounds, starting with round $v = n - 1$ down to round $v = 1$. The Prover queries the formula $\llbracket S_{u,i_1} \leq v \rrbracket$ starting from $i_1 = 1$ up to $i_1 = n$ and reacts the following way to answers of the Delayer:

1. At round $v = n - 1$: The memory of the Prover is either empty - that is in the case $u = 1$ - or contains equalities implying $\tilde{S}(t) = u$ for $t < u$. We ignore this part of the memory as it does not affect what the Prover does in the procedure.
 - (a) as soon as the Delayer answers 1 for some $i_1^* \in [n]$, the Delayer forgets all the previous equalities of the form $\llbracket S_{u,1} \leq n - 1 \rrbracket = 0, \dots, \llbracket S_{u,i_1^*-1} \leq n - 1 \rrbracket = 0$, keeps $\llbracket S_{u,i_1^*} \leq n - 1 \rrbracket = 1$ and moves on to round $v = n - 2$;
 - (b) if the Delayer has answered 0 for all i_1 , then the Prover knows that $\tilde{S}(u) = n$. They then stop the procedure keeping their memory as is.
2. At round $1 \leq v < n - 1$: The memory of the Prover is only comprised of $\llbracket S_{u,i_1^{**}} \leq v + 1 \rrbracket = 1$ for some i_1^{**} . With this memory configuration the Prover knows that $\tilde{S}(u) \leq v + 1$. At this stage, the Prover now wants to ensure whether $\tilde{S}(u) = v + 1$ or $\tilde{S}(u) \leq v$ and goes about this way to distinguish between the two cases by querying the formulae described above.
 - (a) if at some point the Delayer answers $\llbracket S_{u,i_1^*} \leq v \rrbracket = 1$, then the Prover forgets the equalities $\llbracket S_{u,1} \leq v \rrbracket = 0, \dots, \llbracket S_{u,i_1^*-1} \leq v \rrbracket = 0$ as well as $\llbracket S_{u,i_1^{**}} \leq v + 1 \rrbracket = 1$, only keeping $\llbracket S_{u,i_1^*} \leq v \rrbracket = 1$ and moves on to round $v - 1$ if $v > 1$ or stops if $v = 1$ since this implies $\tilde{S}(u) = 1$;
 - (b) if the Delayer has answered 0 for all i_1 , then the Prover memory of the prover should contain the equalities $\llbracket S_{u,i_1^{**}} \leq v + 1 \rrbracket = 1$ for some i_1^{**} and $\llbracket S_{u,i_1} \leq v \rrbracket = 0$ for each i_1 . This clearly encodes the fact that $\tilde{S}(u) = v$ and the Prover stops the auction procedure keeping its memory as is.

We are now ready to describe the strategy in details. Starting with $u = 1$, it goes as follows. The Prover, via the auction procedure described above, finds the node v such that $\tilde{S}(u) = v$ and then proceeds by:

- if $u = v = 1$, the Prover keeps their memory as is and stops;
- if $v < u$, the Prover forgets everything in their memory but the equalities corresponding to $\tilde{S}(u) = v$ and stops;
- if $v = u \neq 1$, then the Prover, prior to computing $\tilde{S}(u)$, already had equalities corresponding to the fact that $\tilde{S}(t) = u$ for some $t < u$. They then stop;
- if $u < v$, then if the Prover had in memory equalities corresponding to $\tilde{S}(t) = u$ for some $t < u$ - which would be the case when $u \neq 1$ - they forget those equalities and repeat the process replacing u by v .

The phase is certain to terminate since u increments by 1 after every step and by the time $u = n$, either the Delayer answers with $\tilde{S}(n) = n$ or $\tilde{S}(n) < n$. Before moving on to the description of the certificate phase, let us count the number of states the Prover can be in. Note that for any pair of nodes $u < v$, the knowledge $\llbracket \tilde{S}(u) = v \rrbracket$ corresponds to any of the following sets of equalities $\{\llbracket S_{u,i_1^*} \leq v+1 \rrbracket = 1\} \cup \{\llbracket S_{u,i_1} \leq v \rrbracket = 0\}_{i_1 \in [n]}$ for any $i_1^* \in [n]$. This represents n possibilities for each such pair of nodes.

- During the auction procedure to compute $\tilde{S}(1)$:
 - $\{\llbracket S_{1,1} \leq n-1 \rrbracket = 0, \dots, \llbracket S_{1,i_1-1} \leq n-1 \rrbracket = 0, \llbracket S_{1,i_1} \leq n-1 \rrbracket = b\}$ for some $i_1 = 0, \dots, n$ and $b \in \{0, 1\}$ which amounts to $2(n+1)$ many possible states;
 - $\{\llbracket S_{1,i_1^*} \leq v+1 \rrbracket = 1, \llbracket S_{1,1} \leq v \rrbracket = 0, \dots, \llbracket S_{1,i_1-1} \leq v \rrbracket = 0, \llbracket S_{1,i_1} \leq v \rrbracket = b\}$ for some $v \in [n-1]$, $i_1^* \in [n]$, $i_1 = 0, \dots, n$ and $b \in \{0, 1\}$ which amounts to $2n(n^2-1)$ many possible states;
 - $\{\llbracket S_{1,i_1^*} \leq 1 \rrbracket = 1\}$ for some $i_1^* \in [n]$ which amounts n possible states.
- with the knowledge $\llbracket \tilde{S}(u) = v \rrbracket$, during the auction procedure to compute $\tilde{S}(v)$:
 - $\{\llbracket S_{1,1} \leq n-1 \rrbracket = 0, \dots, \llbracket S_{1,i_1-1} \leq n-1 \rrbracket = 0, \llbracket S_{1,i_1} \leq n-1 \rrbracket = b\}$ for some $i_1 = 0, \dots, n$. This amounts to $2n(n+1)$ (accounting for the knowledge) possible states for any pair, hence $2n^3(n+1)$ possible states accounting for all possible pairs.
 - $\{\llbracket S_{1,i_1^*} \leq w+1 \rrbracket = 1, \llbracket S_{1,1} \leq w \rrbracket = 0, \dots, \llbracket S_{1,i_1-1} \leq w \rrbracket = 0, \llbracket S_{1,i_1} \leq w \rrbracket = b\}$ for some $w \in [n-1]$, $i_1^* \in [n]$, $i_1 = 0, \dots, n$ and $b \in \{0, 1\}$ which amounts to $2n(n^2-1)$ many possible states for any pair accounting for the knowledge, hence $2n^3(n^2-1)$ many possible states accounting for all possible pairs.
 - $\{\llbracket S_{v,i_1^*} \leq 1 \rrbracket\}$ for some $i_1^* \in [n]$ which amounts to n^2 possible states for any pair accounting for the knowledge, hence n^4 many possible states accounting for all possible pairs.

So we see that the strategy has $O(n^5)$ many possible states for the sink phase. Since each memory state contains at most $2(n+1)$ equalities of the form $\llbracket S_{u,i_1} \leq v \rrbracket = b$ for some nodes u, v , index i_1 and bit b , and the size of such formulae is $O(n^k)$, we get that the size of the strategy is $O(n^{k+6})$.

We now only need to prove that the strategy is valid, i.e. when the Prover stops, at least one axiom of the formula is incompatible with their memory. Let us then consider the different cases:

- $u = v = 1$: then the memory of the Prover contains $\llbracket S_{1,i_1^*} \leq 1 \rrbracket = 1$ for some i_1^* . This is incompatible with the axiom $\llbracket i_1^* \neq \arg \min \vee \tilde{S}_1 \neq 1 \rrbracket$;
- $v < u$: there are two possibilities for the memory of the Prover
 - if $v = 1$: then the memory contains an equality $\llbracket S_{u,i_1^*} \leq 1 \rrbracket = 1$ for some i_1^* which contradicts the axiom $\llbracket i_1^* \neq \arg \min \vee \tilde{S}(u) \neq 1 \rrbracket$;
 - if $v \neq 1$: then the memory contains $\llbracket S_{u,i_1} \leq v-1 \rrbracket = 0$ for all i_1 and $\llbracket S_{u,i_1^*} \leq v \rrbracket = 1$ for some i_1^* . This contradicts the axiom $\llbracket i_1^* \neq \arg \min \vee \tilde{S}_u \neq v \rrbracket$;
- $\tilde{S}(u) = \tilde{S}(v) = v \neq 1$: there are also two possibilities for such a case for the memory of the Prover.
 - if $v \neq n$: then the memory contains the equalities $\llbracket S_{u,i_1^*} \leq v \rrbracket = 1$ and $\llbracket S_{v,j_1^*} \leq v \rrbracket$ for some i_1^*, j_1^* along with all the equalities $\llbracket S_{u,i_1} \leq v-1 \rrbracket = 0$ and $\llbracket S_{v,j_1} \leq v-1 \rrbracket = 0$ for all i_1, j_1 . This contradicts the axiom

$$\llbracket i_1^* \neq \arg \min \vee \tilde{S}(u) \neq v \rrbracket \vee \llbracket j_1^* \neq \arg \min \vee \tilde{S}(v) \neq v \rrbracket.$$
 - if $v = n$: then the memory contains the equalities $\llbracket S_{u,i_1} \leq n-1 \rrbracket = 0$ and $\llbracket S_{v,j_1} \leq n-1 \rrbracket = 0$ for all i_1, j_1 . This contradicts the axioms

$$\llbracket i_1^* \neq \arg \min \vee \tilde{S}_u \neq n \rrbracket \vee \llbracket j_1^* \neq \arg \min \vee \tilde{S}_n \neq n \rrbracket$$

for all possible values of i_1^* and j_1^* . Indeed for any node u , if $\tilde{S}(u) = n$ then i_1^* will always be an argument of the minimum whatever its value.

We can then conclude that we have valid strategy for the Prover of size $O(n^{k+6})$. \square

8 Acknowledgments

Noah Fleming was supported by NSERC. The authors thank Robert Robere, Toniann Pitassi, and Oliver Korten for some helpful discussions.

References

- [AL23] Albert Atserias and Massimo Lauria. Circular (yet sound) proofs in propositional logic. *ACM Trans. Comput. Log.*, 24(3):20:1–20:26, 2023.
- [BB09a] Arnold Beckmann and Samuel R Buss. Characterising definable search problems in bounded arithmetic via proof notations. *Ways of proof theory*, pages 65–134, 2009.
- [BB09b] Arnold Beckmann and Samuel R. Buss. Polynomial local search in the polynomial hierarchy and witnessing in fragments of bounded arithmetic. *Journal of Mathematical Logic*, 09(01):103–138, 2009.
- [BCE⁺98] Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The relative complexity of NP search problems. *J. Comput. Syst. Sci.*, 57(1):3–19, 1998.
- [BFI23] Sam Buss, Noah Fleming, and Russell Impagliazzo. TFNP characterizations of proof systems and monotone circuits. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPIcs*, pages 30:1–30:40. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [BKT14] Samuel R. Buss, Leszek Aleksander Kolodziejczyk, and Neil Thapen. Fragments of approximate counting. *J. Symb. Log.*, 79(2):496–525, 2014.
- [Bus86] Samuel R Buss. *Bounded arithmetic*, volume 86. Bibliopolis Napoli, 1986.
- [CHR24] Lijie Chen, Shuichi Hirahara, and Hanlin Ren. Symmetric exponential time requires near-maximum circuit size. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 1990–1999. ACM, 2024.
- [DR23] Ben Davis and Robert Robere. Colourful TFNP and propositional proofs. In Amnon Ta-Shma, editor, *38th Computational Complexity Conference, CCC 2023, July 17-20, 2023, Warwick, UK*, volume 264 of *LIPIcs*, pages 36:1–36:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [FGPR24] Noah Fleming, Stefan Grosser, Toniann Pitassi, and Robert Robere. Black-box PPP is not turing-closed. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 1405–1414. ACM, 2024.
- [FKP19] Noah Fleming, Pravesh Kothari, and Toniann Pitassi. Semialgebraic proofs and efficient algorithm design. *Found. Trends Theor. Comput. Sci.*, 14(1-2):1–221, 2019.
- [GHJ⁺22a] Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, and Ran Tao. Further collapses in TFNP. *Electron. Colloquium Comput. Complex.*, TR22-018, 2022.
- [GHJ⁺22b] Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, and Ran Tao. Separations in proof complexity and TFNP. *CoRR*, abs/2205.02168, 2022.
- [GKRS19] Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 38:1–38:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [KKMP21] Robert Kleinberg, Oliver Korten, Daniel Mitropolsky, and Christos H. Papadimitriou. Total functions in the polynomial hierarchy. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPIcs*, pages 44:1–44:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [Kor21] Oliver Korten. The hardest explicit construction. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 433–444. IEEE, 2021.

- [KP24] Oliver Korten and Toniann Pitassi. Strong vs. weak range avoidance and the linear ordering principle. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 1388–1407. IEEE, 2024.
- [Kra94] Jan Krajíček. Lower bounds to the size of constant-depth propositional proofs. *The Journal of Symbolic Logic*, 59(1):73–86, 1994.
- [Kra01] Jan Krajíček. On the weak pigeonhole principle. *Fundamenta Mathematicae*, 170(1-2):123–140, 2001.
- [KST07a] Jan Krajíček, Alan Skelley, and Neil Thapen. NP search problems in low fragments of bounded arithmetic. *J. Symb. Log.*, 72(2):649–672, 2007.
- [KST07b] Jan Krajíček, Alan Skelley, and Neil Thapen. NP search problems in low fragments of bounded arithmetic. *The Journal of Symbolic Logic*, 72(2):649–672, 2007.
- [KT21] Leszek Aleksander Kolodziejczyk and Neil Thapen. Approximate counting and NP search problems. 2021.
- [Li24] Zeyong Li. Symmetric exponential time requires near-maximum circuit size: Simplified, truly uniform. In Bojan Mohar, Igor Shinkar, and Ryan O’Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 2000–2007. ACM, 2024.
- [LPR24] Yuhao Li, William Pires, and Robert Robere. Intersection classes in TFNP and proof complexity. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference, ITCS 2024, January 30 to February 2, 2024, Berkeley, CA, USA*, volume 287 of *LIPIcs*, pages 74:1–74:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.
- [PT12] Pavel Pudlák and Neil Thapen. Alternating minima and maxima, nash equilibria and bounded arithmetic. *Annals of Pure and Applied Logic*, 72:604–614, 2012.
- [Pud15] Pavel Pudlák. On the complexity of finding falsifying assignments for herbrand disjunctions. *Arch. Math. Log.*, 54(7-8):769–783, 2015.
- [ST11] Alan Skelley and Neil Thapen. The provably total search problems of bounded arithmetic. *Proceedings of the London Mathematical Society*, 103(1):106–138, 2011.
- [Tha24] Neil Thapen. How to fit large complexity classes into TFNP. 2024.