

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

On Limiting & Limited Non-determinism in NEXP Lower Bounds

Permalink

<https://escholarship.org/uc/item/2w9797pg>

Author

Dhayal, Anant

Publication Date

2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

On Limiting & Limited Non-determinism in NEXP Lower Bounds

A dissertation submitted in partial satisfaction of the
requirements for the degree

Doctor of Philosophy

in

Computer Science

by

Anant Dhayal

Committee in charge:

Professor Russell Impagliazzo, Chair

Professor Samuel Buss

Professor Sanjoy Dasgupta

Professor Ramamohan Paturi

Professor Jacques Verstraete

2021

Copyright

Anant Dhayal, 2021

All rights reserved.

The dissertation of Anant Dhayal is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2021

TABLE OF CONTENTS

	Dissertation Approval Page	iii
	Table of Contents	iv
	List of Tables	vii
	Acknowledgements	viii
	Vita	ix
	Abstract of the Dissertation	x
1	Introduction	1
	1.1 Fast (unambiguous) algorithms imply UEXP circuit lower bounds	4
	1.2 UEXP lower bounds are constructive, useful, and unique	7
	1.3 Gradually increasing the non-determinism in circuits for NEXP lower bounds: with a hope to prove $NEXP \not\subseteq P/poly$	10
	1.4 Our techniques, interesting by-products, and previous work	14
	1.5 Organization of the thesis	22
2	Preliminaries	23
3	EWL and KLT for UTIME and related classes	32
	3.1 Search to decision reduction for UTIME	32
	3.2 EWL and KLT for UTIME	34
	3.3 EWL and KLT for ZUTIME	35
	3.4 EWL and KLT for $\widetilde{FewTIME}$	37
4	UEXP Lower Bounds from Fast Unambiguous Algorithms	39

4.1	Lower bounds from unambiguous tautology and canonization algorithms	40
4.2	Lower bounds from unambiguous Π_2 SAT algorithms	42
4.3	Lower bounds from $\widetilde{\text{Few}}$ tautology algorithms	42
4.4	Lower bounds from unambiguous learning and tautology algorithms	43
4.5	Generalization of lower bound frameworks	44
5	Unique Properties vs. Lower Bounds	47
5.1	ZUE lower bounds vs P-U properties	49
5.2	UE lower bounds vs $P/\log n - u_{\leq 1}$ properties	51
5.3	UE/ n lower bounds vs $P/\log n$ -U properties	53
5.4	ZNE lower bounds vs NP-U properties	54
5.5	ZNE lower bounds vs NP-N properties	56
5.6	NE lower bounds vs NP/ $\log n$ -U properties	58
6	Derandomization Using Unique Properties	60
7	Isolation of properties: EWL & KLT for ZNE	63
8	Lower bounds against prSV non-deterministic circuits	68
8.1	NEXP vs $(\text{NP} \cap \text{Co-NP})/\text{poly}$	68
8.2	NEXP $\not\subseteq (\text{NP} \cap \text{Co-NP})/\text{poly}$ from super-polynomial savings	71
8.3	New gap theorems for CAPP and MA	72
8.4	Fast algorithms imply lower bounds against circuits with limited prSV non-determinism	73
8.5	Uncodntional lower bounds against restricted circuits with limited prSV non-determinism	76
8.6	Unconditional fixed-polynomial lower bounds against unrestricted circuits that use prSV non-determinism	77

9	Conclusions and Open Problems	81
	Bibliography	83

LIST OF TABLES

Table 5.1: Properties vs Lower Bounds	47
-------------------------------------------------	----

ACKNOWLEDGEMENTS

I would like to thank my advisor Prof. Russell Impagliazzo for all his help and encouragement during my Ph.D. I would also like to thank Prof. Ramamohan Paturi and Prof. Jayalal Sarma for teaching me the nuances of research and all the help they provided during various research projects. I am very grateful to Prof. Sam Buss, Prof. Sanjoy Dasgupta, Prof. Jacques Verstraete and Prof. Pavel Pevzner to serve on my exam committees. A special thanks to Prof. Valentine Kabanets, Prof. Dieter van Melkbeek, Prof. Eric Allender and Prof. Ryan Williams for teaching me the basics of circuit lower bounds, and to Sasank Mouli, Vaibhav Krishan, Marco Carmosino, Sam McGuire, Sajin Koroth and Abhishek Bhruhundi for helpful discussions that made computational complexity a bit less complex. I would also like to thank other UCSD theory group students Jiapeng Zhang, Baiyu Li, Sankeerth Rao, Kaave Hosseini, Nick Genise, Jiawei Gao, Stefan Schneider, Ivan Mihajlin, Jessica Sorrell, Mark Schultz, Rex Lei, Max Hopkins, Michael Borkowski, Michael Walter, Daniel Moeller, Joseph Jaeger, Ken Hoover and Michael Jaber for uncountably infinite lunch time gossip.

I am very grateful to my parents Bhanwar Lal and Sunil, brother Anchal, sister-in-law Priyanka, and wife Radha who supported me throughout this Ph.D. program. My soccer teammates and other friends who deserve a special acknowledgement are, Aman Nougrihiya, Yash Gehlot, Saurabh Mogre, Suganth Krishna, Nishant Bhaskar, Shubham Saini, Mainak Biswas, Pulak Sarangi, Thyagarajan Venkatanarayanan, Suruj Deka, William Weston-Dawkes, Raquel Hakes, Geoff Ganzberger, Adit Bhardwaj, Siddhartha Agarwal, Varun Garg, Kshitiz Gupta, Vineet Pandey, Viraj Deshpande, Bharat Sridhar, Madhura Som, Animesh Gupta and Sahil Agarwal.

Chapters 1, 3-6, 8 and 9 contain material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

VITA

- 2012 B. Tech. in Computer Science & Engineering, LNM Institute of Information
Technology, Jaipur, India
- 2015 M. S. by Research in Computer Science, Indian Institute of Technology
Madras, Chennai, India
- 2021 Ph. D. in Computer Science, University of California San Diego

ABSTRACT OF THE DISSERTATION

On Limiting & Limited Non-determinism in NEXP Lower Bounds

by

Anant Dhayal

Doctor of Philosophy in Computer Science

University of California San Diego, 2021

Professor Russell Impagliazzo, Chair

Proving circuit lower bounds is one of the most difficult tasks in computational complexity theory. The NP vs. $P/poly$ problem asks whether there are small non-uniform circuits that can simulate circuit satisfiability. The answer is widely believed to be false, but so far progress has only been made in the case of restricted circuits. In 1980s the progress stalled after it was shown that NP doesn't have non-uniform AC^0 circuits that have MOD- m gates for any prime m . After almost three decades, in 2010s Williams made progress in the relaxed case of NEXP lower bounds. He first showed that non-trivial satisfiability algorithms for a circuit class entail NEXP lower bounds against

that class. Then he designed a fast satisfiability algorithm for ACC circuits (AC^0 circuits with MOD- m gates for any constant m) to show that NEXP doesn't have non-uniform ACC circuits.

We make progress in bringing down the class NEXP, specifically by **limiting non-determinism** (in terms of the number of non-deterministic branches that accept). We show that slightly faster satisfiability algorithms entail lower bounds for UEXP and related classes. We believe this is progress towards making similar connections, and thus proving lower bounds, for EXP and lower complexity classes.

To investigate why progress again stalled around ACC lower bounds, and why TC^0 (AC^0 circuits with majority gates) lower bounds have not been established yet, Williams made rigorous connections between NEXP lower bounds and variations of Natural Proofs. Razborov and Rudich defined Natural Proofs to showcase the limitations of the current lower bound techniques. They showed that any technique that entails Natural Proofs, i.e. Proofs that are (i) constructive, (ii) useful, and (iii) large, fail to prove strong lower bounds. Williams showed that NEXP lower bounds, regardless of the technique, entail Proofs that satisfy the first two of the three conditions of Natural Proofs.

We make Williams connections more rigorous, and show that UEXP lower bounds entail Proofs, that in addition to the first two conditions of Natural Proofs, satisfy a third condition that is exactly the opposite of largeness condition. We call this condition, the uniqueness condition, and these Proofs, the Unique Proofs. These connections showcase that $NEXP \supseteq UEXP \supseteq EXP$ is a viable path to approach EXP lower bounds.

We also discuss an alternate approach to improve NEXP lower bounds. We define a new form of non-determinism to capture the non-uniform circuits from the class $(NP \cap Co-NP)/poly$, and call it promise-Single-Valued non-determinism. We show that in the current NEXP lower bounds, we can allow the non-uniform circuits some **limited non-determinism** (in terms of the number of non-deterministic inputs) of the type promise-Single-Valued. We also discuss that, how small improvements in the amount of this special type of non-determinism, even in the restricted circuits much weaker than ACC, would imply very strong lower bounds such as $NEXP \not\subseteq P/poly$.

Chapter 1

Introduction

The two fundamental problems in theoretical computer science are: (a) design non-trivial algorithms for computational tasks; or (b) prove that such algorithms do not exist, i.e. prove (circuit) lower bounds for the computational task in question. For instance, the famous 1970s question of Cook and Levin, is $NP = P$ or $NP \neq P$ [25, 61]? Or its non-uniform version, is $NP \subset P/poly$ or $NP \not\subset P/poly$ (it would also imply $NP \neq P$)? $P/poly$ is the class of problems that have non-uniform polynomial-size circuits. Non-uniform computation allows the sizes of programs to grow with the sizes of inputs, and can be naturally represented as an infinite family of Boolean circuits (one for each possible input length). They are very powerful and can simulate all undecidable problems when there is no size restriction, and certain undecidable problems even with polynomial size restriction. But it's still not known if $P/poly$ can simulate NP . In fact the answer is believed to be false. One of the main reasons being the collapse of polynomial hierarchy. This is second to the $P = NP$ collapse, in the list of collapses that are widely believed to be false by complexity theorists.

No significant progress after decades of efforts lead to the pursuit of considerably relaxed versions of the above pair of questions. Here we focus on the lower bounds side of the question. Optimists started the bottom-up approach and started proving lower bounds for restricted circuit classes with a hope that gradually they would lift the restrictions over time.

If we just consider size restrictions, the best known lower bound for any NP problem is

$5n - o(n)$, for circuits over the basis of 2-bit AND, OR, and NOT gates [50, 60]. MA (with 1-bit of advice) is the lowest class that contains functions with super-linear and fixed-polynomial (size n^k for any constant k) lower bounds [80]. If we consider restrictions on gate types (basis), monotone circuits, i.e. circuits without negations or NOT gates are the most studied. A super-polynomial lower bound was proved for NP in [76] (for CLIQUE to be specific), and improved to exponential size in [5]. But later in [77] it was shown that these techniques for monotone circuit lower bounds would not extend to general circuits.

Here we only focus on the depth restrictions. Constant depth restrictions are the most studied ones. In 1980s it was shown that the parity function on n bits, which is in P (infact in even lower complexity classes), can't be computed by AC^0 circuits [35, 2]: the class of constant depth circuits over the basis of AND, OR, and NOT gates of arbitrary fan-in. Later the lower bound was improved to exponential size in [103], and eventually an optimal lower bound was established in [38]. Then the next question in line was about the power of AC^0 circuits with parity gates. Let $AC^0[\oplus_p]$ denote the class of AC^0 circuits with MOD- p gates. An exponential size lower bound was proved for $AC^0[\oplus_2]$ in [77] for the majority function on n bits. Majority too lies in P and lower classes. For any two primes $p \neq q$, an exponential size lower bound for the MOD- p function was proved for the $AC^0[\oplus_q]$ circuits in [87]. Based on these lower bounds, the logical next step was to move towards the following two classes that are more expressive: (i) ACC, the class of AC^0 circuits with MOD- m gates for arbitrary constant $m > 1$; (ii) TC^0 , the class of AC^0 circuits with majority (or equivalently, threshold) gates. Note that, TC^0 can simulate ACC.

Failing to prove that NP is not contained in polynomial-size non-uniform ACC circuits, the lower bound question was relaxed further. The next question asked was, whether this lower bound can be established for NEXP, or even EXP^{NP} . Note that, MAEXP is the smallest class known to have super-polynomial lower bounds for unrestricted Boolean circuits [13] (although, in [53] it was shown that $NEXP^{RP}$, that is contained in MAEXP, can't have polynomial-size circuits of both types, Boolean and arithmetic). MAEXP contains NEXP, and is incomparable to EXP^{NP} . Even for EXP^{NP} , the super-polynomial lower bound for ACC (even depth-3 $AC^0[\oplus_6]$) were elusive for about three decades.

In his seminal work in 2010s, Williams [100] showed super-polynomial ACC lower bounds for NEXP, and exponential ACC lower bounds for EXP^{NP} . He used connections between fast algorithms and lower bounds from his prior work [99]. For TC^0 lower bounds, MAEXP is still the smallest class we know.

Pessimists on the other hand started formulating barriers to show that all the techniques used in the bottom-up approach are not good enough to prove stronger lower bounds. There are three main barriers in the literature that any lower bound technique should overcome: (i) Relativization [10], (ii) Algebrization or Algebraic Relativization [1], and (iii) Natural Proofs barrier [78]. The first two barriers essentially say that techniques that relativize, i.e. work even in presence of arbitrary oracles, fail to prove most of the lower bounds. This is because, for most of the lower bound questions, there are some oracles relative to which the lower bound is known to hold, and there are other oracles relative to which the lower bound is not known to hold. In this thesis we will only focus on the third barrier.

The Natural Proofs barrier of Razborov and Rudich [78] argues that almost all known proofs of nonuniform circuit lower bounds, entail algorithms/properties that: (i) are efficient (*constructivity*), (ii) distinguish *hard* functions from *easy* by only accepting hard functions (*usefulness*), (iii) accept many hard functions (*largeness*). Here by hard (*resp.* easy) we mean that the function requires bigger (*resp.* smaller) circuits to compute, typically super-polynomial (*resp.* polynomial) in size. Any such algorithm would refute widely believed cryptography primitives, and thus Natural Proofs are self-limiting in the sense that: in order to prove weak lower bounds, they provide algorithms that refute strong lower bounds that are also believed to be true. Even the small class TC^0 supports cryptography [69, 59, 64], and to prove lower bounds against it we would need un-Natural techniques.

Main results: We focus on three different directions in this thesis: (i) We describe ways in which Williams fast algorithms to lower bounds connection from [99] can be extended to lower bounds for classes smaller than NEXP, by **limiting** the non-determinism used by NEXP (in terms of the number of non-deterministic branches that accept). This can be viewed as a progress towards establishing similar connections, and hence proving lower bounds, for EXP and lower complexity

classes. (ii) Then we discuss how lower bounds for these smaller classes evade the Natural Proofs barrier. (iii) Finally we extend the current NEXP lower bounds by allowing circuit classes a **limited** amount of non-determinism. We also discuss how small increments in our results would lead to much stronger lower bounds such as $\text{NEXP} \not\subseteq \text{P}/\text{poly}$. This can also be viewed as a new bottom-up approach, where gradually lifting the restriction on the amount of non-determinism used by the circuits, would lead to $\text{NEXP} \not\subseteq \text{P}/\text{poly}$.

1.1 Fast (unambiguous) algorithms imply UEXP circuit lower bounds

In the direction of algorithm-design too, the questions were relaxed. The initial question was: do NP-complete languages have polynomial time algorithms? The relaxed version is: do they have algorithms that beat the naive brute-force strategy? For solving circuit satisfiability (Ckt – SAT), the canonical NP-complete language, the naive approach runs in $2^n m$ deterministic time for n -input m -size circuits. The relaxed questions are: Is it possible to design a $2^{cn} \text{poly}(m)$ time algorithm for any constant $c < 1$? Or even $2^n \text{poly}(m)/\text{sp}(n)$ time algorithm for some super-polynomial function sp ? While pessimists formulated many conjunctures [46, 44, 20, 97] believing that no substantial progress is possible, optimists took the bottom-up approach and started designing fast algorithms for restricted classes. The most studied restrictions in this line are: $3\text{-CNF} \subseteq k\text{-CNF} \subseteq \text{CNF} \subseteq \text{AC}^0 \subseteq \text{ACC} \subseteq \text{TC}^0 \subseteq \text{NC}^1 \subseteq \text{NC} \subseteq \text{P}/\text{poly}$ [65, 75, 82, 74, 83, 17, 40, 67, 63, 18, 43, 100, 45, 102, 81, 85, 23].

We often think of algorithm design and lower bounds as being antithetical, for instance: the two statements, $\text{P} = \text{NP}$ and $\text{P} \neq \text{NP}$, can't be simultaneously true. But there have been a series of results showing that efficient algorithm for certain problem in certain computation model, implies a lower bound for related problem in other computation model [42, 99, 100, 102, 68, 53]. The intuition behind these connections is: if there is a fast algorithm for a circuit class, then the algorithm must be exploiting a *simple* structure or pattern that exists in that class, and thus that class can't

simulate *complex* classes.

Williams in his ACC lower bound result, first connects the relaxed versions of $P = NP$ and $NP \not\subseteq P/poly$ that we discussed above [99]: for any super-polynomial function sp , a $2^n/sp(n)$ time Ckt – SAT algorithm for polynomial-size Boolean circuits implies $NEXP \not\subseteq P/poly$. Infact, an equally fast non-deterministic algorithm for circuit tautology suffices. His connection (and many others in the literature) also preserve the nature of circuits. That is, fast satisfiability algorithm for a restricted sub-class of Boolean circuits C , implies lower bounds against C circuits. He developed fast satisfiability algorithm for ACC circuits, and then used this connection to establish $NEXP$ and EXP^{NP} lower bounds [100].

Unfortunately, most of these connections are only known to show circuit lower bounds in relatively large complexity classes such as $NEXP$ or EXP^{NP} (although [68] extends this to scaled-down versions of these classes). Establishment of similar connections for lower classes like EXP , can be seen as the first step towards establishing lower bounds for them.

One of the main focus of this thesis is to extend these connections to non-uniform lower bounds for the class $UEXP$ of languages recognized by unambiguous non-deterministic machines, and to related classes. Since $UEXP$ lies between EXP and $NEXP$, we believe that lower bounds for $UEXP$ based on algorithms would be progress towards making similar connections for EXP . We don't know how far $UEXP$ is from $NEXP$. In [94] a randomized reduction was given from NP to promise-UP which only succeeds with a low probability. In [27] it was shown that derandomizing this reduction, or even increasing the success probability, will have unlikely consequences.

We show that fast unambiguous non-deterministic circuit analysis algorithms imply circuit lower bounds for $UEXP$. In our first result we use fast tautology and canonization algorithms. Roughly speaking, a canonization algorithm for a circuit class, is an algorithm that only accepts one circuit per function or truth-table, from that class (see Section 4.1 for a technical definition).

Theorem 1. *USUBE algorithms for tautology and canonization of C imply $UE/O(n \log n) \not\subseteq C$.*

We use C to denote any non-uniform circuit class from the set $\{AC_0, ACC_0, TC_0, NC_1, NC,$

$P/poly\}$. This set is called *typical* in the literature. We measure the complexity of any circuit analysis algorithm in terms of the input wires, and not the circuit size. We omit the size parameter for circuits if it's polynomial, i.e. $C(poly) = C$. Here by SUBE (sub-E) we denote the class of languages that have $2^{\varepsilon n}$ -time algorithm for each $\varepsilon > 0$. USUBE is the extension to unambiguous non-deterministic algorithms.

Since canonization is not prominently used circuit analysis algorithm in the lower bound literature, we replace the assumption of fast canonization with different circuit analysis algorithms. Based on the definition of canonization, a fast Π_2 SAT algorithm would be an ideal replacement. We derive the following theorem.

Theorem 2. *For every constant k_1 there is a constant k_2 , such that if Π_2 SAT on n variables and n clauses can be solved in $UTIME(2^{n/(\log n)^{k_2}})$, then $UE/n \not\subseteq SIZE(n(\log n)^{k_1})$.*

Relaxing the unambiguity condition a little, helps us totally get rid of the fast canonization assumption. One other well studied variant of $UTIME(t)$ is $FewTIME(t)$, which was first defined in [3]. $L \in FewTIME(t)$, if there exists a constant c and a non-deterministic verifier V , such that the number of accepting certificates on any input is bounded by t^c . We get results for a slightly relaxed version of $FewTIME$, the class $\widetilde{FewTIME}(t)$.

Definition 1.1.1. $\widetilde{FewTIME}(t)$ is the class of problems decidable by $NTIME(t)$ verifiers, where the number of accepting paths are bounded by $2^{2^{(\log \log t)^2}}$.

Actually, even if we bound the number of accepting paths by $2^{2^{sc(t) \log \log t}}$ for any super constant function sc , the definition would satisfy all our results. We use $sc(t) = \log \log t$ for the sake of cleaner presentation. Note that, this definition allows any \widetilde{FewE} verifier to have $2^{2^{(\log n)^2 + O(\log n)}}$ many accepting paths compared to the $2^{2^{\log n + O(1)}}$ upper bound for the $FewE$ verifiers. But it's still very 'few' compared to the maximum possible by an NE verifier, that is $2^{2^{O(n)}}$. Definition of $\widetilde{FewSUBE}$ is analogous to USUBE. We derive the following result.

Theorem 3. *$\widetilde{FewSUBE}$ tautology algorithm for C circuits implies $\forall k \widetilde{FewE}/O(n \log n) \not\subseteq C(n^k)$.*

Lastly, we replace canonization with exact proper learning (that makes membership and equivalence queries) and derive the following result. Here by proper we mean that any hypothesis produced by the learning algorithm is a polynomial-size \mathcal{C} circuit, i.e., belongs to the class being learned (see Section 4.4 for a technical definition).

Theorem 4. *USUBE algorithms for tautology and proper-learning of \mathcal{C} imply $\text{UE}/O(n \log n) \not\subseteq \mathcal{C}$.*

Almost all of our connections work for any typical restricted circuit class \mathcal{C} . But still we derive some translation results to show that, any lower bound framework (that may be different from ours), if uses certain set of fast unambiguous circuit analysis algorithms, and only works for unrestricted Boolean circuits, now can be used for \mathcal{C} . These translations are tight enough to be useful in the scenario where the new framework only requires the algorithms to be $\text{UTIME}(2^n/sp(n))$ for some super-polynomial function $sp(n)$, as opposed to our frameworks that require USUBE algorithms.

Theorem 5. *Either $\text{P} \not\subseteq \mathcal{C}$, or:*

1. *$\text{UTIME}(2^n/n^{\omega(1)})$ tautology and canonization algorithms for \mathcal{C} , imply $\text{UTIME}(2^n/n^{\omega(1)})$ tautology and canonization algorithm for unrestricted Boolean circuits; and*
2. *$\widetilde{\text{FewTIME}}(2^n/n^{\omega(1)})$ tautology algorithm for \mathcal{C} , implies $\widetilde{\text{FewTIME}}(2^n/n^{\omega(1)})$ tautology algorithm for unrestricted Boolean circuits.*

1.2 UEXP lower bounds are constructive, useful, and unique

To inquire why lower bounds for smaller classes like TC^0 are still open after decades of efforts, Williams proved rigorous equivalences between NEXP lower bounds and useful properties that are constructive [101]. He showed that the lower bound $\text{NEXP} \not\subseteq \mathcal{C}$ is equivalent to the existence of $\text{P}/\log n$ constructive property against \mathcal{C} circuits. As it is believed that there can't be any Natural Proofs against TC^0 and more expressive circuit classes, this is a negative result in the sense that any

NEXP lower bound, already satisfies two of the three conditions of Natural Proofs, regardless of the technique used to obtain it.

We extend these results to characterize UEXP lower bounds. Our connections show that the future of UEXP lower bounds is brighter in the sense that, they not only ‘not satisfy’ the third condition of Natural Proofs, but they satisfy a different third condition that is totally opposite of largeness. We introduce a new notion called Unique Proofs. Unique properties are those that contain exactly one function of each input length. Useful unique properties are implicitly proving a circuit lower bound for a specific function: the one function that has the property, but might not explicitly spell out which function the lower bound holds for. We derive the following equivalence.

Theorem 6. $UE/n \not\subseteq C$ if and only if a $P/\log n$ computable unique property exists against C .

In [72] the $NE \cap Co - NE \not\subseteq C$ lower bound was shown to yield a P computable property against C . The equivalence was also conjectured to be true. We prove that conjecture for the case of $UE \cap Co - UE$ lower bounds and P computable unique properties.

Theorem 7. $UE \cap Co - UE \not\subseteq C$ if and only if a P computable unique property exists against C .

As an application of these connections, we get USUBEXP derandomization of BPP from different UEXP lower bounds.

Theorem 8. 1. $UEXP \not\subseteq SIZE(poly) \implies BPP \subset \bigcap_{\epsilon > 0} io-UTIME(n^\epsilon)/n^\epsilon$

2. $UEXP \neq BPP \implies BPP \subset \bigcap_{\epsilon > 0} io-Heur-UTIME(n^\epsilon)/n^\epsilon$

3. $UEXP \cap Co - UEXP \not\subseteq SIZE(poly) \implies BPP \subset \bigcap_{\epsilon > 0} io-UTIME(n^\epsilon)$

4. $UEXP \cap Co - UEXP \neq BPP \implies BPP \subset \bigcap_{\epsilon > 0} io-Heur-UTIME(n^\epsilon)$

We also discuss unique properties with constructivity higher than P . Rudich [79] was the first to study NP computable natural properties. For $Co - NP$ and higher constructivity, natural properties are not that interesting to study. This is because, $Co - NP$ computable natural properties exists

against $P/poly$ circuits or any typical circuit class C , unconditionally (the algorithm just guesses a polynomial size C circuit and verify if its truth-table is equal to the property's input). For useful properties that are not required to satisfy largeness, NP constructivity is shown to be equivalent to P constructivity [4, 72]. So for such properties, both NP constructivity and Co – NP constructivity, are not very interesting.

For the case of unique properties, there is no known equivalence between NP constructivity and P constructivity. Even with Co – NP constructivity, we don't know if unique properties exist against $P/poly$, unconditionally. We only know such properties with P^{NP} constructivity (we can check if the input is the lexicographically first truth-table that requires at least $n^{\log n}$ size circuit) or P/n constructivity (advice itself can provide a unique truth-table with super-polynomial circuit complexity). So it makes sense to explore unique properties with NP constructivity (with and without sub-linear advice). We derive the following equivalence.

Theorem 9. *NE $\not\subseteq C$ if and only if an NP/ $\log n$ computable unique property exists against C*

We also prove the conjecture of [72] for the case of $NE \cap Co - NE$ lower bounds and NP computable unique properties.

Theorem 10. *$NE \cap Co - NE \not\subseteq C$ if and only if an NP computable unique property exists against C*

We extend the techniques to show isolation of properties with different versions of NP constructivity, in the sense that, existence of a property implies existence of a unique property with the same constructivity. These isolation results indicate that NP-computable properties are not too powerful (since their largeness can be diluted).

Theorem 11. *The following properties can be isolated within same constructivity:*

1. *an NP/ $\log n$ computable property against C*
2. *the set of NP/ $O(1)$ computable properties against $C(n^k)$ for each $k \geq 1$*
3. *the set of NP computable promise properties against $C(n^k)$ for each $k \geq 1$*

Here by isolation of a set of properties we mean that, if all the properties in the set exist, then each property has an equivalent unique property. Note that, for the case of $P/\log n$ or $NP/\log n$ computable properties that are not unique or large, an equivalence is known between the existence of a property against C , and the existence of the set of properties against $C(n^k)$ for each $k \geq 1$ [101]. Such equivalence is not known for the case of natural and unique properties.

By promise property we mean that, the property only needs to contain functions for the input lengths where it is useful, and not on all input lengths. Note that, according to the definition of Razborov and Rudich, a useful property is only required to be useful on infinitely many input lengths, but needs to satisfy the largeness condition on all input lengths.

1.3 Gradually increasing the non-determinism in circuits for NEXP lower bounds: with a hope to prove $NEXP \not\subseteq P/poly$

We extend the known NEXP lower bounds by allowing the restricted circuit classes some amount of non-determinism. We also discuss why it would be difficult to increase this non-determinism without proving stronger lower bounds, such as $NEXP \not\subseteq P/poly$. This can also be seen as an approach to prove $NEXP \not\subseteq P/poly$, by gradually increasing the non-determinism in circuits in our current lower bounds.

We use a weaker version of non-determinism (due to technical difficulties that we discuss in the next section), but we'll see that even this kind of non-determinism is very powerful. We first define all the types of non-determinism that circuits in our lower bounds would use.

$P/poly$ is equivalent to the class of non-uniform polynomial size circuits. A language $L \in SIZE(s)$ if L is accepted by a sequence of deterministic Boolean circuits $\{C_n\}_{n \in \mathbb{N}}$ of size $O(s(n))$, where C_n computes L_n (n^{th} -slice of L) and the size is measured by the number of wires.

Similarly $NP/poly$ is equivalent to the class of non-uniform non-deterministic polynomial size circuits. A non-deterministic circuit has extra guess inputs, and the circuit accepts an input, if

there is a setting of these guess inputs that makes the output 1.

Definition 1.3.1. A language $L \in \text{NSIZE}(s)$ if L is accepted by a sequence of non-deterministic Boolean circuits $\{C_n\}_{n \in \mathbb{N}}$ of size $O(s(n))$. C_n receives two inputs, x of length n and guess input y . The function $f_C : \{0, 1\}^n \rightarrow \{0, 1\}$ computed by C satisfies, $f_C(x) = 1 \iff \exists y C(x, y) = 1$.

$\text{NP}/poly \cap \text{Co-NP}/poly$ is equivalent to the class of languages that have non-deterministic circuits for both, the language and its complement. These circuits can also be combined to output an equivalent Single-Valued or SV circuit.

Definition 1.3.2. A language $L \in \text{SVSIZE}(s)$ if L is accepted by a sequence of non-deterministic Single-Valued Boolean circuits $\{C_n\}_{n \in \mathbb{N}}$ of size $O(s(n))$. C_n receives two inputs, x of length n and guess input y . C_n has two outputs, $Flag_{C_n}$ and $Value_{C_n}$. The circuit C_n computes function $f_C : \{0, 1\}^n \rightarrow \{0, 1\}$ if it satisfies the following two promises for any input x : (a) $\exists y Flag_{C_n}(x, y) = 1$; (b) $\forall y Flag_{C_n}(x, y) = 1 \implies Value_{C_n}(x, y) = f_C(x)$.

We define a new type of non-determinism to capture the class $(\text{NP} \cap \text{Co-NP})/poly$. This class is weaker than $\text{NP}/poly \cap \text{Co-NP}/poly$ as both the algorithms need to be complimentary on all the advice sequences. For any language $L \in (\text{NTIME}(t) \cap \text{Co-NTIME}(t))/t$, there is a $\text{DTIME}(t^2)$ algorithm, that on any t size input (which is actually the advice), outputs a pair of non-deterministic circuits (that correspond to the $\text{NTIME}(t)$ and $\text{Co-NTIME}(t)$ algorithms) that accept complimentary set of inputs, and there is an infinite sequence of $t(n)$ -size inputs (one for each $n \in \mathbb{N}$) for which the produced pair of circuits accept L and \bar{L} . We call such pair of circuits promise-SV or prSV, since there is an algorithm that produces these circuits, and the algorithm satisfies the promise of always producing circuits that can be combined to become SV. We call the underlying algorithm, a prSV algorithm.

Definition 1.3.3. A linear-time algorithm \mathcal{A} is called prSV if on each input it outputs a pair of non-deterministic circuits that accepts some n' -bit function $f_{n'}$ and its compliment, for some $n' \leq n$. We say $L \in \text{prSV}^{\mathcal{A}}\text{SIZE}(s)$, if for $n \in \mathbb{N}$, L_n has $O(s(n))$ size non-deterministic circuits C_n and C'_n ,

deciding it and its compliment. Additionally, these circuits are produced by the prSV algorithm \mathcal{A} on some $s(n)$ -length input. $L \in \text{prSVSIZE}(s)$ denotes that $L \in \text{prSV}^{\mathcal{A}}\text{SIZE}(s)$ for some prSV algorithm \mathcal{A} .

The equation $\text{prSVSIZE}(poly) = (\text{NP} \cap \text{Co-NP})/poly$ follows from the definition. We also get, $(\text{NTIME}(n^{k/2}) \cap \text{Co-NP})/n^{k/2} \subseteq \text{prSVSIZE}(n^k) \subseteq (\text{NTIME}(n^k) \cap \text{Co-NP})/n^k$.

These definitions of non-deterministic circuits naturally extend to any restricted circuit class \mathcal{C} of Boolean circuits. For the prSV circuits, the underlying prSV algorithm satisfies an extra promise of always producing \mathcal{C} circuits. This gives us $\mathcal{C}(s) \subseteq \text{prSV-}\mathcal{C}(s) \subseteq \text{SV-}\mathcal{C}(s) \subseteq \text{N-}\mathcal{C}(s)$ for any size parameter $n \leq s(n) \leq 2^n$. We use $\text{prSV}^a\text{-}\mathcal{C}(s)$ to denote $\mathcal{C}(s)$ circuit that uses a amount of prSV non-determinism. $\text{SV}^a\text{-}\mathcal{C}(s)$ and $\text{N}^a\text{-}\mathcal{C}(s)$ are defined similarly.

Now, we discuss why non-determinism, even the prSV type, is very powerful and can lead to big lower bounds.

Theorem 12. *The class $\text{SIZE}(s)$ is contained in $\text{prSV-3-CNF}(s \log n)$.*

Proof. For $L \in \text{SIZE}(s)$ we give a $\text{prSV-3-CNF}(s \log n)$ circuit sequence. The underlying prSV algorithm treats its inputs as $\text{SIZE}(s)$ circuits, converts fan-in of each gate to two by adding more gates, and then converts the input circuit and its compliment into two 3 – CNF circuits with $O(s \log n)$ guess inputs. The algorithm applies Tseitin transformation [91]: for each gate $g(x, y)$ it introduces a new variable y that it labels as guess input, and adds clauses for the equation $y = g(x, y)$. Finally adds a clause with just the guess variable that represents the output gate. \square

Such conversions were discussed in [66] for non-deterministic circuits, we observe that they also extend to prSV circuits. In fact unambiguous prSV non-determinism suffices (the guess inputs introduced in the proof represents gates of the original deterministic circuit, and take unique values on any input). Also, the multiplicative $\log n$ factor can be removed if we start with fan-in two unrestricted Boolean circuits.

The above theorem shows that lower bounds against restricted circuits that contain 3 – CNF and use prSV type of non-determinism, imply lower bounds against unrestricted Boolean circuits.

We derive lower bounds for NE and E^{NP} against such classes, with limited non-determinism. Increasing the non-determinism in our results won't be possible without proving lower bounds like ' E^{NP} is not simulated by linear-size fan-in-two unrestricted Boolean circuits', which are still very far from the reach of current lower bound techniques. Our results also give hope of obtaining TC^0 lower bounds, if one can simulate threshold gates, by the use of less expressive gates and limited non-determinism.

Theorem 13. $\bigcap_{\epsilon > 0} \text{prSV}^{n^\epsilon}\text{-ACC}$ can't simulate NE.

For E^{NP} we get a variety of lower bounds where the amount of non-determinism increases as we go down from ACC to k -CNF.

Theorem 14. $\bigcap_{\epsilon > 0} \text{prSV}^{n/(\log n)^\epsilon}\text{-AC}^0$, $\bigcap_{\epsilon > 0} \text{prSV}^{\epsilon n/(\log n)^2}\text{-k-CNF}$, or $\bigcap_{\epsilon > 0} \text{prSV}^{\epsilon n}\text{-AC}^0(n)$, can't simulate E^{NP} .

One can directly get the lower bound $E^{NP} \not\subseteq \bigcap_{\epsilon > 0} N^{n^\epsilon}\text{-ACC}$ by using Williams sub-exponential size ACC lower bound. But this direct approach won't work for lower bounds against sub-linear non-determinism, and for NE lower bounds (see the next section for full details). Our lower bounds follow from a more general connection that we build by extending Williams's connection.

Theorem 15. For super-polynomial function sp and $s(n) \leq O(n)$:

1. an $\text{NTIME}(2^{n-s(n)^c}/sp(n))$ C-tautology algorithm for every $c > 0$ implies $\text{NE} \not\subseteq \text{prSV}^{s(n)\text{-C}}$
2. an $\text{NTIME}(2^{n-3s(n)}/sp(n))$ C-tautology algorithm implies $E^{NP} \not\subseteq \text{prSV}^{s(n)\text{-C}}$

We also extend Santhanam's [80] lower bound against fixed-polynomial size deterministic circuits to prSV circuits.

Theorem 16. $\forall k \geq 1 \text{ prAM} \not\subseteq \text{prSVSIZE}(n^k)$ and $\forall k \geq 1 \text{ AM}/\omega_n(1) \not\subseteq \text{prSVSIZE}(n^k)$.

We also extend Williams's connection between non-trivial GAP-SAT algorithm and NEXP lower bounds. GAP-SAT is the promise problem, where the positive inputs are tautology circuits, and

negative inputs are s -size circuits that have at most $2^n(1 - 1/s)$ satisfying assignments. Note that, a tautology or a CAPP algorithm also imply a GAP-SAT algorithm. CAPP is the problem of computing the acceptance probability of s -size circuits within an additive error of $\pm 1/s$.

Theorem 17. *An $\text{NTIME}(2^n/sp(n))$ GAP-SAT algorithm for n -input polynomial-size $(\text{NP} \cap \text{Co} - \text{NP})$ -oracle circuits, for any super-polynomial function $sp(n)$, implies $\text{NEXP} \not\subseteq (\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$.*

As there is no complete language in $\text{NP} \cap \text{Co} - \text{NP}$, we need a (possibly different) non-trivial algorithm for each $\text{NP} \cap \text{Co} - \text{NP}$ oracle. Note that, for any $A \in \text{NP} \cap \text{Co} - \text{NP}$, tautology (or CAPP) for poly-size A -oracle circuits has a trivial algorithm that runs in non-deterministic $\text{poly}(n)2^n$ -time: for all 2^n inputs, non-deterministically guess the answers to all the oracle queries, and guess their certificates (for A for any positive answer, for \bar{A} for any negative answer).

One can also view $\text{NP} \cap \text{Co} - \text{NP}$ as $\text{P}^{\text{NP} \cap \text{Co} - \text{NP}}$. In this view, our result works separately for any one $\text{NP} \cap \text{Co} - \text{NP}$ oracle A : a non-trivial GAP-SAT algorithm for A -oracle circuits implies $\text{NEXP} \not\subseteq \text{P}^A/\text{poly}$. So our result is essentially a relativized version of Williams's result upto $(\text{NP} \cap \text{Co} - \text{NP})$ oracles.

1.4 Our techniques, interesting by-products, and previous work

Lower bounds from Karp-Lipton Theorems and fast tautology algorithms:

Previous work: The idea of fast algorithms to lower bounds, can be traced back to the first paper where the non-uniform class P/poly was discussed (by Karp and Lipton [55]), where one of the corollaries (credited to Meyer) is that $\text{P} = \text{NP} \implies \text{EXP} \not\subseteq \text{P}/\text{poly}$. This can be interpreted as: a polynomial time algorithm for $\text{Ckt} - \text{SAT}$ (or any other NP-complete problem), implies $\text{EXP} \not\subseteq \text{P}/\text{poly}$. The connection in [55] was established by first establishing $\text{EXP} \subset \text{P}/\text{poly} \implies \text{EXP} = \Sigma_2^{\text{P}}$. Any such non-uniform to uniform collapse is famously called Karp-Lipton Theorem (or KLT for short) in the literature. The assumed fast SAT algorithm, along with the assumption $\text{EXP} \subset \text{P}/\text{poly}$, implies $\text{EXP} = \Sigma_2^{\text{P}} = \text{P}$ and thus contradicts the deterministic time hierarchy [37, 39].

A similar KLT was proved for NEXP in [42]. Similar KLTs have also been proved for EXP^{NP} [14], PSPACE [9], and related classes [29, 41]. Using the NEXP KLT, an NSUBEXP tautology algorithm contradicts the non-deterministic time hierarchy [26, 84, 104, 33] by showing $\text{NEXP} = \Sigma_2^{\text{P}} = \text{NSUBEXP}$. The collapse was extend from $\Sigma_2^{\text{P}} \cap \Pi_2^{\text{P}}$ to MA in [9]. This collapse yields $\text{NEXP} \not\subseteq \text{P}/\text{poly}$ from an NSUBEXP CAPP algorithm.

Our work: As a starting point for our connection between fast algorithms and UEXP lower bounds, we design similar KLTs for the intermediate classes: $\widetilde{\text{FewEXP}}$, UEXP, and $\text{UEXP} \cap \text{Co} - \text{UEXP}$.

Theorem 18. *For any $k \geq 1$:*

1. $\widetilde{\text{FewE}}/n \subset \text{SIZE}(n^k) \implies \widetilde{\text{FewEXP}} = \text{MA}$
2. $\text{UEXP} \subset \text{SIZE}(\text{poly}) \implies \text{UEXP} = \text{MA}$
3. $\text{UEXP} \cap \text{Co} - \text{UEXP} \subset \text{SIZE}(\text{poly}) \implies \text{UEXP} \cap \text{Co} - \text{UEXP} = \text{MA}$

Unfortunately our KLTs doesn't give the desired lower bound from USUBEXP, or even SUBEXP algorithms. This is because the first quantifiers of Σ_2^{P} and MA are not unambiguous. In any case, SUBEXP tautology or CAPP algorithms are far from the reach of current algorithm design techniques. Soon we'll see that USUBE algorithms are sufficient for UEXP lower bounds.

Easy-witness Lemmas:

Previous work: Williams's work [99] show that even a small progress in SAT and CAPP algorithms will prove super-polynomial NEXP lower bounds. His connection is tighter because he uses the easy-witness lemma (or EWL for short) from [42], which was the core of the collapse of NEXP to EXP in the NEXP KLT. The easy-witness technique was first introduced in [52]. We also derive analogous EWL for UEXP while establishing the UEXP KLT.

EWL for NEXP states that, if NEXP has P/poly circuits, then any NEXP verifier V , for every input x that it accepts, has a certificate y with $V(x, y) = 1$ that is encoded by the truth-table of a P/poly circuit (when the truth-table is seen as a concatenation of the 2^n outputs). In short, if NEXP

has small circuits, than all NEXP verifiers have small witnesses (i.e., witness encoded by small circuits). The proof is given by contradiction: if NEXP witnesses have high circuit complexity, then they can be used to derandomize MA and contradict a certain lower bound.

Our techniques: We prove same connection between the circuit complexity of $\widetilde{\text{FewEXP}}$, UEXP, and $\text{UEXP} \cap \text{Co-UEXP}$, and the witness complexity for corresponding verifiers [28]. Our EWL for $\widetilde{\text{FewEXP}}$ uses the technique from [42], and exploits the fact that derandomization of MA only requires limited non-determinism. For UEXP and $\text{UEXP} \cap \text{Co-UEXP}$ EWLs, we use a simpler technique that gives even stronger results. We show that a specific version of search problem (that searches for the lexicographically smallest certificate) for UEXP verifiers lies within UEXP itself. The circuit complexity of that search version is same as the circuit complexity of witnesses for UEXP verifiers, so we directly get the EWL from this reduction. Note that, such reductions are not possible for NEXP verifiers unless $\text{EXP}^{\text{NP}} = \text{NEXP}$ [47].

Lower bounds from Easy-witness Lemmas and fast tautology algorithms:

Previous work: Even in his approach with NEXP EWL, Williams proves his result by contradicting the non-deterministic hierarchy. He only needs a $\text{NTIME}(2^n/sp(n))$ tautology algorithm for n -input polynomial-size circuits, for any super-polynomial function sp . This is a huge improvement over the NSUBEXP tautology algorithm, that was required if one uses the NEXP KLT directly.

Our work and technique: We use Williams's framework and get UEXP lower bounds from $\text{USUBE}(\cup_{\epsilon>0} \text{UTIME}(2^{\epsilon n}))$ algorithm by using our UEXP EWL. The main reason why we need faster algorithms is that we don't have any hierarchy for UTIME that is as good as the hierarchies for NTIME. Note that, before this no UEXP lower bound was obtained, even from a deterministic SUBEXP tautology algorithm. Although our approach requires a canonization algorithm too, for unambiguously guessing the witness circuits, we get rid of this requirement for the case of $\widetilde{\text{FewEXP}}$ lower bounds.

Related work: Similar results were also proved for BPEXP [22], where they require randomized tautology algorithms (with two-sided error). They require algorithms to run in $2^{n/(\log n)^{\omega(1)}}$,

which is slightly faster than SUBE. Although, they get sharper lower bounds than $\text{BPEXP} \not\subseteq \text{P}/\text{poly}$, namely $\text{BPE} \not\subseteq \text{SIZE}(n(\log n)^{O(1)})$, and only require the fast algorithm to run for quadratic-size circuits, BPEXP is incomparable to NEXP . Only a fast randomized algorithm (with zero-sided error) to ZPEXP lower bound, or a fast randomized algorithm (with one-sided error) to REXP lower bound, would be considered a strict improvement over Williams’s connection for NEXP .

In [68], they extended Williams’s connection to NQP (non-deterministic quasi-polynomial time). There are two major differences with our work. First, the NQP lower bounds require fast algorithms for circuits with size sub-exponential or higher, whereas our results only need fast algorithms for polynomial-size circuits. Second, there is no known comparison between NQP and $\widetilde{\text{FewEXP}}$. Even if we try to distribute all non-deterministic branches of an NQP verifier within exponential time, the new EXP verifier will have more branches than $2^{n^{O(\log n)}}$ (also note that, our results would go through even for slightly strict definition of $\widetilde{\text{Few}}$).

Super-linear lower bounds from fast $\Pi_2\text{SAT}$ algorithms:

Previous work: In [22] they also showed that a randomized algorithm for $\Pi_2\text{SAT}$ with linear-clauses that runs in $2^{n/(\log n)^{\omega(1)}}$ time implies $\text{BPE} \not\subseteq \text{SIZE}(n(\log n)^{O(1)})$.

Our work: We show an analogous result for unambiguous non-deterministic time. We show that, a $\text{UTIME}(2^{n/(\log n)^{\omega(1)}})$ algorithm for $\Pi_2\text{SAT}$ on linear clauses implies $\text{UE}/n \not\subseteq \text{SIZE}(n(\log n)^{O(1)})$.

Our technique: We get this result by generalizing Williams’s connection, whereas in [22] they use altogether different techniques. Note that, one can’t directly get such connections between non-deterministic $\Pi_2\text{SAT}$ algorithms and NE lower bounds, because the NEXP EWL is not as fine-grained as our UEXP EWL.

Lower bounds from fast learning algorithms:

Previous work: The two commonly studied learning models are: “the Angluin’s exact learning model” [6], and “the Valiant’s PAC model” [93]. Fast learning algorithms in these models have been known to yield lower bounds [31, 73, 36, 57, 72].

In [31] it was implicit that, if a circuit (concept) class \mathcal{C} is exact learnable in $\text{SUBEXP}^{\text{NP}}$, then $\text{EXP}^{\text{NP}} \not\subseteq \mathcal{C}$. In [36, 57] it was improved to: SUBEXP learning algorithm implies $\text{EXP} \not\subseteq \mathcal{C}$. In [72] it was implicit that: NSUBEXP learning algorithm (where on any input circuit regardless of its size, there is one branch where the algorithm outputs a hypothesis, and the hypothesis is guaranteed to be correct only for polynomial size circuits) implies $\text{NEXP} \not\subseteq \mathcal{C}$. In [73] it was shown that if \mathcal{C} admits $2^n/n^{\omega(1)}$ randomized (weak) learning algorithm (with membership queries), then $\text{BPEXP} \not\subseteq \mathcal{C}$.

Our work: We show that unambiguous learning algorithm implies UE lower bounds. As far as we know this is the first lower bound result from unambiguous learning. Our results are weak in sense that: the lower bound requires UE to use $O(n \log n)$ bits of advice, we also need a tautology algorithm to assist the learning, and our algorithm only makes membership queries to polynomial-size circuits (it can make queries to sub-exponential size too, but then we would require fast tautology algorithm for sub-exponential size circuits). Even with these small weaknesses, our connection is not implied by any of the previous known connections, and our lower bound is strictly better than $\text{EXP}^{\text{NP}} \not\subseteq \mathcal{C}$ and $\text{NEXP} \not\subseteq \mathcal{C}$, and incomparable to $\text{BPEXP} \not\subseteq \mathcal{C}$. For the $\text{EXP} \not\subseteq \mathcal{C}$ lower bound, they used a clever diagonalization argument in [57] that directly doesn't work for UTIME learning (and also NTIME learning, because it's not clear how the diagonalization process will beat all the non-deterministic branches for all the \mathcal{C} circuits).

Our Technique: Unlike other results, we use Williams's framework itself, and simulate canonization using a learning algorithm that is assisted with a tautology algorithm.

Avoiding the Natural Proofs barrier:

Our work and technique: We avoid the barrier by taking the route $\text{NEXP} \supseteq \text{UEXP} \supseteq \text{EXP}$, and hit the much safer unique properties. Using our UEXP EWL we prove equivalence between UEXP lower bounds and natural properties that are not large, instead are unique. EXP lower bounds were known to yield P computable unique properties, our work extends this to $\text{UEXP} \cap \text{Co}$ – UEXP lower bounds, and also establish the equivalence.

Related work: We put an upper bound on the number of non-deterministic branches that

accept. If one wants to start by putting an lower bound, they would have to take the $\text{NEXP} \supseteq \text{REXP} \supseteq \text{EXP}$ route. Williams work [101] also indicates that this other direction would not be feasible, if one wants to use the easy-witness technique, as certain lower bounds for witnesses of REXP and ZPEXP are equivalent to the existence of natural properties.

Other related work [24] that talks about properties sparser than natural properties, shows that under the same cryptography primitives that indicates the non-existence of natural properties, there is a property with density $1/2^{n^{(\log n)^{O(1)}}}$ (vs. density $1/2^{O(n)}$ of a natural property) against $P/poly$, that establishes $\text{NP} \not\subseteq P/poly$. In [21] they give hope of avoiding the Natural Proofs barrier by establishing equivalence between $\text{NP} \not\subseteq P/poly$ and natural properties that accept SAT and are useful against only those polynomial-size circuits that never error on SAT.

Lower bounds to derandomization:

Previous work: Relationship between derandomization and uniform/non-uniform lower bounds has been studied extensively in the past [70, 71, 9, 49, 42, 101, 92, 89, 48, 19, 8]. In our results we only focus on the lower end of this spectrum, i.e. derandomization that requires sub-exponential time. Note that, BPP is a sub-class of $P/poly$, and we also don't know if $\text{NEXP} \neq \text{BPP}$.

First in a series of work [9, 70, 71] non-uniform lower bounds were shown to yield derandomization of the class BPP . In the lower end of the spectrum it was shown that $\text{EXP} \not\subseteq P/poly$ implies $\text{BPP} \subset \text{io-SUBEXP}$. Later in [48] this connection was extended to the uniform lower bound $\text{EXP} \neq \text{BPP}$. This lower bound was actually shown to be equivalent to $\text{BPP} \subset \text{io-Heur-SUBEXP}$. In [42, 101] these connections and equivalences were extended to NEXP and REXP lower bounds, and derandomization that works in NSUBEXP and ZPSUBEXP .

Our work: We extend these connections to UEXP and $\text{UEXP} \cap \text{Co} - \text{UEXP}$ lower bounds, and derandomization that works in USUBEXP .

Our technique: We use our connections between UEXP lower bounds and unique properties. We only get the lower bounds to derandomization connections, and not the reverse connections. It is due to the lack of complete languages and strong hierarchies for UEXP .

Unconditional super-polynomial lower bounds:

Previous work: Although, TC^0 lower bounds are still untouchable, and at this point we don't even have NEXP lower bound for depth-two circuits with linear-threshold gates, some improvements have been made after Williams's ACC lower bound. In [102] it was shown that NEXP doesn't have ACC circuits where the bottom most layer is allowed to have linear-threshold gates. In [98] the circuit class was further generalized by allowing the top gate to be any sparse symmetric function (exact majority is one example).

Our work: We extend the NEXP lower bound against ACC by allowing the circuits to use sub-polynomial amount of prSV non-determinism. For the case of E^{NP} , our lower bounds allow almost sub-linear amount of non-determinism as we go down to non-uniform k -CNF circuits. Recall that, even k -CNF circuits are very powerful with such type of non-determinism (Theorem 12), and if we increase the amount to linear we will get super-linear lower bounds for E^{NP} .

Our techniques: One can directly get the lower bounds against $\bigcap_{\epsilon>0} N^{n^\epsilon}$ -ACC for E^{NP} and NEXP witnesses, by using Williams sub-exponential size ACC lower bound. Any N^{n^ϵ} -ACC circuit can be converted to a sub-exponential size ACC circuit by OR-ing over all the non-deterministic inputs. Williams used his fast ACC algorithm for sub-exponential size in his result.

There are two drawbacks of this direct approach. First, it requires a very large size. For instance, for sub-linear non-deterministic inputs, the size of the resultant deterministic circuit is $\bigcap_{\epsilon>0} \text{SIZE}(2^{\epsilon n})$. Second, even if we have fast algorithms for these large circuits, the lower bounds don't transfer to NEXP. The fast algorithms can only give lower bounds for NEXP witnesses, or E^{NP} . This is because the NEXP EWL is not that fine-grained.

So we design an NEXP EWL for prSV non-deterministic circuits. Our EWL also works for the case of limited non-determinism. As far as we know, this is the first EWL that talks about any kind of non-deterministic circuits. We use our EWL and combine the non-deterministic and co-non-deterministic circuits for NEXP witnesses in a clever way, to yield $\text{NEXP} \not\subseteq \bigcap_{\epsilon>0} \text{prSV}^{n^\epsilon} \text{ACC}$. Since our technique only requires fast algorithms for polynomial-size circuits, we also get lower

bounds with sub-linear non-determinism, for E^{NP} against circuit classes lower than ACC.

Previous work: Note that, there was already an NEXP KLT for $(NP \cap Co - NP)/poly$ which gives $NEXP = AM$ [96], and can be extended to $NEXP = MA^{NP \cap Co - NP}$ using results from [16]. This KLT implies that an NSUBEXP tautology algorithm would yield $NEXP \not\subseteq (NP \cap Co - NP)/poly$. But we use the EWL to derive lower bounds from much slower algorithms.

Extension of NEXP KLT and new gap theorems:

Our work: While deriving the EWL we also prove the converse of this KLT and extend it to the class $EXP_{||}^{NP}$, where the subscript ‘||’ means that the algorithm is only allowed to make non-adaptive queries. This extension also applies to the NEXP KLT for $P/poly$. The equivalence of non-uniform lower bounds for NEXP and $EXP_{||}^{NP}$ were already known (attributed to Buhrman in [30]). Our result proves an equivalence between uniform lower bounds, and thus results in a better gap theorem for MA than what was previous known [42]. We also get similar gap theorem for $MA^{NP \cap Co - NP}$. As $MA \subseteq MA^{NP \cap Co - NP} \subseteq AM$, this can be seen as an intermediate step for proving a gap theorem for AM.

Using the EWL we also get a gap/speed-up theorem for CAPP for $(NP \cap Co - NP)$ -oracle circuits. We first extend Williams connection to show that, non-trivial CAPP algorithm for $(NP \cap Co - NP)$ -oracle circuits imply $NEXP \not\subseteq (NP \cap Co - NP)/poly$. As this lower bound is equivalent to NSUBEXP CAPP algorithm (that works infinitely often, and uses sub-polynomial advice), we get that non-trivial savings in CAPP for $(NP \cap Co - NP)$ -oracle circuits imply sub-exponential savings.

Our technique: In the extension of the KLT to $EXP_{||}^{NP}$, in some sense we derive an EWL for $EXP_{||}^{NP}$, where the witness circuit captures all the NP-oracle queries on all n -length inputs.

Unconditional fixed-polynomial size lower bounds:

Previous work: This work of fixed-polynomial size lower bounds was started by Kannan in 1982 [54]. He used the low-end KLT for NP, which collapses the polynomial hierarchy to $\Sigma_2^P \cap \Pi_2^P$ [55], to prove fix-polynomial lower bounds for $\Sigma_2^P \cap \Pi_2^P$. Better lower bounds were proved using

improved low-end KLTs [56, 12, 15, 95, 62], before Santhnam [80] gave the lower bound for MA/1 and prMA, using the high-end KLT for PSPACE from [9].

For non-deterministic and SV non-deterministic circuits, the best high-end KLT was given in [8], which collapses PSPACE to $M(AM||Co - NP)$. In [88] this KLT was used to give fixed-polynomial lower bounds for prM(AM||Co - NP) against non-deterministic and SV non-deterministic circuits. The class $M(AM||Co - NP)$ lies in the third-level of the polynomial hierarchy and contains AM and Σ_2^P .

Our work and technique: We establish fixed-polynomial lower bounds for prAM against prSV non-deterministic circuits. Note that, in this lower bound, the class has to have one language that beats all of the $(NTIME(n^k) \cap Co - NTIME(n^k))/n^k$ algorithms for each $k \geq 1$. The main technical difficulty in proving this lower bound was the lack of complete problems for $NTIME(n^k) \cap Co - NTIME(n^k)$. Prior to our result, the best known lower bounds for prAM were against fixed-polynomial deterministic circuits.

1.5 Organization of the thesis

In the Chapter 2 we discuss and define all the technical definitions that we use. In the Chapter 3 we derive the EWL and KLT for UEXP and related classes. In the Chapter 4 we derive the connections between fast unambiguous circuit analysis algorithms and circuit lower bounds. In the Chapter 5 we derive connections between unique properties and UEXP lower bounds. In the Chapter 6 we derive fast unambiguous derandomization results from UEXP lower bounds. In the Chapter 7 we discuss the isolation of the NP computable properties. In the Chapter 8 we give all the results regarding the prSV non-deterministic circuits. Finally in the Chapter 9 we give concluding remarks and discuss some open problems.

Chapter 1 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Chapter 2

Preliminaries

Basic notations: Unless a new range is declared during the usage, we use t for time-constructible functions $n \leq t(n) \leq 2^{n^{O(1)}}$, a for advice functions $0 \leq a(n) \leq \text{poly}(n)$, s for circuit sizes (number of wires) $n \leq s(n) \leq 2^n$. For language L we use $L_n = \{x \mid x \in L \wedge |x| = n\}$ to denote the n^{th} -slice of L (or the characteristic function of L on n -length inputs). For circuit C , we use $tt(C)$ to denote its truth-table, and $|C|$ to denote its size.

Uniform classes: We assume that the reader is familiar with the standard complexity classes such as P, NP, RP, UP, BPP, ZPP, MA, AM, Σ_2^P , Π_2^P , PH (see [7]) and their corresponding complexity measures, DTIME, NTIME, RTIME, UTIME, BPTIME, ZPTIME, MATIME, AMTIME, Σ_2 TIME, Π_2 TIME. For the special cases of Σ_2^P, Π_2^P , we omit the superscript P and simply write Σ_2, Π_2 . For $C = D, N, R, U, BP, ZP, MA, AM, \Sigma_2, \Pi_2$: $\text{CTIME}(t)$ denotes the class of languages accepted by CTIME machines that run in $O(t)$ time. CE, CEXP, CSUBE, CSUBEXP, denote the classes $\cup_{c \geq 0} \text{CTIME}(2^{cn}), \cup_{c \geq 0} \text{CTIME}(2^{n^c}), \cap_{c \geq 0} \text{CTIME}(2^{cn}), \cap_{c \geq 0} \text{CTIME}(2^{n^c})$ respectively. We assume familiarity with Ckt – SAT (circuit satisfiability), Ckt – TAUT (circuit tautology), k – SAT, k – TAUT, CNF – SAT, DNF-TAUT, Σ_2 -SAT and Π_2 -SAT.

Zero-error classes: We extend the concept of zero-error class to non-deterministic and

unambiguous classes. We do this for the sake of clarity in certain arguments, and specially for distinguishing between certain non-uniform classes.

$L \in \text{ZCTIME}(t)$ if there exists a Turing machine M , that for input (x, y) with $|x| = n$ and $|y| = c \cdot t(n)$ for some constant c , runs in time $c \cdot t(n)$ for $\forall n \in \mathbb{N}$, and whose output lies in $\{1, ?\}$ if $x \in L$, and in $\{0, ?\}$ if $x \notin L$. Additionally, the quantity $\sum_{y: M(x, y) \in \{0, 1\}} 1$ is equal to: 1 for $\mathbb{C} = \text{U}$ (*uniqueness*), $\geq \frac{1}{2} \times 2^{c \cdot t(n)}$ for $\mathbb{C} = \text{R}$ (*largeness*), ≥ 1 for $\mathbb{C} = \text{N}$ (*existence*). The verifier/predicate corresponding to M is called zero-error non-deterministic. For the special cases of $\mathbb{C} = \text{U}$ and $\mathbb{C} = \text{R}$, its called zero-error unambiguous and zero-error randomized respectively.

Remark : $\text{ZRTIME} = \text{ZPTIME}$, and for $\mathbb{C} = \text{N, R, U}$, $\text{ZCTIME}(t) = \text{CTIME}(t) \cap \text{Co} - \text{CTIME}(t)$ follows by a similar argument that shows $\text{ZPTIME}(t) = \text{RTIME}(t) \cap \text{Co} - \text{RTIME}(t)$.

Circuit classes: We assume basic familiarity with Boolean circuits and their sub-classes. We use \mathcal{C} to denote any *typical* non-uniform circuit class, i.e., any class from the set $\{\text{AC}_0, \text{ACC}_0, \text{TC}_0, \text{NC}_1, \text{NC}, \text{P/poly}\}$. All these circuit classes are of polynomial size. We use $\mathcal{C}(s)$ to denote the class of $O(s)$ -size \mathcal{C} circuits. For truth-table tt , we use $\text{ckt}_{\mathcal{C}}(tt)$ to denote its exact \mathcal{C} circuit complexity, i.e. the minimum size of any \mathcal{C} circuit C whose truth-table (when concatenated to make the string $C(00\dots 0) \dots C(11\dots 1)$) is tt . In the case of unrestricted Boolean circuits, instead of $\mathcal{C}(s)$ and $\text{ckt}_{\mathcal{C}}(tt)$, we use $\text{SIZE}(s)$ and $\text{ckt}(tt)$ respectively. $\text{ckt}^M(tt)$ denotes the minimum size of any M -oracle circuit whose truth-table is tt .

Non-uniform classes: $L \in \Gamma/a$ if there exists an advice-taking Γ Turing Machine M , and advice sequence $\{a_n\}_{n \in \mathbb{N}}$ satisfying $\forall n |a_n| = a(n)$, such that: $x \in L \iff M(x)/a_{|x|} = 1$. For semantic classes, the machine M only needs to satisfy the semantic promise on the correct advice sequence $\{a_n\}_{n \in \mathbb{N}}$ (and not on all advice sequences). Below we define an exception for $\mathbb{C} = \text{N, R, U}$:

1. $L \in (\text{CTIME}(t) \cap \text{Co} - \text{CTIME}(t))/a$: If there are $\text{NTIME}(t)$ Turing Machines M and M' , and advice sequence $\{a_n\}_{n \in \mathbb{N}}$ satisfying $\forall n |a_n| = a(n)$, such that: (i) $x \in L \iff M(x)/a_{|x|} = 1$;

- (ii) both M and M' satisfy the semantic promise on $\{a_n\}_{n \in \mathbb{N}}$ (for $\mathbb{C} = \mathbb{N}$ there is no promise); and (iii) both accept complement languages. For the other advice sequences, M and M' are not required to satisfy the semantic promise, but are required to accept complement languages.
2. $L \in \text{ZCTIME}(t)/a$: It's the same as (1), except that in the “other advice sequences” part, M and M' are not required to accept complement languages. That is, both M and M' are required to accept complement languages just for some correct advice sequence, and simultaneously satisfy the semantic promise. Equivalently, there is a $\text{ZCTIME}(t)$ Turing Machine N , and advice sequence $\{a_n\}_{n \in \mathbb{N}}$ satisfying $\forall n |a_n| = a(n)$, such that: $x \in L \iff N(x)/a_{|x|} = 1$; and N satisfy the semantic promise. For other advice sequences N might: (i) fail to provide uniqueness/largeness/existence; (ii) for some input, output both 0 and 1 (on different non-deterministic branches); or (iii) for some input, output just ‘?’ (on all non-deterministic branches).
3. $L \in \text{CTIME}(t)/a \cap \text{Co-CTIME}(t)/a$: It's further relaxed than (2). We need two advice sequences $\{a_n\}_{n \in \mathbb{N}}$ and $\{b_n\}_{n \in \mathbb{N}}$, satisfying $\forall n |a_n| = a(n)$ and $\forall n |b_n| = b(n)$ (these sequences need not be the same). M satisfy the semantic promise on $\{a_n\}_{n \in \mathbb{N}}$ and accept L . M' satisfy the semantic promise on $\{b_n\}_{n \in \mathbb{N}}$ and accept \bar{L} . There are no other conditions.

Remark: $L \in \text{ZCTIME}(t)/a$ is equivalent to L having $\text{CTIME}(t)/a$ and $\text{Co-CTIME}(t)/a$ algorithms that both use the same advice. This shows:

$$(\text{CTIME}(t) \cap \text{Co-CTIME}(t))/a \subseteq \text{ZCTIME}(t)/a \subseteq \text{CTIME}(t)/a \cap \text{Co-CTIME}(t)/a \subseteq \text{ZCTIME}(t)/2a$$

So the difference between $\text{ZCTIME}(t)/a$ and $\text{CTIME}(t)/a \cap \text{Co-CTIME}(t)/a$ only matters when the amount of advice is precise.

Heuristic classes: For uniform/non-uniform class Λ , $L \in \text{Heur-}\Lambda$ if $\exists L' \in \Lambda$, such that for all polynomially samplable distributions \mathcal{D} , $\forall n \Pr_{x \sim \mathcal{D}, |x|=n} [L_n(x) = L'_n(x)] \geq 1 - \frac{1}{n}$.

Infinitely-often classes: For uniform/non-uniform, heuristic/non-heuristic class Λ , $L \in \text{io-}\Lambda$ if $\exists L' \in \Lambda$, and an infinite subset $S \subset \mathbb{N}$, such that $n \in S \implies L_n = L'_n$.

Variety of witness complexities for CTIME: We define different ways of measuring complexity of witnesses for non-deterministic verifiers. We will later see that lower-bounds based on these measures are not always the same (see Table5.1).

1. **Witnesses:** A non-deterministic verifier V for L , has witnesses in s -size C circuits, if for every $x \in L$, there is an $s(|x|)$ -size C circuit C_x , such that $V(x, tt(C_x)) = 1$. If V uses a amount of advice, then we say that V/a has witnesses in s -size C circuits, if for some correct advice sequence $\{a_n\}_{n \in \mathbb{N}}$ satisfying $\forall n |a_n| = a(n)$, for every $x \in L$, there is an $s(|x|)$ -size C circuit C_x , such that $V(x, tt(C_x))/a_{|x|} = 1$.
2. **Hitting-sets for witnesses (all witnesses in one):** A non-deterministic verifier V for L has l -size hitting-sets in s -size C circuits, if $\forall n \in \mathbb{N}$, there is an $s(n)$ -size C circuit C_n such that $tt(C_n)$ when partitioned into l strings $\{str_1, \dots, str_l\}$ of equal lengths, satisfies $\forall(x : |x| = n \wedge x \in L) \exists(i \in [1, l]) V(x, str_i) = 1$. The default value of l is 2^n . If V uses advice, hitting-sets are defined analogous to witnesses in the advice setting.
3. **Oblivious witnesses (ordered hitting-sets for witnesses):** Let y_1, \dots, y_{2^n} denote the n -length strings arranged in the lexicographical order. A non-deterministic verifier V for L has oblivious witnesses in s -size C circuits, if $\forall n \in \mathbb{N}$, there is an $s(n)$ -size C circuit C_n such that $tt(C_n)$ when partitioned into 2^n strings $\{str_1, \dots, str_{2^n}\}$ of equal lengths, satisfies $\forall(i \in [1, 2^n]) y_i \in L \implies V(y_i, str_i) = 1$. For i with $y_i \notin L$, str_i is the all 0s string. If V uses advice, oblivious witnesses are defined analogous to the witnesses in advice setting.

We use $\text{CTIME}(t)/a \subset_w C(s)$ to say that, any $\text{CTIME}(t)$ verifier, for some correct advice sequence, has witnesses in $C(s)$ circuits. $\text{CTIME}(t)/a \not\subset_w C(s)$ means that, there is a $\text{CTIME}(t)$ verifier that: for any correct advice sequence, doesn't have witness in $C(s)$ circuits infinitely often.

\subset_{hw} ($\not\subset_{hw}$) and \subset_{ow} ($\not\subset_{ow}$) are defined similarly for “hitting-sets for witnesses” and “oblivious witnesses” respectively.

Variety of seed complexities for ZCTIME: Any language in ZCTIME has two, a CTIME algorithm and a Co – CTIME algorithm deciding it. So instead of witnesses, we define a stronger notion: seeds, which is nothing but a technical way of combining witnesses from the two algorithms.

1. **Seeds:** A zero-error non-deterministic verifier V for L has seeds in s -size C circuits, if for every x , there is an $s(|x|)$ -size C circuit C_x , such that $V(x, tt(C_x)) \in \{0, 1\}$. If V uses a amount of advice, then we say that V/a has seeds in s -size C circuits, if for some correct advice sequence $\{a_n\}_{n \in \mathbb{N}}$ satisfying $\forall n |a_n| = a(n)$, for every x , there is an $s(|x|)$ -size C circuit C_x , such that $V(x, tt(C_x))/a_{|x|} \in \{0, 1\}$.
2. **Hitting-sets for seeds (all seeds in one):** A zero-error non-deterministic verifier V for L has l -size hitting-sets in s -size C circuits, if $\forall n \in \mathbb{N}$, there is an $s(n)$ -size C circuit C_n such that $tt(C_n)$ when partitioned into l strings $\{str_1, \dots, str_l\}$ of equal lengths, satisfies $\forall(x : |x| = n) \exists(i \in [1, l]) V(x, str_i) \in \{0, 1\}$. The default value of l is 2^n . If V uses advice, hitting-sets are defined analogous to seeds in the advice setting.
3. **Oblivious seeds (ordered hitting-sets for seeds):** Let y_1, \dots, y_{2^n} denote the n -length strings arranged in the lexicographical order. A zero-error non-deterministic verifier V for L has oblivious seeds in s -size C circuits, if $\forall n \in \mathbb{N}$, there is an $s(n)$ -size C circuit C_n such that $tt(C_n)$ when partitioned into 2^n strings $\{str_1, \dots, str_{2^n}\}$ of equal lengths, satisfies $\forall(i \in [1, 2^n]) V(y_i, str_i) \in \{0, 1\}$. If V uses advice, oblivious seeds are defined analogous to seeds in the advice setting.

We use $ZCTIME(t)/a \subset_w C(s)$ to say that, any $ZCTIME(t)$ verifier, for some correct advice sequence, has seeds in s -size C circuits. $ZCTIME(t)/a \not\subset_w C(s)$ means that, there is a $ZCTIME(t)$ verifier that: for any correct advice sequence, doesn't have seeds in s -size C circuits infinitely often.

\subset_{hs} ($\not\subset_{hs}$) and \subset_{os} ($\not\subset_{os}$) are defined similarly for “hitting-sets for seeds” and “oblivious seeds” respectively.

Useful properties: We define a generalized version of the natural properties.

Definition 2.0.1 (Useful uniform properties). A uniform Γ algorithm \mathcal{A} is a Γ -C property if it satisfies the first condition stated below, on the inputs that are powers of 2 (interpreted as truth-tables of Boolean functions). \mathcal{A} is said to be useful against s -size C circuits if it satisfies the second condition stated below.

1. *Size restrictions:*

$$(a) \text{ Uniqueness for } C = U: \forall n \in \mathbb{N} \sum_{x:|x|=2^n \wedge \mathcal{A}(x)=1} 1 = 1$$

$$(b) \text{ Largeness for } C = R: \forall n \in \mathbb{N} \Pr_{x:|x|=2^n}[\mathcal{A}(x) = 1] \geq \frac{1}{2^n}$$

$$(c) \text{ Existence for } C = N: \forall n \in \mathbb{N} \sum_{x:|x|=2^n \wedge \mathcal{A}(x)=1} 1 \geq 1$$

2. *Usefulness:* for infinitely many $n \in \mathbb{N}$, $\forall(x: |x| = 2^n) \mathcal{A}(x) = 1 \implies ckt_C(x) > s(n)$

Note that, in the case where s is $poly(n)$, a single algorithm \mathcal{A} should be useful against n^k -size C circuits for all k . That is, for each k , there should be infinitely many $n \in \mathbb{N}$, such that $\forall(x: |x| = 2^n) \mathcal{A}(x) = 1 \implies ckt_C(x) > n^k$.

Definition 2.0.2 (Useful properties that use advice). A Γ/a algorithm \mathcal{A} is a Γ/a -C property if it satisfies the first condition of Definition 2.0.1 on an advice sequence $\{a_n\}_{n \in \mathbb{N}}$ that satisfies $\forall n |a_n| = a(n)$. \mathcal{A} is said to be useful against s -size C circuits if it satisfies the second condition of Definition 2.0.1 on the same advice sequence $\{a_n\}_{n \in \mathbb{N}}$. For $C = U$, based on how \mathcal{A} behaves on the advice sequences other than $\{a_n\}_{n \in \mathbb{N}}$, it has two special categories:

$$1. \text{ } \Gamma/a\text{-strong-unique or } \Gamma/a\text{-}u=1: \forall n \in \mathbb{N} \forall(b_n: |b_n| \leq a(n)) \sum_{x:|x|=2^n \wedge \mathcal{A}(x)/b_n=1} 1 = 1$$

$$2. \text{ } \Gamma/a\text{-mild-unique or } \Gamma/a\text{-}u \leq 1: \forall n \in \mathbb{N} \forall(b_n: |b_n| \leq a(n)) \sum_{x:|x|=2^n \wedge \mathcal{A}(x)/b_n=1} 1 \leq 1$$

Definition 2.0.3 (Promise useful properties). A Γ/a algorithm \mathcal{A} is a Γ/a -prC property useful against s -size C circuits if there is an advice sequence $\{a_n\}_{n \in \mathbb{N}}$ satisfying $\forall n |a_n| = a(n)$ such that: for infinitely many $n \in \mathbb{N}$, \mathcal{A} satisfies the first and second conditions of Definition 2.0.1. Informally, the only condition property satisfies is: it is simultaneously non-trivial/large/unique and useful for infinitely many input lengths, and no conditions for the other input lengths (property is even allowed to be empty).

For $\square \in \{\mathbb{N}, \mathbb{R}, \mathbb{U}, u_{=1}, u_{\leq 1}, \text{pr}\mathbb{N}, \text{pr}\mathbb{R}, \text{pr}\mathbb{U}\}$, we use $\Gamma/a-\square \not\subseteq_{tt} C$ to say that there is a $\Gamma/a-\square$ property useful against $C(s)$ circuits. In other words, there is a Γ/a algorithm, whose set of accepting inputs satisfies the size restrictions of \square , and on infinitely many input lengths 2^n , algorithm accepts no truth-tables tt with $\text{ckt}_C(tt) \leq s(n)$.

For $\square \in \{\mathbb{N}, \mathbb{R}, \mathbb{U}, u_{=1}, u_{\leq 1}\}$ we use $\Gamma/a-\square \not\subseteq_{tt} \text{io-}C$ to denote properties that are useful on all but finitely many input lengths (i.e., there is a point after which the property is always useful).

Promise problems: A promise problem $\Pi = (\Pi_Y, \Pi_N)$ is a pair of disjoint sets Π_Y and Π_N . In the special case where $\Pi_Y \cup \Pi_N = \{0, 1\}^*$, Π is also a language. We say that a language L agrees with Π if: $x \in \Pi_Y \implies x \in L$ and $x \in \Pi_N \implies x \notin L$.

Infinitely-often promise classes: For semantic class Λ , a promise problem L is in the class $\text{ip-}\Lambda$ if: (i) the set of promise inputs include an infinite subset $S \subseteq \mathbb{N}$ of input lengths (contains the entire n^{th} -slice for any $n \in S$, and nothing from the other slices); (ii) there exists a Turing machine M such that: for $n \in S$ and for n -length input x , M is a Λ -type machine on x , and $M(x) = 1 \iff L(x) = 1$.

Remark: An infinitely-often promise class is different from an infinitely-often class in the sense that, in the definition of the latter, the Turing machine M should be of Λ -type for all input lengths. For example, io-MA is the class of languages that match any MA language for infinitely many input lengths. But, ip-MA is the class of languages that have an MA algorithm for infinitely many

input lengths, and for all the other input lengths the algorithm is not required to satisfy the semantic promise of MA.

Special ip-classes: We say that $L \in \text{ip-ZNE}$ or $L \in \text{ip}-(\text{NE} \cap \text{Co} - \text{NE})$ if there is an infinite subset $S \subseteq \mathbb{N}$, $L_1 \in \text{NE}$ and $L_2 \in \text{NE}$ such that: for $n \in S$ and n -length input x , $L(x) = 1 \iff L_1(x) = 1 \iff L_2(x) = 0$. We define the classes $\text{ip}-(\text{AM} \cap \text{Co} - \text{AM})$ and $\text{ip}-(\text{MA} \cap \text{Co} - \text{MA})$ similarly (for the MA case, the protocols for L_1 and L_2 are not required to satisfy the semantic promise on inputs lengths not in the set S).

Upper and lower bounds for ip-classes: We say that $\text{ip-}\Lambda \subset C(s)$ if any $L \in \text{ip-}\Lambda$, has $C(s)$ circuits for inputs lengths where the promise is met. We say that $\text{ip-}\Lambda \not\subset C(s)$ if there is an $L \in \text{ip-}\Lambda$, that for infinite subset of promise input lengths (where the promise is met) does not have $C(s)$ circuits. The upper and lower bounds for the $\text{ip-}\Lambda$ seeds / hitting-sets for seeds / oblivious-seeds are defined in the similar fashion.

Lower bounds against $\text{prSV}\Sigma_i$ circuits:

The promise $\text{SV}\Sigma_i$ circuits need special care in regards of lower bounds, as we need to deal with each prSV algorithm separately. For any class Γ , we define the following upper bounds and lower bounds against prSV circuits:

- Let \mathcal{A} be a prSV algorithm. For any 2^n -length truth-table tt , we use $\text{ckt}_{\text{SV}(\mathcal{A})}(tt)$ to denote the minimum size $s(n)$ such that: \mathcal{A} outputs C with $tt(C) = tt$ on an $s(n)$ -size input. We use $\text{ckt}_{\text{prSV}}^f(tt)$ to denote the minimum size $s(n)$ such that: any prSV algorithm with description length at most $f(n)$, outputs C with $tt(C) = tt$ on an $s(n)$ -size input.
- $\Gamma \not\subset \text{prSVSIZE}(s)$: There is an $L \in \Gamma$, such that for every prSV algorithm \mathcal{A} , there is an infinite subset $S \subset \mathbb{N}$ of input lengths, such that $n \in S \implies \forall x : |x| = s(n) \text{ } tt(\mathcal{A}(x)) \neq L_n$.
- $\Gamma \subset_{\text{ow/os/hw/hs/w/s}} \text{prSVSIZE}(s)$: For $L \in \Gamma$ and Γ verifier V for L , there is a prSV algorithm

\mathcal{A} , there is an input sequence $\{x_{s(n)}\}_{n \in \mathbb{N}}$ such that $\forall n \in \mathbb{N} \text{ tt}(\mathcal{A}(x_{s(n)}))$ is the ‘oblivious witness / oblivious seed / hitting-set for witnesses / hitting-set for seeds’ for V on n -length inputs. For the case of ‘witnesses / seeds’, $\forall n \in \mathbb{N} \forall y : |y| = n$ there is an input $x_{s(n)}$ such that $\text{tt}(\mathcal{A}(x_{s(n)}))$ is the ‘witness / seed’ for V on input y .

- $\Gamma \not\subseteq_{ow/os/hw/hs} \text{prSVSIZE}(s)$: There is an $L \in \Gamma$ and Γ verifier V for L , such that for every prSV algorithm \mathcal{A} , there is an infinite subset $S \subset \mathbb{N}$ of input lengths, such that $n \in S \implies \forall x : |x| = s(n) \text{ tt}(\mathcal{A}(x))$ is not the ‘oblivious witness / oblivious seed / hitting-set for witnesses / hitting-set for seeds’ for V on n -length inputs. For the case of ‘witnesses / seeds’, $n \in S \implies \exists y : |y| = n \forall x : |x| = s(n) \text{ tt}(\mathcal{A}(x))$ is not the ‘witness / seed’ for V on input y .

Hardness vs randomness: The process of using a function that is hard for a circuit class Λ (i.e. requires large size of Λ circuits) to yield a pseudo random generator (PRG) that fools Λ circuits (i.e. creates a sparse subset of inputs with roughly same fraction of inputs resulting in 1) is well known in the literature. A PRG G creates this sparse subset by mapping a small input length to the required larger output length (same as the input length of the circuit).

A PRG $G : s(n) \rightarrow n$ is computable in Γ if the language $L_G = \{(s, i, b) \mid \text{the } i^{\text{th}}\text{-bit of } G(s) \text{ is } b\}$ is in Γ . Inputs to G are called seeds, and their size (here $s(n)$) is called the seed length of G .

A PRG G is fooling a circuit C means: the fraction of inputs from the $2^{s(n)}$ size image of G that C accepts, is same as the fraction of all the inputs that C accepts (within error $\pm 1/n$). We use the following theorem in all our derandomization results.

Theorem 19. [71, 58, 92, 86] *There exists a universal constant g such that the following holds for any class O of oracles and oracle M , and any constants $\epsilon > 0$ and $d \geq 1$: if a Boolean function family $f = \{f_n\}_{n \in \mathbb{N}}$ computable in E^O that satisfies $\forall n \in \mathbb{N} \text{ ckt}^M(f(n)) \geq n^{gd/\epsilon}$, then there exists a PRG family $G = \{G_n\}_{n \in \mathbb{N}}$ computable in E^O , such that $G_n : n^\epsilon \rightarrow n^d$ fools n^d -size B -oracle circuits. Moreover, if circuit lower bound holds infinitely often, then G fools circuits infinitely often.*

Chapter 3

EWL and KLT for UTIME and related classes

We first give a specific search to decision reduction for UTIME (Section 3.1). Using this reduction we give the EWL and KLT for UTIME (Section 3.2). Then we describe similar results for ZUTIME (Section 3.3) and $\widetilde{\text{FewTIME}}$ (Section 3.4).

3.1 Search to decision reduction for UTIME

For $L \in \text{NP}$ and verifier V for L , there is a standard P^{NP} algorithm for the corresponding search problem. This algorithm implicitly decides the following language:

$$L_{ewl(V)} = \{(x, i) \mid \exists y [V(x, y) = 1 \wedge (i^{\text{th}}\text{-bit of } y \text{ is } 1) \wedge \forall (z <_{l.o.} y) V(x, z) = 0]\} \quad (3.1)$$

where $z <_{l.o.} y$ stands for “ z is lexicographically smaller than y ”, and the subscript $ewl(V)$ in $L_{ewl(V)}$ stands for “easy-witness language for V ”.

So if $\text{P} = \text{NP}$, then $L_{ewl(V)} \in \text{P}$. For $L \in \text{NEXP}$ and verifier V for L , such results are not known. In particular, it is not known whether $\text{NEXP} = \text{EXP}$ yields an EXP algorithm for the corresponding search problem, let alone $L_{ewl(V)}$.

$\text{NEXP} \subset_w \text{SIZE}(poly)$ yields EXP algorithms for the NEXP search problems, by a simple

brute-force argument. It is known that $\text{NEXP} \subset_w \text{SIZE}(poly)$ is equivalent to $\text{NEXP} \subset \text{SIZE}(poly)$ [42, 101], and to $\text{NEXP} = \text{MA}$ [42] (reverse implication was attributed to van Melkebeek). In [42] it was also shown that a weaker collapse, namely $\text{NEXP} = \text{AM}$, is sufficient to give EXP algorithms for NEXP search problems. This is the weakest collapse known so far.

From [47] we get: $\forall (L \in \text{NEXP}) \forall (\text{NEXP verifier } V \text{ for } L) L_{\text{ewl}(V)} \in \text{EXP} \iff \text{EXP}^{\text{NP}} = \text{EXP}$.

In this section we show that for $L \in \text{UTIME}(t)$ and unambiguous verifier V for L , $L_{\text{ewl}(V)} \in \text{UTIME}(t)$. $\text{UTIME}(t)$ languages also have $O(t)$ -time verifiers that are ambiguous. We show why it would be difficult to extend this result to all $O(t)$ -time ambiguous verifiers for $\text{UTIME}(t)$ languages.

Theorem 20. *For $L \in \text{UTIME}(t)$ and unambiguous verifier V for L , $L_{\text{ewl}(V)} \in \text{UTIME}(t(n))$ (where n is the input size for L and not $L_{\text{ewl}(V)}$). Moreover if this statement is true for every $O(t)$ -time non-deterministic verifier (ambiguous and unambiguous) for every $\text{UTIME}(t)$ language, then $\text{ZNTIME}(t) = \text{ZUTIME}(t)$.*

Proof. Algorithm for $L_{\text{ewl}(V)}$: For input (x, i) , guess a certificate y and simulate $V(x, y)$. Accept if V accepts and the i^{th} -bit of y is 1, otherwise reject. This algorithm is correct and unambiguous as V is unambiguous. It runs in time $O(t(|x|))$.

The moreover part: For $L \in \text{ZNTIME}(t)$, let V_1 and V_0 be its $\text{NTIME}(t)$ and $\text{Co-NTIME}(t)$ verifiers respectively. Consider the $\text{UTIME}(t)$ language $L' = \{0, 1\}^*$.

Using V_1 and V_0 we first construct a verifier V' for L' : if the first bit of the certificate is i , V' simulates V_i using the rest of the certificate.

Now using a $\text{UTIME}(t(n))$ algorithm \mathcal{A} for $L'_{\text{ewl}(V')}$ we give a $\text{UTIME}(t)$ algorithm for L : on input x , simulate \mathcal{A} on $(x, 1)$. Accept iff \mathcal{A} accepts.

If \mathcal{A} accepts then we know that $x \in L$ because there is no positive certificate for V' that starts with 0 (or in other words, no positive certificate for V_0). If \mathcal{A} rejects, then $x \in L$ because there is a positive certificate for V' that starts with 0 (or in other words, a positive certificate for V_0).

Similarly, there is a $\text{UTIME}(t)$ algorithm for \bar{L} , and thus $L \in \text{ZUTIME}(t)$. □

For the advice setting same proof goes through for the following adaptation of $L_{\text{ewl}(V)}$. For

non-deterministic verifier V/a , that uses a amount of advice to decide a language L , for any correct advice sequence $\{a_n\}_{n \in \mathbb{N}}$:

$$L_{ewl(V/a)} = \{(x, i) \mid \exists y [V(x, y)/a_{|x|} = 1 \wedge (i^{th}\text{-bit of } y \text{ is } 1) \wedge \forall (z <_{l.o.} y) V(x, z)/a_{|x|} = 0]\} \quad (3.2)$$

Using this adaptation we get the following stronger corollary.

Corollary 3.1.1. *For $L \in \text{UTIME}(t)/a$ and unambiguous verifier V/a for L , $L_{ewl(V/a)} \in \text{UTIME}(t)/a$.*

3.2 EWL and KLT for UTIME

Using the search to decision reduction from Theorem 20 we derive EWL for unambiguous verifiers of languages in $\text{UTIME}(t)$. Here again we see why it might be difficult to extend this to all ambiguous verifiers. Using the EWL we also get a KLT for UTIME.

Theorem 21. *For time-constructible $t \in 2^{O(n)}$, and constants c and k :*

1. $\text{UTIME}(t) \subseteq \mathcal{C}(n^k) \implies \text{UTIME}(t) \subseteq_{ow} \mathcal{C}(n^k)$. *Moreover if this statement is true for every $O(t)$ -time non-deterministic verifier (ambiguous and unambiguous) for every $\text{UTIME}(t)$ language, even just for witnesses (let alone oblivious-witnesses), then $\text{ZNTIME}(t) \subseteq \text{DTIME}(2^{n^{k+1}})$.*
2. $\text{UTIME}(t)/a \subseteq \mathcal{C}(n^k) \implies \text{UTIME}(t)/a \subseteq_{ow} \mathcal{C}(n^k)$.
3. $\text{UTIME}(2^{n^c})/a \subseteq \mathcal{C}(n^k) \implies \text{UTIME}(2^{n^c})/a \subseteq_{ow} \mathcal{C}(n^{ck})$.
4. $\text{UEXP}/a \subseteq \text{SIZE}(\text{poly}) \implies \text{UEXP}/a = \text{MA}/a$.

Proof. Proof of (1): For $L \in \text{UTIME}(t)$, let $x \in L$ be an n -length input, and V be an unambiguous verifier for L whose certificate length is $\leq d \cdot t$ for some constant d . The $\text{UTIME}(t)$ algorithm for $L_{ewl(V)}$ from Theorem 20 puts it into $\mathcal{C}(m^k)$ for input size m . The \mathcal{C} circuit for input length $m = (|x| + \log t + \log d) \in O(n)$ is the oblivious witness circuit for n -length inputs.

The moreover part: For $L \in \text{ZTIME}(t)$, construct the same verifier V' for the language $L' = \{0, 1\}^*$ as in the proof of Theorem 20. As $L' \in \text{UTIME}(t)$, V' will have witness in $\mathcal{C}(n^k)$. Now a $\text{DTIME}(2^{n^{k+1}})$ algorithm for L is: for n -length input x , go through all the circuits in $\mathcal{C}(n^k \log n)$ one at a time, compute their truth-tables tt , and then compute $V'(x, tt)$. Due to the way V' is constructed, all of its positive certificates have the same first bit. If V accepts on any tt whose first bit is 1, then $x \in L$. Else $x \notin L$.

Proofs of (2) & (3): They are analogous to the proof of (1), except that they use Corollary 3.1.1.

Proof of (4): Let $L \in \text{UEXP}/a$, and V/a be an unambiguous (given the correct advice) verifier V for L that runs in time $O(2^{n^c})$ for some constant c . Since $\exists k L \in \text{SIZE}(n^k)$, from the proof of part (3) we know that V/a has witnesses in $\text{SIZE}(n^{ck})$ for some constant k .

Using this we first give an EXP/a algorithm for L . On n -length input x , go through all the circuits in $\text{SIZE}(n^{ck} \log n)$ one at a time, compute their truth-tables tt , and then compute $V(x, tt)/a$. Accept if V/a accepts for any tt , else reject. This is an EXP/a algorithm as simulation of V/a needs the original advice.

Once we get $\text{UEXP}/a = \text{EXP}/a$, $\text{EXP}/a \subseteq \text{SIZE}(\text{poly})$ gives $\text{UEXP}/a = \text{MA}/a$ [55]. \square

3.3 EWL and KLT for ZUTIME

We extend the techniques from the previous section to give similar results for ZUTIME. The main difference in the proof of our search to decision reduction is that, we adapt our definition of $L_{ewl(V)}$ to capture seeds of zero-error non-deterministic verifiers. First let's define this adaptation. For zero-error non-deterministic verifier V for language L :

$$L_{ewl(V)} = \{(x, i) \mid \exists y [V(x, y) \in \{0, 1\} \wedge (i^{\text{th}}\text{-bit of } y \text{ is } 1) \wedge \forall (z <_{l.o.} y) V(x, z) = ?]\} \quad (3.3)$$

The difference is that $L_{ewl(V)}$ captures the lexicographically first certificate that gives the

correct answer (doesn't matter whether the answer is 1 or 0). Once the search to decision reduction is established, the EWL and KLT follow from similar arguments as in the previous section. Here again we see why it might be difficult to extend these results to all ambiguous zero-error verifiers.

Theorem 22. *For time-constructible $t \in 2^{O(n)}$, and constants c and k :*

1. *For $L \in \text{ZUTIME}(t)$ and zero-error unambiguous verifier V for L , $L_{\text{ewl}(V)} \in \text{ZUTIME}(t(n))$ (where n is the input size for L). Moreover if this statement is true for all $O(t)$ -time zero-error non-deterministic verifiers (ambiguous and unambiguous) for every $\text{ZUTIME}(t)$ language, then $\text{ZNTIME}(t) = \text{ZUTIME}(t)$.*
2. *$\text{ZUTIME}(t) \subset C(n^k) \implies \text{ZUTIME}(t) \subset_{\text{os}} C(n^k)$. Moreover if this statement is true for all $O(t)$ -time zero-error non-deterministic verifiers (ambiguous and unambiguous) for every $\text{ZUTIME}(t)$ language, even just for seeds (let alone oblivious-seeds), then $\text{ZNTIME}(t) \subseteq \text{DTIME}(2^{n^{k+1}})$.*
3. *$\text{ZUTIME}(2^{n^c}) \subset C(n^k) \implies \text{ZUTIME}(2^{n^c}) \subset_{\text{os}} C(n^{ck})$.*
4. *$\text{ZUEXP} \subset \text{SIZE}(\text{poly}) \implies \text{ZUEXP} = \text{MA}$.*

Proof. Proof of (1): Algorithm for $L_{\text{ewl}(V)}$: For input (x, i) , guess a certificate y and simulate $V(x, y)$. Output '?' if V outputs '?'. Output 1 if V outputs in $\{0, 1\}$ and the i^{th} -bit of y is 1. Output 0 if V outputs in $\{0, 1\}$ and the i^{th} -bit of y is 0. This algorithm is correct and zero-error unambiguous as V is zero-error unambiguous. It runs in time $O(t(|x|))$.

The moreover part: For $L \in \text{ZNTIME}(t)$, let V be its $\text{ZNTIME}(t)$ verifier. Consider the $\text{ZUTIME}(t)$ language $L' = \{0, 1\}^*$.

Using V we first construct a $\text{ZNTIME}(t)$ verifier V' for L' : V' ignores the first bit of the certificate and simulates V using the rest of the certificate. V' outputs '?' if V outputs '?', it outputs 1 if V outputs in $\{0, 1\}$ and its output matches the first bit of the certificate.

Now using a $\text{ZUTIME}(t(n))$ algorithm \mathcal{A} for $L'_{\text{ewl}(V')}$ we give a $\text{ZUTIME}(t)$ algorithm for L : on input x , simulate \mathcal{A} on $(x, 1)$. Output whatever \mathcal{A} outputs.

Proofs of (2), (3) & (4): They are analogous to the proofs in Theorem 21. □

3.4 EWL and KLT for $\widetilde{\text{FewTIME}}$

We use the following folklore result to translate our results for Boolean circuits to any typical circuit class.

Lemma 3.4.1. *If $P \subset C$, then there exists a constant c such that: for large enough n , any s -size circuit has an equivalent s^c -size C circuit.*

Proof. $\text{Ckt} - \text{Eval}$ is a problem in P whose input is a Boolean circuit C and a string x , and the output is the output of C on x . If $P \subset C$, then there is a constant c such that $\text{Ckt} - \text{Eval}$ has $n^{c/2}$ -size C circuits.

Let B be a P/poly circuit of size s . Let E be $(n + s \log s)^{c/2}$ -size circuit corresponding to the $(n + s \log s)^{\text{th}}$ -slice of $\text{Ckt} - \text{Eval}$. Define $D(x) = E(B, x)$. It is easy to check that: (i) D is an s^c -size C circuit; and (ii) D is equivalent to B . \square

Now we give the EWL and KLT for $\widetilde{\text{FewE}}$.

Theorem 23. *For constant $k \geq 1$:*

1. $\widetilde{\text{FewE}}/(a+n) \subset C(n^k) \implies \exists k' \widetilde{\text{FewE}}/a \subset_w C(n^{k'})$
2. $\widetilde{\text{FewE}}/(a+n) \subset \text{SIZE}(n^k) \implies \widetilde{\text{FewE}}/a \subseteq \text{MA}/a$

Proof. Proof of (1): We prove the result for unrestricted Boolean circuits. The result for the circuit class C follows from the Lemma 3.4.1. The assumption implies $P \subset C$, thus any $\text{SIZE}(n^k)$ circuit has an equivalent $C(n^{ck'})$ circuit, for some constant c .

Contradiction: $\widetilde{\text{FewE}}/(a+n) \subset \text{SIZE}(n^k)$ implies $\text{EXP} \subset \text{SIZE}(\text{poly})$ and thus $\text{EXP} = \text{MA}$ [9]. Now we show that, if $\forall k' \geq 1 \widetilde{\text{FewE}}/a \not\subset_w \text{SIZE}(n^{k'})$, then $\text{MA} \subset \text{io-}\widetilde{\text{FewE}}/(a+n)$. Combined with the above statement it leads to the contradiction $\text{EXP} \subset \text{io-SIZE}(n^k)$ (since we can diagonalize against fixed-polynomial size circuits in EXP).

Hardness tester: $\forall k \geq 1 \widetilde{\text{FewE}}/a \not\subset_w \text{SIZE}(n^k)$ implies that for every $k \geq 1$, there is a $2^{\sqrt{n}}$ -time $\widetilde{\text{FewE}}/a$ verifier V_k/a that has infinite set of inputs S_k that it accepts, and for $x \in S_k$

and $2^{\sqrt{|x|}}$ -length certificate y such that $V_k(x, y)/a = 1$ (using the correct advice), the constraint $ckt(y) \geq n^k$ is true (where y is truth-table of a \sqrt{n} -input circuit). If V_k with its original a amount of advice, is also given elements of S_k as advice (one element per input length, and all 0s string for the input lengths for which S_k contains no element), V_k becomes a $\widetilde{\text{FewTIME}}(2^{\sqrt{n}})/(a+n)$ hardness tester.

Derandomization: For $L \in \text{MA}$, we derandomize the MA protocol for L using the above described $\widetilde{\text{Few}}$ verifiers. After including the non-determinism of Merlin into Arthur's input, let the size of the circuit C that captures the BP computation of Arthur for L be bounded by n^l (for some constant l). We use the verifier V_k for $k = lg$, where g is the constant from Theorem 19. Our algorithm guesses a 2^n -bit string Y and simulates V_k on Y , using the $a+n$ amount of advice as described above. It rejects if V_k rejects, else it uses the certificate that V_k accepted. Note that, infinitely often the accepted certificates Y will satisfy $ckt(Y) \geq n^{lg}$. We use the certificates to construct a PRG $G : n \rightarrow n^l$ using the Theorem 19, that fools n^l -size circuits. We brute-force through the seeds of G to compute the acceptance probability of the circuit C in $2^{O(n)}$ -time (within $\pm 1/n^l$ error). If the acceptance probability is greater than $1/2$, our algorithm accepts, else it rejects. The running time of our non-deterministic algorithm is bounded by 2^n , and the number of accepting branches is bounded by $2^{2^{(\log n)^2/4}} \times 2^{n^l}$, which is less than $2^{2^{(\log n)^2}}$ for large enough n .

Proof of (2): $\widetilde{\text{FewE}}/(a+n) \subset \text{SIZE}(n^k)$ combined with (1) gives $\widetilde{\text{FewE}}/a \subset_w \text{SIZE}(n^{k'})$ for some constant k' . This gives $\widetilde{\text{FewE}}/a \subset \text{EXP}/a$: by brute-forcing through the truth-tables of all $\text{SIZE}(n^{k'})$ circuits to find accepting certificates (if there are any). Finally we get $\widetilde{\text{FewE}}/a \subseteq \text{MA}/a$ since $\text{EXP}/a = \text{MA}/a$ by [55]. \square

Chapter 3 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Chapter 4

UEXP Lower Bounds from Fast

Unambiguous Algorithms

First we show how to get UEXP lower bounds from fast unambiguous algorithms for canonization and tautology (Section 4.1). Then we show how to replace canonization and tautology by Π_2 SAT and get more fine-grained results (Section 4.2). Then we show how to completely get rid of canonization for the case $\widetilde{\text{FewE}}$ lower bounds (Section 4.3). Then we show how to simulate canonization using proper learning (Section 4.4). Finally, we show how to generalize certain lower bound frameworks for unrestricted Boolean circuits, to typical circuit classes, even when the algorithms are very slow (Section 4.5). We use the following hierarchy for semantic classes in our proofs.

Theorem 24 (Hierarchy for Semantic Classes [34]). *For any time bound t such that $n \leq t \leq 2^n$, there is a constant $\varepsilon > 0$ and an advice bound $a \in O(\log(t) \log(\log(t)))$ such that $\text{UTIME}(t)/a \not\subseteq \text{UTIME}(t^\varepsilon)/(a+1)$ (resp. $\widetilde{\text{FewTIME}}(t)/a \not\subseteq \widetilde{\text{FewTIME}}(t^\varepsilon)/(a+1)$).*

4.1 Lower bounds from unambiguous tautology and canonization algorithms

We use the following ‘tight reductions to 3 – USAT’ in this section.

Theorem 25 (Efficient local reductions [51, 90, 32]). *Every language $L \in \text{UTIME}(2^n)$ can be reduced to 3 – USAT (uniquely satisfiable 3 – SAT) instances of $2^n n^c$ -size, for some constant c . Moreover, given an instance of L there is an n^c -size C (P-uniform) circuit that, on an integer $i \in [2^n n^c]$ in binary as input, outputs the i^{th} -clause of the resulting 3 – USAT formula.*

We first formally define canonization and related notations.

Canonization : A subset S of circuits is called $\text{CAN}_{(s,C,p)}$, if for any s -size C circuit C , there exists a unique circuit $C' \in S$ with $tt(C) = tt(C')$, and $|C'| \leq p(\text{ckt}_C(tt(C')))$. $\text{CAN}_{(s,C,p)} \in \Gamma/a$ means there is a Γ/a algorithm that decides $\text{CAN}_{(s,C,p)}$.

$\text{TAUT}_{(s,C)}$ (resp. $\text{SAT}_{(s,C)}$) denotes the TAUT (resp. SAT) for s -size C circuits.

In these definitions we omit, the parameter s when it is $\text{poly}(n)$, and the circuit class when $C = \text{Boolean}$.

The main idea is to guess the witness circuit unambiguously using the canonization algorithm, and then combine the witness circuit with the reduction circuit in the same manner that Williams did [99]. The existence of the witness circuit follows from the UTIME EWL.

Theorem 26. *For $\delta \leq 1$, let a, c and ϵ be the parameters of Theorems 24 and 25 for the time bound $t = 2^{\delta n}$. Then for constant k and function $p(n) \geq n$, $\text{UTIME}(2^{\delta n})/a \not\subseteq C(n^k)$ if:*

1. $\text{TAUT}_{(p(n^{k+1})n+n^c, C)} \in \text{UTIME}(2^{\epsilon n})$ and $\text{CAN}_{(n^{k+1}, C, p)} \in \text{UTIME}(2^{\epsilon n})/1$; or
2. $\text{TAUT}_{(p(n^{k+1})n+n^c, C)} \in \text{UTIME}(2^{\epsilon n})/1$ and $\text{CAN}_{(n^{k+1}, C, p)} \in \text{UTIME}(2^{\epsilon n})$.

Proof. Using the assumptions (1 or 2), we will contradict the UTIME hierarchy (Theorem 24) by designing a $\text{UTIME}(2^{\epsilon n})/(a+1)$ algorithm for arbitrary $L \in \text{UTIME}(2^{\delta n})/a$.

Reduction circuit: For $L \in \text{UTIME}(2^{\delta n})/a$ and input x , let F_x be the $2^n n^c$ -size 3 – USAT formula we get by reducing from x (Theorem 25). There is an n^c -size (P-uniform) C circuit D with $n + c \log n$ input wires, that outputs the i^{th} -clause of F when given the input $i \in [1, 2^n n^c]$.

Special verifier: Let V be the verifier for L that first reduces input x to the 3 – USAT formula F_x , and then non-deterministically guesses a satisfying assignment for F_x .

Easy-witness circuit: From UTIME EWL (Theorem 21) and the assumption $\text{UTIME}(2^{\delta n})/a \subset C(n^k)$ we know that V has witness circuits in $C(n^k)$. Let E be a witness circuit of this verifier for the input length $|x| = n$.

Final circuit C : Combining D and E we construct a circuit C that satisfies: “ C is a tautology $\iff x \in L$ ”. On input i , the output of D is $3n + 3c \log n + 3$ bits long. The first $3n + 3c \log n$ bits are the three variables of the i^{th} -clause of F . Plug these output bits to three separate copies of E . The last three bits indicate whether the corresponding literals are positive or negative. Use these three bits and the three output bits from the three copies of E to compute the value of the i^{th} -clause (based on the assignment encoded by $tt(E)$).

Contradicting the first assumption: Non-deterministically guess a $p(n^{k+1})$ -size C circuit E . Simulate the $\text{CAN}_{(n^{k+1}, C, p)}$ algorithm on E . This requires $\text{UTIME}(2^{\epsilon n})/1$. Reject if the answer is negative. Continue if it’s positive, and construct C as described above. $|C| \leq p(n^{k+1})n + n^c$. Note that, for any truth-table only one non-deterministic branch will lead to a non-rejecting path. Now simulate the $\text{TAUT}_{(p(n^{k+1})n + n^c, C)}$ algorithm on C . This requires $\text{UTIME}(2^{\epsilon n})$. Note that, C is accepted if and only if, $x \in L$, and $tt(E)$ is the unique witness of V . This whole process requires the advice used in the $\text{UTIME}(2^{\delta n})/a$ algorithm for L . So we get a $\text{UTIME}(2^{\epsilon n})/(a + 1)$ algorithm.

Contradicting the second assumption: The algorithm is exactly the same, expect that the extra 1-bit of advice is used by the tautology algorithm, and not by the canonization algorithm. \square

We get the following corollary that is cleaner in presentation.

Corollary 4.1.1. $\text{UE}/O(n \log n) \not\subset C$, if $\text{TAUT}_C \in \text{USUBE}$ and $\text{CAN}_{(C, p)} \in \text{USUBE}$, for $p(n) \in \text{poly}(n)$.

4.2 Lower bounds from unambiguous Π_2 SAT algorithms

Here the idea is to simulate canonization using a Π_2 SAT algorithm.

Theorem 27. *For every constant k_1 there is a constant k_2 , such that if Π_2 SAT on n variables and n clauses can be solved in $\text{UTIME}(2^{n/(\log n)^{k_2}})$, then $\text{UE}/n \not\subseteq \text{SIZE}(n(\log n)^{k_1})$.*

Proof. From the Theorem 24 we know that there is an $a \leq n$ such that $\text{UTIME}(2^{n/(\log n)^2})/a \not\subseteq \text{UTIME}(2^{n/(\log n)^3})/(a+1)$. $\text{UE}/n \subset \text{SIZE}(n(\log n)^{k_1})$ gives $\text{UTIME}(2^{n/(\log n)^2})/a \subset \text{SIZE}(n(\log n)^{k_1})$. From the UTIME EWL we get $\text{UTIME}(2^{n/(\log n)^2})/a \subset_w \text{SIZE}(n(\log n)^{k_1})$.

Now the proof is similar to the proof of Theorem 26, except few changes. We prove that any $L \in \text{UTIME}(2^{n/(\log n)^2})/a$ has an $\text{UTIME}(2^{n/(\log n)^3})/a$ algorithm. After guessing a $\text{SIZE}(n(\log n)^{k_1})$ witness circuit E , we use a fast Π_2 SAT algorithm on it (from our assumption) to make sure that we move forward unambiguously. We check that, for all lexicographically small (in some fixed encoding scheme) circuits D , there is at least one input z , such that $E(z) \neq D(z)$. For the final circuit C , we use three copies of E and a reduction circuit that is (P-uniform) linear-size (the n^c -size circuit in the Theorem 54 can be made linear [51]). Finally we run a fast tautology algorithm (from our assumption) on the circuit C . □

4.3 Lower bounds from $\widetilde{\text{Few}}$ tautology algorithms

The idea is that there can only be exponential many possibilities for the witness circuit, and thus the number of positive non-deterministic branches of the final algorithm remain within the limits of a $\widetilde{\text{Few}}$ verifier.

Theorem 28. $\text{TAUT}_C \in \widetilde{\text{FewSUBE}} \implies \forall k \widetilde{\text{FewE}}/O(n \log n) \not\subseteq C(n^k)$.

Proof. For the sake of contradiction, assume that $\exists k \widetilde{\text{FewE}}/O(n \log n) \subset C(n^k)$. From the Theorem 23 ($\widetilde{\text{FewTIME}}$ EWL) we get that $\exists k' \widetilde{\text{FewE}}/a \subset_w C(k')$, where a is the advice parameter of the Theorem

24 for time bound $t = 2^n$. Using the fast $\widetilde{\text{Few}}$ algorithm from our assumption we contradict the Theorem 24 by showing that any arbitrary $L \in \text{FewE}/a$ has an $\widetilde{\text{FewSUBE}}/(a+1)$ algorithm.

Now again our proof follows the structure of Theorem 26 with some modifications. We construct the final circuit C without using any canonization. There can only be $2^{O(n^{k'})}$ many witness circuits E . For designing a $\widetilde{\text{FewTIME}}(2^{n^\epsilon})/a$ algorithm for any $0 < \epsilon < 1$, we use a $\widetilde{\text{FewTIME}}(2^{n^\delta})$ tautology algorithm for some $\delta \leq \epsilon$. This makes the number of accepting paths for the final algorithm at most $2^{O(n^{k'})} \times 2^{(\delta n)^{\log(\delta n)}} \leq 2^{(\epsilon n)^{\log(\epsilon n)}}$. \square

4.4 Lower bounds from unambiguous learning and tautology algorithms

We first show how exact (proper) learning along with tautology algorithm imply canonization. Then we plug this connection in the Theorem 26 to get lower bounds from learning and tautology algorithms. Before we give our result, we first formally define the UTIME exact learning algorithm that we use in our results.

Exact UTIME learning with membership and equivalence queries: Let s be the size of the target concept C (the circuit to be learned). A $\text{UTIME}(t)$ algorithm is called $\text{LRN}_{(s,C,p)}$, if for any s -size C circuit C , it outputs a circuit C' of size at most $p(s)$ in time at most $t(n)$ (where n is the number of input wires) with $tt(C) = tt(C')$, on exactly one of its non-deterministic branches, and rejects all the other branches. The algorithm is allowed to make “membership” and “equivalence” queries. A membership query is of the type: “What is the value of $C(x)$?”. An equivalence query is of the type: “Is the current hypothesis (H) equal to C ?”. On any positive equivalence query, it halts and outputs the current hypothesis. On any negative query, it gets x from the oracle, such that $H(x) \neq C(x)$.

If the output, and the equivalence queries are all polynomial-size C circuits, the algorithm is called $\text{P-LRN}_{(s,C,p)}$ (proper learning).

Here again, we omit the size parameter when $s(n) = \text{poly}(n)$, and the circuit class when $C = \text{Boolean}$. Here we omit $p(n)$ too, if it is $\text{poly}(n)$. Unlike in $\text{CAN}_{(C,p)}$, in $\text{LRN}_{(C,p)}$ p decides the size of the output (and not the input).

Theorem 29. *For any polynomial $p(n)$:*

1. $\text{P-LRN}_{(C,p)} \in \text{UTIME}(t) \wedge \text{TAUT}_C \in \text{UTIME}(t') \implies \text{CAN}_{(C,p)} \in \text{UTIME}(t(t' + \text{poly}(n)))$
2. $\text{P-LRN}_{(C,p)} \in \text{USUBE} \wedge \text{TAUT}_C \in \text{USUBE} \implies \text{UE}/O(n \log n) \not\subseteq C$

Proof. Proof of (1): In an exact proper learning algorithm, if we have access to the circuit C that we are learning, then we can get a canonization algorithm for C (because the learning algorithm only cares about the truth-table of the circuit that it is learning, and outputs the same hypothesis for all the circuits that have same truth-tables). As the final hypothesis will be of size at most $p(s)$ for s -size C circuits, we get a UTIME algorithm for $\text{CAN}_{(C,p)}$. The membership queries can be handled directly since we have the circuit with us. For the equivalence queries, we non-deterministically guess the faith for the hypothesis H .

If we guess $H \equiv C$, we use the tautology algorithm to verify it. If our guess is wrong, we reject. If our guess is right, we accept if only if, H 's description is same as C .

If we guess $H \not\equiv C$, we have to output an input z , such that $H(z) \neq C(z)$. We try to guess the lexicographically smallest such z to keep the whole process unambiguous. After guessing z , we check that $H(z) \neq C(z)$ and use our tautology algorithm to check that $\forall (z' <_{l.o.} z) H(z') = C(z')$, where l.o. stands for lexicographical ordering (the test $z' <_{l.o.} z$ can be encoded by any typical circuit of linear size). We return z if both the checks pass, else we reject.

Proof of (2): We get this directly from (1) and the Corollary 4.1.1. □

4.5 Generalization of lower bound frameworks

In the above sections we saw that fast UTIME algorithms for certain circuit analysis algorithms for C circuits were fed to certain frameworks to yield lower bounds for UTIME against C . Consider

the scenario where: a framework is altogether different, or is a fine-grained version of one of the current ones (in terms of size and depth of the circuits and running time of the algorithms), and works for Boolean circuits, but not for some restriction \mathcal{C} . Also consider that, the assumptions of these frameworks are satisfied for that \mathcal{C} , but not for unrestricted Boolean circuits. Do we get any lower bounds? In this section we prove that this question has a positive answer.

We use win-win type arguments analogous to the ones used in [72] for fast NTIME algorithms. We show that, either $P \not\subseteq \mathcal{C}$ (i.e., a stronger lower bound exists against \mathcal{C}), or fast unambiguous algorithms for \mathcal{C} circuits imply fast unambiguous algorithms for Boolean circuits (i.e., frameworks that only work for Boolean circuits can now be used). To prove our results, we use the Lemma 3.4.1.

Theorem 30. *Either $P \not\subseteq \mathcal{C}$, or $\exists c$, for $p(m) = m^k$ for any $k \geq 1$, and $t, t', t'' \leq 2^n$:*

1. $\text{CAN}_{(\mathcal{C}, p)} \in \text{UTIME}(t) \implies \text{CAN}_{p^c} \in \text{UTIME}(t)$
2. $\text{CAN}_{(\mathcal{C}, p)} \in \text{UTIME}(t) \wedge \text{TAUT}_{\mathcal{C}} \in \text{UTIME}(t') \implies \text{TAUT} \in \text{UTIME}((t + t')\text{poly}(n))$
3. $\text{CAN}_{(\mathcal{C}, p)} \in \text{UTIME}(t) \wedge \text{TAUT}_{\mathcal{C}} \in \text{UTIME}(t') \wedge \text{SAT}_{\mathcal{C}} \in \text{UTIME}(t'')$
 $\implies \text{SAT} \in \text{UTIME}((t + t')\text{poly}(n) + t'')$

Proof. If $P \subset \mathcal{C}$, from the Lemma 3.4.1 we know there exists a constant c such that: for each s -size Boolean circuit B , there is an equivalent s^c -size \mathcal{C} circuit C (for large enough n).

Proof of (1): By a simple modification of an algorithm \mathcal{A} for $\text{CAN}_{(\mathcal{C}, p)}$, we obtain an algorithm \mathcal{A}' for CAN_{p^c} . On input B , the algorithm \mathcal{A}' first checks whether B belongs to \mathcal{C} . It rejects if the answer is negative. If the answer is positive it simulates \mathcal{A} on B and accepts if and only if \mathcal{A} accepts.

Proof of (2): Let \mathcal{A} be a $\text{UTIME}(t)$ algorithm for $\text{CAN}_{(\mathcal{C}, p)}$, \mathcal{A}' be a $\text{UTIME}(t')$ algorithm for $\text{TAUT}_{\mathcal{C}}$. Using \mathcal{A} and \mathcal{A}' , we construct a UTIME algorithm \mathcal{A}'' for TAUT .

For input B to \mathcal{A}'' , for each gate g of B , let B_g be the circuit corresponding to the output wire of gate g . For the output gate o , \mathcal{A}'' first guesses an equivalent \mathcal{C} circuit C'_o . To make sure that its

guess is unambiguous, it simulates \mathcal{A} on C'_o and rejects if \mathcal{A} rejects. Then it simulates \mathcal{A}' on C'_o (to check if C'_o is a tautology) and rejects if it rejects. The only thing left to check is that C'_o is actually equivalent to C_o .

For checking the consistency of C'_o , \mathcal{A}'' first guesses C circuit C'_g , for each gate g . It then simulates \mathcal{A} on each C'_g and rejects if \mathcal{A} rejects on any of them. Finally it simulates \mathcal{A}' on C''_g for each g , where C''_g is the circuit that captures the tautology “ $C'_g = op(C'_{g_1}, \dots, C'_{g_l})$ ” for $g = op(g_1, \dots, g_l)$. It accepts if and only if \mathcal{A} accepts on all of them.

Proof of (3): For input B , with the same strategy as in the proof of 2, we first unambiguously construct an equivalent C circuit C . Then, on this C we simulate a $UTIME(t'')$ algorithm for SAT_C . □

For the case of \widetilde{Few} algorithms, we get the following theorem where we don't need canonization. The proof is same as of the above theorem, except that now we can skip all the canonization steps. This change will not increase the number of positive non-deterministic branches of the final algorithm by much, and thus the constraints of a \widetilde{Few} verifier are not violated.

Theorem 31. *Either $P \not\subseteq C$, or:*

1. $TAUT_C \in \widetilde{FewTIME}(2^n/n^{\omega(1)}) \implies TAUT \in \widetilde{FewTIME}(2^n/n^{\omega(1)})$
2. $TAUT_C \in \widetilde{FewTIME}(2^n/n^{\omega(1)}) \wedge SAT_C \in \widetilde{FewTIME}(2^n/n^{\omega(1)}) \implies SAT \in \widetilde{FewTIME}(2^n/n^{\omega(1)})$

Chapter 4 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Chapter 5

Unique Properties vs. Lower Bounds

In this chapter we establish relationships between different types of unique properties and lower bounds against $UTIME$, $ZUTIME$, $NTIME$ and $ZNTIME$. Main results of this chapter are summarized in the Table 5.1.

Table 5.1: Properties vs Lower Bounds

Properties useful against \mathcal{C}	Witness LB	HS LB	Ob-witness LB	Set LB
$P/\log n - u_{=1} \equiv P-U$	$ZUE \not\subseteq_s \mathcal{C}$	$ZUE \not\subseteq_{hs} \mathcal{C}$	$ZUE \not\subseteq_{os} \mathcal{C}$	$ZUE \not\subseteq \mathcal{C}$
$P/\log n - u_{\leq 1}$	$UE \not\subseteq_w \mathcal{C}$			
		$UE \not\subseteq_{hw} \mathcal{C}$		
			$UE \not\subseteq_{ow} \mathcal{C}$	$UE \not\subseteq \mathcal{C}$
$P/\log n - U$	$UE/n \not\subseteq_w \mathcal{C}$	$UE/n \not\subseteq_{hw} \mathcal{C}$	$UE/n \not\subseteq_{ow} \mathcal{C}$	$UE/n \not\subseteq \mathcal{C}$
$NP/\log n - u_{=1} \equiv NP-U$				$ZNE \not\subseteq \mathcal{C}$
$NP-N \equiv P-N$	$ZNE \not\subseteq_s \mathcal{C}$	$ZNE \not\subseteq_{hs} \mathcal{C}$	$ZNE \not\subseteq_{os} \mathcal{C}$	
$NP/\log n - U \equiv NP/\log n - N$ $\equiv P/\log n - N$	$NE \not\subseteq_w \mathcal{C}$	$NE \not\subseteq_{hw} \mathcal{C}$	$NE \not\subseteq_{ow} \mathcal{C}$	$NE \not\subseteq \mathcal{C}$
$NP-prU$				$ip-ZNE \not\subseteq \mathcal{C}$
$NP-prN \equiv P-prN$	$ip-ZNE \not\subseteq_s \mathcal{C}$	$ip-ZNE \not\subseteq_{hs} \mathcal{C}$	$ip-ZNE \not\subseteq_{os} \mathcal{C}$	

Lower bounds in any particular column are all equivalent, and lower bounds from any column imply the lower bounds in the column just below it (except the columns that are separated by

double lines). Note that, as the lower bounds get weaker, the properties become less restrictive (or the constructivity goes higher). Also note that, at any place in the Table 5.1 we can remove the $\log n$ advice by assigning each advice a different input-length. But then the property no more remains a property, instead gets converted into an useful algorithm (see [101]). An useful algorithm, unlike an useful property, accepts inputs of all lengths and not just powers of 2. It appends zeros on the inputs that are not powers of 2, to make it a truth-table. We stick to properties in our presentation.

In almost all of our connections we use the following connection (Lemma 5.0.1) between UP-U and P-U properties. The proof of this connection is along the same lines as the original connection [4, 72, 101]: an useful NP (*resp.* RP-natural) property yields an useful P (*resp.* P-natural) property.

Lemma 5.0.1. *UP/a property \mathcal{U} can be converted into a P/a property \mathcal{P} such that:*

1. \mathcal{U} is UP/a-U property $\implies \mathcal{P}$ is P/a-U property;
2. \mathcal{U} is UP/a- $u_{=1}$ property $\implies \mathcal{P}$ is P/a- $u_{=1}$ property;
3. \mathcal{U} is UP/a- $u_{\leq 1}$ property $\implies \mathcal{P}$ is P/a- $u_{\leq 1}$ property;
4. \mathcal{U} is useful against $C \implies \mathcal{P}$ is useful against C .

Proof. Let V be the unambiguous verifier corresponding to \mathcal{U} 's algorithm. Let c be a constant such that $2^{cn} - 2^n$ is the length of the certificates that V guesses for the inputs of size 2^n . Now we design \mathcal{P} which satisfies the promises of the theorem statement. For m which is not a multiple of c , among all the inputs of length 2^m , \mathcal{P} only accepts the all 0s string. For $m = cn$ for some n , for any input xy where $|x| = 2^n$ and $|y| = 2^{cn} - 2^n$, \mathcal{P} simulates V on (x, y) , and accepts if and only if V accepts. For any $n \in \mathbb{N}$, \mathcal{P} uses the same advice for 2^{cn} -size inputs, that \mathcal{U} uses for 2^n -size inputs.

Proofs of (1), (2) & (3): The construction of \mathcal{P} ensures this for the inputs of size 2^m , where m is not a multiple of c . For all the other input sizes this is ensured by the fact that \mathcal{U} is a UP property, and the behavior of \mathcal{U} on different advice strings. For any $n \in \mathbb{N}$, and any advice string, the number of 2^{cn} -size inputs \mathcal{P} accepts, is same as the number of 2^n -size inputs \mathcal{U} accepts.

Proof of (4): If \mathcal{U} is useful against \mathcal{C} , then for each k there exists an infinite subset S_k such that for each $n \in S_k$, $\mathcal{U}(x) = 1 \implies \text{ckt}_{\mathcal{C}}(x) > n^k$. For any x , let y be the unique certificate such that $V(x, y) = 1$. Since $\text{ckt}_{\mathcal{C}}(x) > n^k \implies \text{ckt}_{\mathcal{C}}(xy) > n^k \geq (cn)^{k-1}$ for each k , \mathcal{P} is useful against n^{k-1} -size \mathcal{C} circuits for each k , and hence is useful against \mathcal{C} . \square

5.1 ZUE lower bounds vs P-U properties

Theorem 32. [Row 1 of Table 5.1] *The following statements are equivalent:*

1. $\text{ZUE} \not\subseteq \mathcal{C}$
2. $\text{ZUE} \not\subseteq_{os} \mathcal{C}$
3. $\text{ZUE} \not\subseteq_{hs} \mathcal{C}$
4. $\text{ZUE} \not\subseteq_s \mathcal{C}$
5. $\text{P-U} \not\subseteq_{tt} \mathcal{C}$
6. $\text{P}/\log n\text{-u}_{=1} \not\subseteq_{tt} \mathcal{C}$

Proof. (1) \implies (5) Let $L \in \text{ZUE} \setminus \mathcal{C}$, and let V be $2^{O(n)}$ -time zero-error unambiguous verifier for L . For any n , L_n can be viewed as a function f_n , where $f_n^{-1}(1) = \{x \in L \mid |x| = n\}$.

Now using V we give a UP-U property \mathcal{U} that is useful against \mathcal{C} . Then the result follows from the Lemma 5.0.1.

For any input y of length 2^n , \mathcal{U} simulates V on all the n -length strings, one by one. For each i , it matches the i^{th} bit of y , and the output of V on the i^{th} n -length string. \mathcal{U} accepts if and only if it succeeds in all 2^n verifications.

Constructivity & uniqueness: For $n \in \mathbb{N}$, \mathcal{U} unambiguously accepts the truth table of function f_n , and rejects all the other strings. As it runs for $2^{O(n)}$ -time on 2^n -length inputs, it is UP-U (as V is ZUE).

Usefulness: As $L \notin C$, for each k , there are infinitely many input lengths n , such that f_n doesn't have n^k -size C circuits. Thus \mathcal{U} is useful against C .

(5) \implies (4) Let \mathcal{P} be a P-unique property useful against C . Using \mathcal{P} we construct a zero-error unambiguous verifier V for the ZUE language $\{0, 1\}^*$ such that V doesn't have seeds in C .

For any n -length input x , V guesses a string y of length 2^n and accepts if and only if \mathcal{P} accepts y . Since \mathcal{P} is P-unique property useful against C , the unique accepting witnesses of V are not in C .

(4) \implies (3) This follows from the definitions.

(3) \implies (2) This follows from the definitions.

(2) \implies (1) The contrapositive follows from the ZUTIME EWL (Theorem 22).

(5) \iff (6) The forward direction is trivial. For the reverse direction, for P/ $\log n$ -u=1 property \mathcal{P} , we convert \mathcal{P} to a P-U property \mathcal{P}' .

For odd m , among all the inputs of length 2^m , \mathcal{P}' only accepts the all 0s string. For $m = 2n$ for some n , for any 2^m -length input $x_1 x_2 \dots x_{2^{2n}}$ where $\forall i |x_i| = 2^n$, \mathcal{P}' accepts if and only if for each i : \mathcal{P} accepts input x_i with the advice y_i (i^{th} n -length string in the lexicographical order).

Constructivity & uniqueness: These both follow for \mathcal{P}' directly from the fact that \mathcal{P} is a strong-unique P-property.

Usefulness: If \mathcal{P} is useful against C with advice sequence $\{a_n\}_{n \in \mathbb{N}}$, then for each k there exists an infinite subset S_k such that for each $n \in S_k$, $\mathcal{P}(x, a_n) = 1 \implies \text{ckt}_C(x) > n^k$. Among all the 2^{2n} -length strings, let y be the unique string that \mathcal{P}' accepts. $y = x_1 \dots x_{b_n} \dots x_{2^{2n}}$, where $\forall i \in [1, 2^{2n}] |x_i| = 2^n$, b_n is the lexicographical rank of a_n among all the n -length strings, and $x = x_{b_n}$. Since $\text{ckt}_C(x) > n^k \implies \text{ckt}_C(y) > n^k \geq (2n)^{k-1}$ for each k , \mathcal{P}' is useful against n^{k-1} -size C circuits for each k , and hence is useful against C . \square

5.2 UE lower bounds vs $P/\log n-u_{\leq 1}$ properties

We use a fine-grained version of the techniques from [101], to prove the following two theorems: (i) “witness lower bound $\iff P/\log n-u_{\leq 1}$ useful property” and (ii) “oblivious witness lower bound \iff UE lower bound”. Unfortunately, the “oblivious witness lower bound \implies witness lower bound” connection of NTIME doesn’t go through in the case of UTIME. If we try to establish a “oblivious witness lower bound $\implies P/\log n-u_{\leq 1}$ useful property” connection, we don’t get a mild-unique property, but just a unique property. In the next section we will see that this connection can be established in the presence of advice.

Theorem 33. [Row 2 of Table 5.1] *The following statements are equivalent:*

1. $UE \not\subseteq_w C$
2. $P/\log n-u_{\leq 1} \not\subseteq_{tt} C$

Proof. (1) \implies (2) Let L' be a UE language, and V' be an unambiguous verifier for L' that doesn’t have witnesses in C . By a simple padding argument we can construct $L \in UTIME(2^n)$, and an unambiguous verifier V for L with certificate length 2^n , that doesn’t have witnesses in C .

If the inputs are given as advice, and the certificates are given as inputs, then V becomes a $P/\log n$ property \mathcal{P} , that is useful against C .

For \mathcal{P} to be a $u_{\leq 1}$ property, it should be unique with respect to the same advice that makes it useful. At this point, all we know is that for every input length and every advice, \mathcal{P} accepts at most one truth-table (since V is unambiguous). The advice that makes \mathcal{P} useful may not be present for all input lengths. For some of these input lengths n where no such advice is present, it is also possible that L_n is empty (i.e. no advice is present that makes the property non-empty).

We will be done if L is non-empty for all input lengths. Consider the two modifications of V : (i) V_0 , that changes its behavior on the all 0s strings and always accepts them (unambiguously); and (ii) V_1 , that changes its behavior on the all 1s strings and always accepts them (unambiguously). The modified languages and their corresponding verifiers are also $UTIME(2^n)$, and have 2^n -length

certificates. We show that at least one of these two modifications doesn't have witness in \mathcal{C} . If V_0 has witnesses in \mathcal{C} , then V 's witnesses corresponding to the all 0s strings must be the ones that were not in \mathcal{C} (at least infinitely often), so then V_1 doesn't have witnesses in \mathcal{C} (infinitely often).

(2) \implies (1) Let \mathcal{P} be a $\mathbb{P}/\log n - u_{\leq 1}$ property that is useful against \mathcal{C} . Define $L = \{x \mid \exists y \mathcal{P}(y)/x = 1\}$. Let V be the verifier for L , that on any n -length input x , guesses a string y of length 2^n , and simulates \mathcal{P} on y using x as advice. V is a UE verifier since \mathcal{P} is a mild-unique property. V doesn't have witnesses in \mathcal{C} since \mathcal{P} is useful against \mathcal{C} . \square

Theorem 34. [Row 4 of Table 5.1] *The following statements are equivalent:*

1. $\text{UE}/a \not\subseteq \mathcal{C}$
2. $\text{UE}/a \not\subseteq_{\text{ow}} \mathcal{C}$

Proof. $\neg(1) \implies \neg(2)$ This follows from the UTIME EWL (Theorem 21).

$\neg(2) \implies \neg(1)$ Assume that UE/a has oblivious witnesses (for all verifiers that are unambiguous given the correct advice) in \mathcal{C} . Let L be a UE/a language and V/a be a UE/a verifier for L . By our assumption, V/a has oblivious witnesses in n^k -size \mathcal{C} circuits. Now we show that $L \in \mathcal{C}$.

Using V/a we construct an unambiguous verifier V'/a for the UEXP/a language $\{0, 1\}^*$ such that: for $n \in \mathbb{N}$, an oblivious witness circuit for V'/a on n -length inputs, computes L_n .

For $x \in L$ (this can be verified by brute forcing through all the n^{k+1} -size circuits), $V'(x, y)/a = 1$ only when y is the all 1s string. For $x \notin L$, $V'(x, y)/a = 1$ only when y is the all 0s string. Since UE/a has oblivious witnesses in \mathcal{C} , by a simple padding argument UEXP/a too has oblivious witnesses in \mathcal{C} . Let $\{C_n\}_{n \in \mathbb{N}}$ be the \mathcal{C} circuit family encoding the oblivious witnesses of V' . Then, the \mathcal{C} circuit family defined by $D_n(x) = C_n(x, 1)$ encodes the language L (since the first bit of the unique accepting certificate of V'/a on input x , dictates whether $x \in L$ or not). \square

5.3 UE/ n lower bounds vs P/ $\log n$ -U properties

The arguments from Section 5.1, when extended to the advice setting, circumvent the problems from Section 5.2, and yield the following theorem.

Theorem 35. [Row 5 of Table 5.1] *The following statements are equivalent for any constant $k \geq 1$:*

1. $\text{UE}/n^k \not\subseteq \mathcal{C}$
2. $\text{UE}/n^k \not\subseteq_{\text{ow}} \mathcal{C}$
3. $\text{UE}/n^k \not\subseteq_{\text{hw}} \mathcal{C}$
4. $\text{UE}/n^k \not\subseteq_w \mathcal{C}$
5. $\text{P}/(\log n)^k\text{-U} \not\subseteq_{\text{tt}} \mathcal{C}$

Proof. (1) \implies (5) Let $L \in \text{UE}/n^k \setminus \mathcal{C}$, and let V be $2^{O(n)}$ -time unambiguous verifier for L . For any n , L_n can be viewed as a function f_n , where $f_n^{-1}(1) = \{x \in L \mid |x| = n\}$.

Now using V we give a $\text{UP}/\log^k m\text{-U}$ property \mathcal{U} that is useful against \mathcal{C} . Then the result follows from the lemma 5.0.1.

For odd m , among all the inputs of length 2^m , \mathcal{U} only accepts the all 0s string. For $m = 2n$ for some n , for any 2^m -length input yz with $|y| = 2^n$ and $|z| = 2^{2n} - 2^n$, \mathcal{U} goes through all the n -length strings, one by one. If the i^{th} bit of y is 0, it does nothing. If the i^{th} bit of y is 1, it simulates V on the i^{th} n -length string in the lexicographical order (to verify its inclusion in L). The first n^k bits of advice is the advice required for the simulation of V . The rest of the $(\log 2^{2n})^k - n^k = (2n)^k - n^k \geq n$ bits of advice encodes the number of n -length inputs that V accepts. \mathcal{U} accepts if and only if: (i) it succeeds in all 2^n verifications; (ii) the hamming weight of y is equal to the number encoded by the last $(2n)^k - n^k$ bits of advice; and (iii) z is an all 0s string.

Constructivity & uniqueness: For $n \in \mathbb{N}$, \mathcal{U} unambiguously accepts the truth table of function f_n (followed by an all 0s string of length $2^{2n} - 2^n$), and rejects all the other strings. As it runs for $2^{O(n)}$ -time for 2^n -length inputs with n^k -size advice, it is $\text{UP}/(\log n)^k\text{-U}$ (as V is UE).

Usefulness: As $L \notin C$, for each l , there are infinitely many input lengths n , such that f_n doesn't have n^{l+1} -size C circuits. Corresponding to each such n , for the inputs of length 2^{2n} , \mathcal{U} accepts strings y that doesn't have $(2n)^l$ -size C circuits because: any $(2n)^l$ -size circuit C with $tt(C) = y$, decides L_n after we fix the first half of its input wires to 1s, and $(2n)^l \leq n^{l+1}$.

(5) \implies (4) Let \mathcal{P} be a $P/(\log n)^k$ -U property useful against C . We construct an unambiguous verifier V for the UE/n^k language $\{0, 1\}^*$, that doesn't have witnesses in C . For any n -length input x , guess a 2^n -length string y and simulate \mathcal{P} on y , and accept if and only if \mathcal{P} accepts.

Since \mathcal{P} is useful against C , V doesn't have witnesses in C . As \mathcal{P} is unique, V is UE/n^k .

(4) \implies (3) This follows from the definitions.

(3) \implies (2) This follows from the definitions.

(2) \implies (1) The contrapositive follows from the UTIME EWL (Theorem 21). □

5.4 ZNE lower bounds vs NP-U properties

In [72] it was conjectured, “ $ZNE \not\subseteq C \iff \exists$ P-N (or NP-N) property useful against C ” while only forward direction was proved. We use a fine-grained version of the proof to establish the equivalence in the case of unique properties. So if this conjecture is true, then any NP-N property has an equivalent NP-U property.

In the Section 5.5 we establish the equivalence between: NP-N properties and lower bounds for ZNE seeds. Since NP-U properties imply NP-N properties, this result can be viewed as an reverse-EWL for ZNE. Moreover, if the conjecture of [72] is true, then we also get an EWL for ZNE.

In the Section 5.6 we show an equivalence between NP-N and NP-U properties, when they are allowed to use $\log n$ amount of advice. This equivalence uses the EWL for NE. In the chapter 7 we reduce this advice to $O(1)$, on the expense of making the lower bounds fixed-polynomial. This gives us an EWL and reverse-EWL for $ZNE/O(1)$ for fixed-polynomial upper/lower bounds.

Theorem 36. [Row 6 of Table 5.1] *The following statements are equivalent:*

1. $\text{ZNE} \not\subseteq \mathcal{C}$

2. $\text{NP}/\log n\text{-}u=1 \not\subseteq_{tt} \mathcal{C}$

3. $\text{NP-U} \not\subseteq_{tt} \mathcal{C}$

Proof. (2) \iff (3) The reverse direction is trivial. For the forward direction, we convert any $\text{NP}/\log n\text{-}u=1$ property \mathcal{P} into an NP-U property \mathcal{P}' . The conversion from the proof of Theorem 32 works for NP properties as well.

(1) \implies (3) Let $L \in \text{ZNE} \setminus \mathcal{C}$, and let V be $2^{O(n)}$ -time zero-error non-deterministic verifier for L . For any n , L_n can be viewed as a function f_n , where $f_n^{-1}(1) = \{x \in L \mid |x| = n\}$.

Now using V we give an NP-U property \mathcal{U} that is useful against \mathcal{C} . For any input y of length 2^n , \mathcal{U} simulates V on all the n -length strings, one by one. If the i^{th} bit of y is 0, it verifies the inclusion of the i^{th} (lexicographically) n -length string in \bar{L} . If the i^{th} bit of y is 1, it verifies the inclusion of the i^{th} (lexicographically) n -length string in L . \mathcal{U} accepts if and only if it succeeds in all 2^n verifications.

Constructivity & uniqueness: For $n \in \mathbb{N}$, \mathcal{U} accepts the truth table of function f_n , and rejects all the other strings. As it runs for $2^{O(n)}$ -time on 2^n -length inputs, it is NP-U.

Usefulness: As $L \notin \mathcal{C}$, for each k , there are infinitely many input lengths n , such that f_n doesn't have n^k -size \mathcal{C} circuits. Thus \mathcal{U} is useful against \mathcal{C} .

(3) \implies (1) Let \mathcal{U} be an NP-unique property useful against \mathcal{C} . We construct a language L in $\text{ZNE} \setminus \mathcal{C}$, whose ZNE verifier uses \mathcal{U} .

For any n -length input x , V guesses a string y of length 2^n and simulates \mathcal{U} on it. Let the lexicographical rank of x (among all n -bit strings) be i . V outputs '?' if \mathcal{U} rejects, else it proceeds further. It outputs 1, if y 's i^{th} -bit is equal to 1. Else it outputs 0.

Since \mathcal{U} is NP-unique property, for each n the 2^n -length string y_n it accepts is unique. Thus V accepts the language whose slices are represented by the strings y_n (let's call this language L), and satisfies the promises of a ZNE verifier.

Since \mathcal{U} is useful against \mathcal{C} , for each k , there are infinitely many values of n where y_n doesn't have n^k -size \mathcal{C} circuits. Thus $L \notin \mathcal{C}$. \square

The above proof also gives an equivalence between NP-promise-unique properties and lower bounds against ip-ZNE (the promise version of ZNE).

Theorem 37. [Row 9 of Table 5.1] *The following statements are equivalent:*

1. $\text{ip-ZNE} \not\subseteq \mathcal{C}$
2. $\text{NP-prU} \not\subseteq_{tt} \mathcal{C}$

Proof Idea: (\implies) Any $L \in \text{ip-ZNE} \setminus \mathcal{C}$ that satisfies the ZNE or $(\text{NE} \cap \text{Co} - \text{NE})$ -promise on n -length inputs for some n , yields a property that is unique on the input lengths 2^n . Since L satisfies the lower bound on the promise inputs, the property is useful on the inputs on which it is unique.

(\impliedby) Any NP-unique useful property \mathcal{U} that is unique on 2^n -length inputs for some n , yields a ZNE verifier that satisfies the ZNE-promise on n -length inputs. Since useful inputs of \mathcal{U} also satisfy the promise, the verifier satisfies the lower bound on the promise inputs. \square

5.5 ZNE lower bounds vs NP-N properties

Theorem 38. [Row 7 of Table 5.1] *The following statements are equivalent:*

1. $\text{ZNE} \not\subseteq_{os} \mathcal{C}$
2. $\text{ZNE} \not\subseteq_{hs} \mathcal{C}$
3. $\text{ZNE} \not\subseteq_s \mathcal{C}$
4. $\text{NP-N} \not\subseteq_{tt} \mathcal{C}$
5. $\text{P-N} \not\subseteq_{tt} \mathcal{C}$

Proof. (4) \iff (5) It is proved in [72].

(3) \implies (2) This follows from the definitions.

(2) \implies (1) This follows from the definitions.

(1) \implies (5) Let V be a ZNE verifier that doesn't have oblivious seeds in \mathcal{C} . Let V 's certificate length be 2^{cn} , for some constant c . Using V we construct a P-N property \mathcal{P} useful against \mathcal{C} . View any $2^{(c+1)n}$ -length input as a collection of 2^n certificates. \mathcal{P} accepts if and only if, for each $i \in [1, 2^n]$, V outputs in $\{0, 1\}$ on the i^{th} (lexicographically) n -length input when given the i^{th} certificate from the collection. Clearly, \mathcal{P} is a P-N property. It is useful against \mathcal{C} as it only accepts oblivious witnesses of V .

(5) \implies (3) Let \mathcal{P} be a P-N property useful against \mathcal{C} . For each k , let S_k be the infinite set of inputs where \mathcal{P} only accepts strings str with $ckt_{\mathcal{C}}(str) \geq n^k$. Using \mathcal{P} we construct a ZNE verifier V for $\{0, 1\}^*$ that doesn't have seeds in \mathcal{C} . For n -length input x , V guesses a string y of length 2^n . V outputs 1, if \mathcal{P} accepts the string y . Otherwise V outputs '?'. For each k , for any n with $2^n \in S_k$, due to the way it is constructed, V doesn't have seeds in n^k -size \mathcal{C} . Thus V doesn't have seeds in \mathcal{C} . \square

The above proof also gives an equivalence between NP-promise properties and lower bounds for ip-ZNE seeds. The proof idea is similar to the one given for the Theorem 37.

Theorem 39. [Row 10 of Table 5.1] *The following statements are equivalent:*

1. $ip\text{-ZNE} \not\subseteq_{os} \mathcal{C}$
2. $ip\text{-ZNE} \not\subseteq_{hs} \mathcal{C}$
3. $ip\text{-ZNE} \not\subseteq_s \mathcal{C}$
4. $NP\text{-prN} \not\subseteq_{tt} \mathcal{C}$
5. $P\text{-prN} \not\subseteq_{tt} \mathcal{C}$

5.6 NE lower bounds vs NP/log n-U properties

Theorem 40. [Row 8 of Table 5.1] *The following statements are equivalent:*

1. $NE \not\subseteq C$
2. $NE \not\subseteq_{ow} C$
3. $NE \not\subseteq_{hw} C$
4. $NE \not\subseteq_w C$
5. $NP/\log n\text{-U} \not\subseteq_{tt} C$
6. $NP/\log n\text{-N} \not\subseteq_{tt} C$
7. $P/\log n\text{-N} \not\subseteq_{tt} C$

Proof. *Equivalence of (1), (2), (4), (6) & (7):* It is proved in [101, 72].

Equivalence with (3): The implications, (4) \implies (3) and (3) \implies (2), follow from the definitions.

Equivalence with (5): The implication (5) \implies (6) follows from the definitions. We will now show the implication (1) \implies (5).

Let $L \in NE \setminus C$, and let V be $2^{O(n)}$ -time non-deterministic verifier for L . For any n , L_n can be viewed as a function f_n , where $f_n^{-1}(1) = \{x \in L \mid |x| = n\}$.

Now using V we give an NP/log n-U property \mathcal{U} that is useful against C . For any input y of length 2^n , \mathcal{U} goes through all the n -length strings, one by one. If the i^{th} bit of y is 0, it does nothing. If the i^{th} bit of y is 1, it simulates V on the i^{th} n -length string in the lexicographical order (to verify its inclusion in L). \mathcal{U} accepts if and only if it succeeds in all 2^n verifications and the hamming weight of y is equal to the number encoded by the advice. If the advice is equal to the size of L_n , \mathcal{U} accepts the truth table corresponding to the function f_n , and rejects all the other strings. Since it runs in $2^{O(n)}$ -time on 2^n -length inputs, it is NP/log n-U. As $L \notin C$, \mathcal{U} is useful against C . \square

Chapter 5 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Chapter 6

Derandomization Using Unique Properties

Here we extend the “lower-bounds to derandomization” connection to UEXP and ZUEXP. We use the following two connections from the Table 5.1:

1. $ZUE \not\subseteq_s C \iff P-U \not\subseteq_{tt} C$
2. $UE \not\subseteq_w C \iff P/\log n-U \not\subseteq_{tt} C$

The idea is to obtain unique properties from UEXP and ZUEXP lower bounds, and then use these properties to unambiguously obtain hard functions, which then yield the desired derandomization.

Theorem 41. *[Unambiguous derandomization from UEXP and ZUEXP lower bounds]*

1. $ZUEXP \neq EXP \implies BPP \subset \bigcap_{\epsilon > 0} io-ZUTIME(2^{n^\epsilon})$
2. $ZUEXP \not\subseteq SIZE(poly) \implies BPP \subset \bigcap_{\epsilon > 0} io-ZUTIME(2^{n^\epsilon})$
3. $ZUEXP \neq MA \implies BPP \subset \bigcap_{\epsilon > 0} io-ZUTIME(2^{n^\epsilon})$
4. $ZUEXP \neq BPP \implies BPP \subset \bigcap_{\epsilon > 0} io-Heur-ZUTIME(2^{n^\epsilon})$
5. $UEXP \neq EXP \implies BPP \bigcap_{\epsilon > 0} \subset io-ZUTIME(2^{n^\epsilon})/n^\epsilon$

$$6. \text{ UEXP} \not\subset \text{SIZE}(\text{poly}) \implies \text{BPP} \subset \bigcap_{\epsilon > 0} \text{io-ZUTIME}(2^{n^\epsilon})/n^\epsilon$$

$$7. \text{ UEXP} \neq \text{MA} \implies \text{BPP} \subset \bigcap_{\epsilon > 0} \text{io-ZUTIME}(2^{n^\epsilon})/n^\epsilon$$

$$8. \text{ UEXP} \neq \text{BPP} \implies \text{BPP} \subset \bigcap_{\epsilon > 0} \text{io-Heur-ZUTIME}(2^{n^\epsilon})/n^\epsilon$$

Proof. Proof of (1): Let's assume that $\text{ZUEXP} \neq \text{EXP}$. Then ZUE can't have seeds in $\text{SIZE}(\text{poly})$, because brute-forcing through the seeds will prove $\text{ZUEXP} = \text{EXP}$. Thus, there exists a P-U property \mathcal{P} useful against $\text{SIZE}(\text{poly})$ (from the Table 5.1). For each c , let S_c be the infinite set of input lengths where \mathcal{P} only accept strings str satisfying $\text{ckt}(str) \geq n^c$. These strings are truth-tables of hard functions, and can be computed in UE using the constructivity of \mathcal{P} .

For $k, \epsilon > 0, \epsilon' > \epsilon'/2$ and $L \in \text{BPTIME}(n^{k/2})$, set $c = gk/\epsilon'$ (where g is the constant from Theorem 19). We give a $\text{ZCTIME}(2^{n^\epsilon})$ algorithm for L that works for any input length n with $2^n \in S_c$. For n -length input x of L , let C_x be the $\text{SIZE}(n^k)$ circuit capturing the BP computation of L .

Non-deterministically guess a string Y of length $m = 2^{n^{\epsilon'}}$. Output '?' if $\mathcal{P}(Y) = 0$. Once we have access to Y with $\mathcal{P}(Y) = 1$ (or $\text{ckt}(Y) \geq n^k$), we can construct PRG $G : n^\epsilon \rightarrow n^k$ from Y (using the Theorem 19) that is computable in E. We brute-force through all the $n^{\epsilon'}$ -length seeds, and on each of the output strings of length n^k , compute the circuit C_x to calculate its acceptance probability in time 2^{n^ϵ} (within $\pm 1/n^k$ error). Output 1 if this value is $1/2$ or more, else output 0. Since $\mathcal{P}(Y) = 1$ holds for unique Y , the whole process is unambiguous.

Proofs of (2): We prove the contrapositive. Assume that $\exists \epsilon > 0$ such that $\text{BPP} \not\subset \text{io-ZUTIME}(2^{n^\epsilon})$. This gives us $\text{EXP} \subset \text{SIZE}(\text{poly})$ [9, 70, 71], and $\text{ZUEXP} = \text{EXP}$ from (1). Thus, $\text{ZUEXP} \subset \text{SIZE}(\text{poly})$.

Proof of (3): Using (1), (2) and EXP KLT we get a series of implications that conclude the

proof.

$$\begin{aligned} \text{ZUEXP} \neq \text{MA} &\implies \text{ZUEXP} \neq \text{EXP} \text{ or } \text{EXP} \neq \text{MA} \\ &\implies \text{BPP} \subset \bigcap_{\epsilon > 0} \text{io-ZUTIME}(2^{n^\epsilon}) \text{ or } \text{EXP} \not\subset \text{SIZE}(\text{poly}) \\ &\implies \text{BPP} \subset \bigcap_{\epsilon > 0} \text{io-ZUTIME}(2^{n^\epsilon}) \text{ or } \text{ZUEXP} \not\subset \text{SIZE}(\text{poly}) \\ &\implies \text{BPP} \subset \bigcap_{\epsilon > 0} \text{io-ZUTIME}(2^{n^\epsilon}) \end{aligned}$$

Proof of (4): Its the same as above, except $\text{EXP} \neq \text{MA}$ is replaced with $\text{EXP} \neq \text{BPP}$ and the derandomization from [49] is used.

Proof of (5): It's analogous to the proof of (1), except that the property we get is $\text{P}/\log n$ and not P constructive. The $\log n$ -bit advice for this property is precisely the n^ϵ -bit advice for the $\text{ZUTIME}(2^{n^\epsilon})$ algorithm we get.

Proofs of (6), (7) & (8): They are analogous to the proofs of (2), (3) and (4). The advice from the proof of (5) travels to them as well. \square

Chapter 6 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Chapter 7

Isolation of properties: EWL & KLT for ZNE

In this chapter we discuss the consequences of isolating properties with different constructivity. By isolation we mean: extracting (*resp.* proving existence of) an useful unique property from (*resp.* from the existence of) an arbitrary useful property. From the Table 5.1 we know that:

1. Isolation of P-properties is equivalent to: $ZUE \subset C \iff ZNE \subset_{os} C$.
2. Isolation of P/ $\log n$ -properties is equivalent to: $UE/n \subset C \iff NE \subset C$.
3. Isolation of NP-properties is equivalent to: $ZNE \subset C \iff ZNE \subset_{os} C$.
4. Isolation of NP-promise-properties is equivalent to: $ip-ZNE \subset C \iff ip-ZNE \subset_{os} C$.
5. Isolation of NP/ $\log n$ -properties was already achieved in the Theorem 40.

In this chapter we focus on the points (3) and (4). For the case of fix-polynomial lower bounds: we merge the rows 6 and 7 (in presence of $O(1)$ advice), and rows 9 and 10 of the Table 5.1. Most of the equivalences follow from the arguments from the previous chapter: Theorems 36, 37, 38 and 39. The main technical results of this chapter are the implications: (i) $\forall k \geq 1 ZNE \not\subset_{os} C(n^k) \implies \forall k \geq 1 ZNE/1 \not\subset C(n^k)$; (ii) $\forall k \geq 1 ZNE \not\subset_{os} io-C(n^k) \implies \forall k \geq 1 ZNE \not\subset C(n^k)$; and (iii) $\forall k \geq 1 ip-ZNE \not\subset_{os} C(n^k) \implies \forall k \geq 1 ip-ZNE \not\subset C(n^k)$. Contrapositive of these can be viewed

as EWLs for ZNE. Using these EWLs we derive KLTs for ZNE, and isolation results for NP-properties. We also use the Lemma 3.4.1 to make our EWLs work for typical circuit classes.

First we prove the ZNE EWL and KLT.

Theorem 42. *For constant $k \geq 1$:*

1. $\text{ZNE}/1 \subset \mathcal{C}(n^k) \implies \exists k' \geq 1 \text{ ZNE} \subset_{os} \mathcal{C}(n^{k'})$
2. $\text{ZNE}/O(1) \subset \mathcal{C}(n^k) \implies \exists k' \geq 1 \text{ ZNE}/O(1) \subset_{os} \mathcal{C}(n^{k'})$
3. $\text{ZNE} \subset \mathcal{C}(n^k) \implies \exists k' \geq 1 \text{ ZNE} \subset_{os} \text{io-}\mathcal{C}(n^{k'})$
4. $\text{ip-ZNE} \subset \mathcal{C}(n^k) \implies \exists k' \geq 1 \text{ ZNE} \subset_{os} \mathcal{C}(n^{k'})$
5. $\text{ip-ZNE} \subset \mathcal{C}(n^k) \implies \exists k' \geq 1 \text{ ip-ZNE} \subset_{os} \mathcal{C}(n^{k'})$

Proof. We prove these results for the unrestricted Boolean circuits. The result for the circuit class \mathcal{C} follows from the Lemma 3.4.1. All the assumptions imply $\text{P} \subset \mathcal{C}$, thus any $\text{SIZE}(n^k)$ circuit has an equivalent $\mathcal{C}(n^{ck'})$ circuit, for some constant c .

Proof of (1): $\text{ZNE}/1 \subset \text{SIZE}(n^k)$ implies $\text{EXP} \subset \text{SIZE}(poly)$, and thus $\text{EXP} = \text{MA} \cap \text{Co-MA}$ [9].

Now we show that, if $\forall k' \geq 1 \text{ ZNE} \not\subset_{os} \text{SIZE}(n^{k'})$, then $\text{MA} \cap \text{Co-MA} \subset \text{io-ZNE}/1$. Combined with the above statement it leads to the contradiction $\text{EXP} \subset \text{io-SIZE}(n^k)$ (since we can diagonalize against fix-polynomial size circuits in EXP).

$\forall k \geq 1 \text{ ZNE} \not\subset_{os} \text{SIZE}(n^k) \implies \forall k \geq 1 \text{ P-N} \not\subset_{H} \text{SIZE}(n^k)$ (the arguments from Theorem 38 apply to the fix-polynomial lower bounds as well). For any $L \in \text{MA} \cap \text{Co-MA}$: we derandomize the MA protocols for L and \bar{L} using the useful P-N properties to give an ZNE/1 algorithm that works infinitely often. For any constant p let \mathcal{N}_p be a P-N property useful against n^p -size circuits.

The ZNE/1 algorithm: After including the non-determinism of Merlin into Arthur's input: let the size of the circuit C (resp. C') that captures the BP computation of Arthur for L (resp. L') be bounded by n^l , for some constant l . We use the property \mathcal{N}_{lg} , where g is the constant from Theorem

19. The 1-bit of advice indicates whether the property is useful or not. If it's 0, the algorithm always outputs 0. If it's 1, the algorithm guesses a 2^n -bit string Y and simulates \mathcal{N}_g on Y . It outputs '?' if \mathcal{N}_g rejects Y , else it proceeds further and guesses another bit b . If $b = 0$: it derandomizes C' (after guessing Merlin's non-determinism) and outputs 0 if the acceptance probability is $\geq 1/2$, else outputs '?'. If $b = 1$: it derandomizes C (after guessing Merlin's non-determinism) and outputs 1 if the acceptance probability is $\geq 1/2$, else outputs '?'.

Derandomization: The property \mathcal{N}_g yields truth-tables that don't have n^l -size circuits, infinitely often. Once we have access to these truth-tables, we construct a PRG $G : n \rightarrow n^l$ using the Theorem 19, that fools n^l -size circuits. We brute-force through the seeds of G to compute the acceptance probability of the circuits C and C' in $2^{O(n)}$ -time (within $\pm 1/n^l$ error).

Proof of (2): It's same as (1), except we use the fact that the arguments from Theorem 38 also apply in the advice setting and yield: $\forall k \geq 1 \text{ ZNE}/O(1) \not\subseteq_{os} \text{SIZE}(n^k) \implies \forall k \geq 1 \text{ P}/O(1)\text{-N} \not\subseteq_{tt} \text{SIZE}(n^k)$. The extra 1-bit of advice used to indicate the usefulness of the property during derandomization, now hides in the $O(1)$ advice.

Proof of (3): It's same as (1), except we use the fact that the arguments from Theorem 38 also yield: $\forall k \geq 1 \text{ ZNE} \not\subseteq_{os} \text{io-SIZE}(n^k) \implies \forall k \geq 1 \text{ P-N} \not\subseteq_{tt} \text{io-SIZE}(n^k)$. That is, if we start with a ZNE verifier V whose oblivious-seeds have high circuit complexity on all input lengths, then we get a P-N property that is useful everywhere. Now, when we derandomize any $L \in \text{MA} \cap \text{Co-MA}$, we don't need that one bit of advice.

Proof of (4): It's same as (1), except if we don't use that 1-bit of advice to encode the usefulness of the property during derandomization, we get an ip-ZNE algorithm. The ZNE-promise is met only when the property is useful.

Proof of (5): Since in (4) we don't use advice to encode the usefulness, we might as well use promise property. So we use the Theorem 39 instead, to get: $\forall k \geq 1 \text{ ip-ZNE} \not\subseteq_{os} \text{SIZE}(n^k) \implies \forall k \geq 1 \text{ P-prN} \not\subseteq_{tt} \text{SIZE}(n^k)$. □

Using the above EWL we give the following KLT.

Theorem 43. For constant $k \geq 1$:

1. $\text{ZNE}/1 \subset \text{SIZE}(n^k) \implies \text{ZNE} \subset \text{MA}$
2. $\text{ZNE}/O(1) \subset \text{SIZE}(n^k) \implies \text{ZNE}/O(1) \subset \text{MA}/O(1)$
3. $\text{ZNE} \subset \text{SIZE}(n^k) \implies \text{ZNE} \subset \text{io-MA}$
4. $\text{ip-ZNE} \subset \text{SIZE}(n^k) \implies \text{ZNE} \subset \text{MA}$
5. $\text{ip-ZNE} \subset \text{SIZE}(n^k) \implies \text{ip-ZNE} \subset \text{MA}$

Proof Idea: All the assumptions give $\text{EXP} \subset \text{SIZE}(\text{poly})$ or $\text{EXP}/O(1) \subset \text{SIZE}(\text{poly})$. This gives $\text{EXP} = \text{MA}$ or $\text{EXP}/O(1) = \text{MA}/O(1)$ from [9]. From the Theorem 42, all these assumptions give circuit upper bounds on the oblivious-seeds of ZNE, ZNE/O(1) or ip-ZNE. Brute-forcing through these circuits that encode the seeds, we get collapses to EXP, EXP/O(1) or io-EXP. \square

Using the above EWL we also give the following isolation results.

Theorem 44. For constant $k \geq 1$:

1. $\forall k \geq 1 \text{ NP-N} \not\subseteq_{tt} \mathcal{C}(n^k) \implies \forall k \geq 1 \text{ NP}/1\text{-U} \not\subseteq_{tt} \mathcal{C}(n^k)$
2. $\forall k \geq 1 \text{ NP}/O(1)\text{-N} \not\subseteq_{tt} \mathcal{C}(n^k) \iff \forall k \geq 1 \text{ NP}/O(1)\text{-U} \not\subseteq_{tt} \mathcal{C}(n^k)$
3. $\forall k \geq 1 \text{ NP-N} \not\subseteq_{tt} \text{io-}\mathcal{C}(n^k) \implies \forall k \geq 1 \text{ NP-U} \not\subseteq_{tt} \mathcal{C}(n^k)$
4. $\forall k \geq 1 \text{ NP-N} \not\subseteq_{tt} \mathcal{C}(n^k) \implies \forall k \geq 1 \text{ NP-prU} \not\subseteq_{tt} \mathcal{C}(n^k)$
5. $\forall k \geq 1 \text{ prN-N} \not\subseteq_{tt} \mathcal{C}(n^k) \iff \forall k \geq 1 \text{ NP-prU} \not\subseteq_{tt} \mathcal{C}(n^k)$

Proof Idea: The result follows if we replace the hypotheses and the conclusions, in the contrapositive of these statements, by equivalent hypotheses and conclusions from the Theorems 36, 37, 38 and 39. \square

The points (2) and (5) of the above theorem also yield the following more general result.

Theorem 45. *The following statements are equivalent:*

1. $\forall k \geq 1 \text{ ZNE}/O(1) \not\subseteq C(n^k)$ (*resp.* $\forall k \geq 1 \text{ ip-ZNE} \not\subseteq C(n^k)$)
2. $\forall k \geq 1 \text{ ZNE}/O(1) \not\subseteq_{os} C(n^k)$ (*resp.* $\forall k \geq 1 \text{ ip-ZNE} \not\subseteq_{os} C(n^k)$)
3. $\forall k \geq 1 \text{ ZNE}/O(1) \not\subseteq_{hs} C(n^k)$ (*resp.* $\forall k \geq 1 \text{ ip-ZNE} \not\subseteq_{hs} C(n^k)$)
4. $\forall k \geq 1 \text{ ZNE}/O(1) \not\subseteq_s C(n^k)$ (*resp.* $\forall k \geq 1 \text{ ip-ZNE} \not\subseteq_s C(n^k)$)
5. $\forall k \geq 1 \text{ NP-U}/O(1) \not\subseteq_{tt} C(n^k)$ (*resp.* $\forall k \geq 1 \text{ NP-prU} \not\subseteq_{tt} C(n^k)$)
6. $\forall k \geq 1 \text{ NP-N}/O(1) \not\subseteq_{tt} C(n^k)$ (*resp.* $\forall k \geq 1 \text{ NP-prN} \not\subseteq_{tt} C(n^k)$)
7. $\forall k \geq 1 \text{ P-N}/O(1) \not\subseteq_{tt} C(n^k)$ (*resp.* $\forall k \geq 1 \text{ P-prN} \not\subseteq_{tt} C(n^k)$)

Chapter 8

Lower bounds against prSV non-deterministic circuits

Here we discuss all the results regarding lower bounds against prSV non-deterministic circuits. First we derive the EWL for the case of NEXP and $(\text{NP} \cap \text{Co-NP})/\text{poly}$ (Section 8.1) and use it to derive the connection between non-trivial GAP-SAT algorithm and the lower bound $\text{NEXP} \not\subseteq (\text{NP} \cap \text{Co-NP})/\text{poly}$ (Section 8.2). Then we derive new gap theorems for MA and CAPP (Section 8.3). Then we derive connections between fast algorithms and NE and E^{NP} lower bounds against circuits that use limited amount of prSV non-determinism (Section 8.4) and use that connection to derive unconditional lower bounds (Section 8.5). Finally we show unconditional lower bounds against fixed-polynomial prSV non-deterministic circuits (Section 8.6).

8.1 NEXP vs $(\text{NP} \cap \text{Co-NP})/\text{poly}$

Here we give NEXP EWL and KLT for $(\text{NP} \cap \text{Co-NP})/\text{poly}$, and the converses. We also extend the results to $\text{E}_{\parallel}^{\text{NP}}$. These results work even if replace $\text{NP} \cap \text{Co-NP}$ with P (in circuit classes and as oracles).

Theorem 46. *The following statements are equivalent:*

1. $NE \not\subseteq MA^{NP \cap Co-NP}$
2. $E_{||}^{NP} \not\subseteq MA^{NP \cap Co-NP}$
3. $E_{||}^{NP} \not\subseteq (NP \cap Co-NP)/poly$
4. $NE \not\subseteq (NP \cap Co-NP)/poly$
5. $NE \not\subseteq_w (NP \cap Co-NP)/poly$
6. $NE \not\subseteq_{hw} (NP \cap Co-NP)/poly$
7. $NE \not\subseteq_{ow} (NP \cap Co-NP)/poly$
8. $prMA^{NP \cap Co-NP} \subset \bigcap_{\epsilon > 0} io-NTIME(2^{n^\epsilon})/n^\epsilon$

Proof. (1) \implies (2) This is trivial.

(2) \implies (3) For the sake of contradiction, assume that $E_{||}^{NP} \not\subseteq MA^{NP \cap Co-NP}$ and $E_{||}^{NP} \subset (NP \cap Co-NP)/poly$. The latter implies $EXP = AM = MA^{NP \cap Co-NP}$. Thus the former implies $E_{||}^{NP} \not\subseteq EXP$ and $\exists k NE/O(n) \subset NSIZE(n^k)$ (using a linear time NE-complete language). $AM \subset io-NE/O(n)$, and these implications, gives us the contradiction $\exists k EXP \subset io-NSIZE(n^k)$.

We show $AM \subset io-NE/O(n)$ by using a language $L \in E_{||}^{NP}$, such that $L \notin EXP$ (which again follows from the assumptions). For any $E_{||}^{NP}$ algorithm \mathcal{A} deciding L , for any constant k , for infinitely many n , there can't be any NP-oracle n^k -size circuits encoding the witnesses for all the positive oracle queries that \mathcal{A} makes on all n -length inputs. This is because, brute-forcing through such circuits will give an EXP algorithm for L . Now using \mathcal{A} we get an $NE/O(n)$ algorithm \mathcal{B} , that for each k , for infinitely many n , produces $2^{O(n)}$ -length strings tt with $ckt^{NP}(tt) > n^k$. The advice encodes the number of positive oracle queries that \mathcal{A} makes on that input length. For any n , \mathcal{B} simulates \mathcal{A} on all n -length inputs and using advice guesses that many queries to be positive. It verifies its guesses by non-deterministically guessing certificates for the positive oracle queries. After all the verification steps, it outputs the concatenation of all its non-deterministic certificates. This concatenated string can't have NP-oracle n^k -size circuits, for any n where its sub-strings that

represent the positive oracle queries of \mathcal{A} doesn't have NP-oracle n^{k+1} -size circuits (because a circuit for the whole string, can be projected down to get a circuit for any sub-string). Now \mathcal{B} gives us $\text{AM} \subset \text{io-NE}/O(n)$, using the hardness to derandomization connection from [58].

(3) \implies (4) This follows from the result $\text{E}_{\parallel}^{\text{NP}} \subset \text{NE}/O(n)$, where the advice gives the count of the number of positive oracle queries on all n -length inputs.

(4) \implies (5) Let $L \in \text{NE} \setminus (\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$. We construct a NE verifier V for the language $\{0, 1\}^*$ that doesn't have witnesses in $(\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$. V accepts any n -length string only on the 2^n -length witness that represents the characteristic function of L_n . Since $L \notin (\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$, witness of V are also not in $(\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$.

(5) \implies (6) This follows from the definitions.

(6) \implies (7) This follows from the definitions.

(7) \implies (8) The NE verifier V that doesn't have oblivious-witnesses in $(\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$, yields a function sequence computable in $\text{NE}/O(n)$ that, for any constant k , for any $(\text{NP} \cap \text{Co} - \text{NP})$ -oracle A , infinitely often, doesn't have A -oracle circuits of size n^k . The advice is used to encode the number of inputs that the NE verifier accepts, and output sequence is just the oblivious-witnesses of V . Moreover, by a simple padding argument, for any $\varepsilon > 0$, the function sequence can be computed in $\text{NTIME}(2^{n^\varepsilon})/n^\varepsilon$.

Any language $L \in \text{prMA}^{\text{NP} \cap \text{Co} - \text{NP}}$, has MA protocols where Arthur uses some $A \in \text{NP} \cap \text{Co} - \text{NP}$ as oracle. After including the non-determinism of Merlin into the input, Arthur's computation can be converted into an A -oracle n^d -size circuit C for some constant d . This conversion only takes NP. Now for any $\varepsilon > 0$, we derandomize these circuits for infinitely many input lengths n , in $\text{NTIME}(2^{n^\varepsilon})/n^\varepsilon$. This will establish $\text{prMA}^{\text{NP} \cap \text{Co} - \text{NP}} \subset \bigcap_{\varepsilon > 0} \text{io-NTIME}(2^{n^\varepsilon})/n^\varepsilon$.

For any input length n , we first compute the function from the function sequence that doesn't have $n^{gd/\varepsilon}$ -size A -oracle circuits, and then use that function and the Theorem 19, to construct a PRG $G : n^\varepsilon \rightarrow n^d$. This PRG fools n^d -size A -oracle circuits, and thus brute-forcing through its inputs, we can estimate the acceptance probability of C , and output accordingly.

(8) \implies (1) If $\text{NE} \subset \text{MA}^{\text{NP} \cap \text{Co} - \text{NP}}$ and $\text{MA}^{\text{NP} \cap \text{Co} - \text{NP}} \subset \bigcap_{\varepsilon > 0} \text{io-NTIME}(2^{n^\varepsilon})/n^\varepsilon$, then we get

$\text{EXP} = \text{NEXP} \subset \bigcap_{\epsilon > 0} \text{io-NTIME}(2^{n^\epsilon})/n^\epsilon$. This gives us $\text{EXP} \subset \bigcap_{\epsilon > 0} \text{io-TIME}(2^{n^\epsilon})/n$ for some constant c . This is false due to the diagonalization result given in [42]. \square

8.2 $\text{NEXP} \not\subseteq (\text{NP} \cap \text{Co-NP})/\text{poly}$ from super-polynomial savings

In this section we show that super-polynomial savings in non-deterministic algorithms for GAP-SAT for $(\text{NP} \cap \text{Co-NP})$ -oracle circuits, imply $\text{NEXP} \not\subseteq (\text{NP} \cap \text{Co-NP})/\text{poly}$. We first state the following PCP verifier for NEXP, and hierarchy theorem for NTIME, that we will need in our result.

Theorem 47 (see [11, 99]). *For any $L \in \text{NTIME}(2^n)$, there exists a PCP verifier $V(x, y, r)$ with soundness $1/2$, perfect completeness, randomness complexity $n + c \log n$, query complexity n^c , and verification time n^c , for some constant c . That means:*

- *V has random access to x and y , uses at most $|r| = n + c \log n$ random bits in any execution, makes n^c queries to the candidate proof y , and runs in at most n^c steps.*
- *if $x \in L$, $\exists y : |y| = n^c \Pr_r[V(x, y, r) = 1] = 1$.*
- *if $x \notin L$, $\forall y : |y| = n^c \Pr_r[V(x, y, r) = 1] \leq 1/2$.*

Theorem 48 (NTIME Hierarchy [104]). *Let t_1 and t_2 be time constructible functions that satisfy $t_1(n+1) \in o(t_2(n))$. There is a unary language in $\text{NTIME}(t_2(n))$ that is not in $\text{NTIME}(t_1(n))$.*

Now we prove our result. Recall that, a CAPP or tautology algorithm can also solve GAP-SAT.

Theorem 49. *For any super-polynomial function sp , an $\text{NTIME}(2^n/sp(n))$ GAP-SAT algorithm for n -input $\text{poly}(n)$ -size A -oracle circuits, for every $A \in (\text{NP} \cap \text{Co-NP})$, implies $\text{NEXP} \not\subseteq (\text{NP} \cap \text{Co-NP})/\text{poly}$.*

Proof. For $L \in \text{NTIME}(2^n)$ we design an $\text{NTIME}(2^n/sp(n))$ algorithm, under the assumption $\text{NEXP} \subset (\text{NP} \cap \text{Co-NP})/\text{poly}$. This will contradict the NTIME hierarchy from Theorem 48.

Reduction circuit: Let V be a PCP verifier for L from the Theorem 47. On any input x , $V(x, y, r)$ receives $|r| = n + c \log n$ random bits, makes oracle queries to the proof y of size $2^n n^c$, and runs for n^c -time. Let C_x be an oracle circuit capturing this computation. For the oracle gates, we will use copies of the following described easy-witness circuit B_x for a special verifier V' .

Special NE verifier: On input x and certificate y , $V'(x, y)$ computes $V(x, y, r)$ on each value of r and outputs 1 if and only if $\forall r V(x, y, r) = 1$.

Easy-witness circuit: Since $\text{NEXP} \subset (\text{NP} \cap \text{Co-NP})/\text{poly} \implies \text{NEXP} = \text{AM}$, from [42] we get that the search problem for V' is in EXP. Thus, there is an algorithm \mathcal{A} that: on any input $x \in L$ outputs y such that $V'(x, y) = 1$; on any input $x \notin L$ outputs an all zeros string. Now define a new language $L' = \{(x, i) \mid i^{\text{th}} \text{ output bit of } \mathcal{A} \text{ on input } x \text{ is } 1\}$. $L' \in \text{EXP}$ and thus $L' \in \text{P}^A/\text{poly}$ for some $A \in \text{NP} \cap \text{Co-NP}$. Let B_x be the A -oracle circuit whose truth-table is the witness for V' on input x that is produced by \mathcal{A} .

Final circuit F_x : $(n + c \log n)$ -bits long input r is given to C_x . The oracle gates are replaced by the circuit B_x . The final output is the output of C_x .

Final algorithm: On input x , we get C_x , non-deterministically guess B_x , construct F_x and run the fast GAP-SAT algorithm on F_x .

Correctness: The GAP-SAT algorithm on F_x checks if the non-deterministic guess B_x satisfies $\Pr_r [V(x, \text{tt}(B_x), r) = 1] \geq 1/2$, or equivalently $V'(x, \text{tt}(B_x)) = 1$. If $x \notin L$, this is not possible for any B_x due to the definition of V . If $x \in L$, as argued above, this is true for a poly-size A -oracle circuit B_x that captures witnesses for V' . □

8.3 New gap theorems for CAPP and MA

Results from the previous two sections also gives us gap theorems for CAPP and MA. First we saw that $\text{NEXP} \not\subset (\text{NP} \cap \text{Co-NP})/\text{poly}$ is equivalent to the derandomization of CAPP for $(\text{NP} \cap \text{Co-NP})$ -oracle circuits in NSUBEXP (infinitely often, with sub-polynomial advice). Then we saw that a non-trivial derandomization is sufficient to prove $\text{NEXP} \not\subset (\text{NP} \cap \text{Co-NP})/\text{poly}$. So we get the

following gap theorem for CAPP.

Theorem 50 (Gap theorem for CAPP on $(\text{NP} \cap \text{Co-NP})$ -oracle circuits). *Let sp be any super-polynomial function, then an $\text{NTIME}(2^n/sp(n))$ CAPP algorithm for n -input $\text{poly}(n)$ -size oracle circuits, for every $(\text{NP} \cap \text{Co-NP})$ -oracle, implies a $\bigcap_{\epsilon>0} \text{io-NTIME}(2^{n^\epsilon})/n^\epsilon$ algorithm for n -input $\text{poly}(n)$ -size oracle circuits, for every $(\text{NP} \cap \text{Co-NP})$ -oracle.*

In [42] they used NEXP KLT and its converse to establish a gap theorem for MA: either MA is as powerful as NEXP, or can be derandomized in NSUBEXP (infinitely often, with sub-polynomial advice). From the arguments in Section 8.1 we can get an improved gap theorem where $\text{MA} = \text{EXP}_{||}^{\text{NP}}$ in the first case.

Theorem 51 (Gap theorem for MA). *Exactly one of the following statements is true:*

1. $\text{MA} = \text{EXP}_{||}^{\text{NP}}$
2. $\text{MA} \subset \bigcap_{\epsilon>0} \text{io-NTIME}(2^{n^\epsilon})/n^\epsilon$

We also get a similar gap theorem for $\text{MA}^{\text{NP} \cap \text{Co-NP}}$: either $\text{MA}^{\text{NP} \cap \text{Co-NP}}$ is as powerful as $\text{EXP}_{||}^{\text{NP}}$, or can be derandomized in NSUBEXP (infinitely often, with sub-polynomial advice).

Theorem 52 (Gap theorem for $\text{MA}^{\text{NP} \cap \text{Co-NP}}$). *Exactly one of the following statements is true:*

1. $\text{MA}^{\text{NP} \cap \text{Co-NP}} = \text{EXP}_{||}^{\text{NP}}$
2. $\text{MA}^{\text{NP} \cap \text{Co-NP}} \subset \bigcap_{\epsilon>0} \text{io-NTIME}(2^{n^\epsilon})/n^\epsilon$

8.4 Fast algorithms imply lower bounds against circuits with limited prSV non-determinism

Here we show how fast tautology algorithms imply lower bounds for NE and E^{NP} , against circuits that use limited amount of prSV non-determinism.

Theorem 53. For $s(n) \in O(n)$:

$$1. \text{NE} \subset \text{prSV}^{s(n)\text{-C}} \implies \text{NE} \subset_{\text{ow}} \text{prSV}^{s(n)^{O(1)}\text{-C}}$$

$$2. \text{E}^{\text{NP}} \subset \text{prSV}^{s(n)\text{-C}} \implies \text{NE} \subset_{\text{ow}} \text{prSV}^{s(n)\text{-C}}$$

Proof. Let $L \in \text{NE}$, and V be an NE verifier for L .

Proof of (1): Since $\text{NEXP} \subset (\text{NP} \cap \text{Co-NP})/\text{poly} \implies \text{NEXP} = \text{AM}$, from [42] we get that the search problem for V is in EXP. Thus, there is an algorithm \mathcal{A} that: on any input $x \in L$ outputs y such that $V(x, y) = 1$; on any input $x \notin L$ outputs an all zeros string. Now define a new language $L' = \{(x, i) \mid i^{\text{th}} \text{ output bit of } \mathcal{A} \text{ on input } x \text{ is } 1\}$. $L' \in \text{EXP}$ and thus $L' \in \text{prSV}^{s(n)^{O(1)}\text{-C}}$. The circuit sequence for L' captures oblivious-witnesses for V .

Proof of (2): Let \mathcal{A} that is defined above, output the lexicographically smallest witnesses for V . Then its already known that the corresponding L' language is in E^{NP} (the algorithm does a binary search over all the witnesses, and use the NP-oracle to check if there is any positive witness smaller than the current witness). □

Before giving our main result, we state the local reductions that we will use in our proof.

Theorem 54 (Efficient local reductions [51, 90, 32]). *Every language $L \in \text{NTIME}(2^n)$ can be reduced to 3 – SAT instances of $2^n n^d$ -size, for some constant c . Moreover, given an instance of L there is an P-uniform deterministic circuit C that, on an integer $i \in [2^n n^d]$ in binary as input, output the i^{th} -clause of the resulting 3 – SAT formula. Each output bit of C depends on at most d input bits.*

Now we prove our main result.

Theorem 55. For super-polynomial function sp and $s(n) \leq O(n)$:

$$1. \text{an } \text{NTIME}(2^{n-s(n)^c}/sp(n)) \text{ C-tautology algorithm for every } c > 0 \text{ implies } \text{NE} \not\subset \text{prSV}^{s(n)\text{-C}}$$

$$2. \text{an } \text{NTIME}(2^{n-3s(n)}/sp(n)) \text{ C-tautology algorithm implies } \text{E}^{\text{NP}} \not\subset \text{prSV}^{s(n)\text{-C}}$$

Proof. Assume that $\text{NE} \subset \text{prSV}^{s(n)\text{-}C}$ or $\text{E}^{\text{NP}} \subset \text{prSV}^{s(n)\text{-}C}$. We contradict the NTIME hierarchy by giving an $\text{NTIME}(2^n/sp(n))$ algorithm for arbitrary $L \in \text{NE}$.

Reduction circuit: From the Theorem 54 we get: any input x for L uniformly reduces to a 3-SAT instance ϕ_x , where the number of variables and clauses in ϕ_x are bounded by $n^d 2^n$ for some constant d . Moreover the reduction is local in the sense that: it can be uniformly converted to a deterministic circuit C_x that on $(n + d \log n)$ -bits input i outputs the three variables x_{i1}, x_{i2}, x_{i3} ($3n + 3d \log n$ bits) from the i^{th} -clause of ϕ_x , along with three extra bits z_1, z_2, z_3 that indicate for each of these three variables, whether it appears as a positive literal or a negative literal.

Special verifier: Let V be a non-deterministic verifier for L , that first reduces L to 3 – SAT, and then non-deterministically guesses a satisfying assignment for the 3 – SAT formula.

Witness Circuits B_x : We construct two witness circuits (after guessing the advice of the $(\text{NP} \cap \text{Co} - \text{NP})/\text{poly}$ algorithm \mathcal{A} that has V 's oblivious-witnesses): one non-deterministic B_x^1 , and one co-non-deterministic B_x^2 .

Final Circuit F_x : Take the reduction circuit C_x . C_x outputs three literals. Plug any positive literal into a copy of the co-non-deterministic circuit B_x^2 , and any negative literal into a copy of the co-non-deterministic circuit $\overline{B_x^1}$. Output is the logical-or of the three copies used. To make the circuit deterministic, include the non-deterministic inputs of the copies of B_x^2 and $\neg B_x^1$ into the actual input.

Final algorithm: Get C_x . Non-deterministically guess the advice for \mathcal{A} , and get B_x^1 and B_x^2 (that are guaranteed to have complementary truth-tables). Construct the deterministic circuit F_x as described above. Run the the fast TAUT algorithm on F_x .

Correctness: $x \in L \iff F_x$ is a tautology. The tautology algorithm on F_x checks if the pair (B_x^1, B_x^2) satisfy, $V(x, tt(B_x^1)) = 1$. If $x \notin L$, this is not possible for any B_x^1 and B_x^2 . If $x \in L$, this is true for the witness circuits that exists due the easy-witness lemma proved in the above theorem. While constructing F_x , we use the fact that tautology of a co-non-deterministic circuit, is same as the tautology of the deterministic circuit we get after including the non-deterministic inputs into the actual input.

Final contradiction: The final input size is increased by $s(n)^c + O(\log n)$ if we use the EWL from assumption $NE \not\subseteq \text{prSV}^{s(n)}\text{-C}$, and is increased by $3s(n) + O(\log n)$ if we use the EWL from assumption $E^{\text{NP}} \not\subseteq \text{prSV}^{s(n)}\text{-C}$. So algorithms from our assumptions are fast enough to contradict the NTIME hierarchy. \square

8.5 Unconditional lower bounds against restricted circuits with limited prSV non-determinism

Using the Theorem 55 from previous section we get unconditional lower bounds against restricted circuits that use limited prSV non-determinism. The following theorem follows from $\text{TIME}(2^{n-n^\epsilon})$ tautology algorithm for ACC circuits [100], where the constant ϵ depends on the depth and the modulus function used by the circuits.

Theorem 56. $NE \not\subseteq \bigcap_{\epsilon > 0} \text{prSV}^{n^\epsilon}\text{-ACC}$

The following theorem follows from the $\text{ZPTIME}(2^{n(1-1/(\log n)^\epsilon)})$ tautology algorithm for AC^0 circuits [43], where the constant ϵ increases as the size or depth of the circuits increases.

Theorem 57. $E^{\text{NP}} \not\subseteq \bigcap_{\epsilon > 0} \text{prSV}^{n/(\log n)^\epsilon}\text{-AC}^0$

In the proof of the Theorem 55, the final circuit is constructed by giving the output bits of the reduction circuit as input to the witness circuit. Each output bit of the reduction circuit of Theorem 54 only depends on constant number of inputs, so can be represented by a set of constant-width clauses or terms, and thus can be plugged without increasing the depth. Thus the depth of the final circuit is only increased by the top OR-gate. For the case of E^{NP} , the final circuit also preserves the size of the witness circuit upto a constant factor. So we get the following result using fast AC^0 algorithms for different size and depth ranges [43].

Theorem 58. $E^{\text{NP}} \not\subseteq \bigcap_{\epsilon > 0} \text{prSV}^{\epsilon n / (\log n)^2}\text{-k-CNF}$ and $E^{\text{NP}} \not\subseteq \bigcap_{\epsilon > 0} \text{prSV}^{\epsilon n}\text{-AC}^0(n)$

Note that, $O(n)$ amount of prSV non-determinism in any of the above two lower bounds, will give super-linear lower bounds for E^{NP} .

8.6 Unconditional fixed-polynomial lower bounds against unrestricted circuits that use prSV non-determinism

Here we give unconditional lower bounds for prAM against fixed-polynomial size prSV non-deterministic circuits. We use the following PSPACE-complete language of Santhanam [80], which was also a crucial technical step in his celebrated MA lower-bound.

Lemma 8.6.1. *There is a PSPACE-complete language L^S and probabilistic polynomial-time oracle Turing machines M and M' such that the following holds for any n -length input x :*

1. M and M' only query their oracle on strings of length n .
2. If M (resp. M') is given L^S as its oracle and $x \in L^S$ (resp. $x \notin L^S$), then M (resp. M') accepts with probability 1.
3. If $x \notin L^S$ (resp. $x \in L^S$), then irrespective of the oracle, M (resp. M') rejects with probability at least $2/3$.

Like Santhanam's proof, our proof is also split into two cases: (i) The easier case where $PSPACE \subset (NP \cap Co - NP)/poly$, we use the KLT from [16]. (ii) The difficult case where PSPACE doesn't have poly-size prSV circuits, we design AM protocol for a padded version of L^S that doesn't have fixed-polynomial prSV circuits. We first prove an auxiliary lemma that we use for the second case.

Lemma 8.6.2. *For $k \geq 1$ and super-constant function sc , using L^S from Lemma 8.6.1 we define:*

$$L^k = \{x1^y \mid x \in L^S \wedge \exists(z \in \mathbb{N}) y = 2^z \geq |x| > 0, (2y + |x|)^{k+1} \geq ckt_{prSV}^{sc}(L_{|x|}^S) > (y + |x|)^{k+1}\}$$

If $PSPACE \not\subset (NP \cap Co - NP)/poly$, then $L^k \notin prSVNSIZE(n^k)$ for every $k \geq 1$.

Proof. For the sake of contradiction, let's assume that $L^k \in \text{prSVN}(n^k)$. That means, there is a prSVN algorithm \mathcal{A} , that produces an n^k -size SV non-deterministic circuit sequence, that decides L^k . We modify this sequence to yield a sequence for L^S (used in the definition of L^k). Any input length n can be broken into unique m and $y = 2^z$ such that $y \geq m$ and $m + y = n$. If y satisfies $(2y + m)^{k+1} \geq \text{ckt}_{\text{prSV}}^{sc}(L_m^S) \geq (y + m)^{k+1}$, then a circuit for the n^{th} -slice of L^k can be used to yield a circuit for the m^{th} -slice of L^S (by fixing the last y input bits to 1s). Moreover for any m , there is a unique y that satisfies $(2y + m)^{k+1} \geq \text{ckt}_{\text{prSV}}^{sc}(L_m^S) \geq (y + m)^{k+1}$ (since y is a power of 2).

For any input length m , the size of the circuit from this sequence will be $(m + y)^k$ for the unique y that is paired with m . This leads to the contradiction $(m + y)^k \geq \text{ckt}_{\text{prSV}(\mathcal{A})}(L_m^S) \geq \text{ckt}_{\text{prSV}}^{sc}(L_m^S) > (m + y)^{k+1}$ on input lengths m where L^S requires more than m^{k+1} size (due to $L^S \notin (\text{NP} \cap \text{Co-NP})/\text{poly}$) and thus a positive y exists. The first inequality follows from the fact that the circuit sequence is produced by \mathcal{A} . The second inequality uses the fact that the measure $\text{ckt}_{\text{prSV}}^{sc}$ beats the measure $\text{ckt}_{\text{prSV}(\mathcal{A})}$ for any prSVN algorithm \mathcal{A} , after a certain input length (because \mathcal{A} 's description is only of constant length, i.e. less than $sc(m)$). The third inequality follows from the definition of L^k . \square

Now we prove one of the two main results of this section.

Theorem 59. *For any super-constant function sc , $\forall k \text{ AM}/sc(n) \not\subseteq \text{prSVNSIZE}(n^k)$.*

Proof. If $\text{PSPACE} \subset (\text{NP} \cap \text{Co-NP})/\text{poly}$, then $\text{PSPACE} = \text{MA}^{\text{NP} \cap \text{Co-NP}}$. As in PSPACE we can diagonalize against any fixed-polynomial size circuit class, we get the desired fixed-polynomial circuit lower-bound for $\text{MA}^{\text{NP} \cap \text{Co-NP}}$ (without any advice). $\text{MA}^{\text{NP} \cap \text{Co-NP}}$ is contained $\text{MAM} = \text{AM}$ (Idea: after Arthur guesses its random bits, it sends them to Merlin, who then computes all the $\text{NP} \cap \text{Co-NP}$ queries Arthur will make, and sends Arthur the replies along with the certificates for the queries and their compliments).

If $\text{PSPACE} \not\subseteq (\text{NP} \cap \text{Co-NP})/\text{poly}$. From the Lemma 8.6.2 we get languages $(L^k$ for $k \geq 1$) with the desired lower bounds. We design $\text{AM}/sc(n)$ protocols for these languages. Arthur rejects everything if the first advice bit is 0. The first advice bit is 1 exactly for the input lengths n that

split into valid (m, y) pairs (see the proof of the Lemma 8.6.2 for the notion of valid pairs), when L^k is defined using the measure ckt_{prSV}^{sc-1} . Arthur rejects if the input is not in the format $x1^y$. Else, it simulates the machine M from the Lemma 8.6.1 to check if $x \in L^S$ or not. It accepts if and only if $x \in L^S$. It uses the circuit C , that it computes from Merlin's reply and the rest of the $sc(n) - 1$ bits of advice, as an oracle to M (from the Lemma 8.6.1).

The last $sc(n) - 1$ bits of advice encodes a prSVN algorithm \mathcal{A} . Correct advice encodes the most efficient one of the most efficient prSVN algorithms for that input length, i.e. an algorithm \mathcal{A} such that $ckt_{\text{SV}(\mathcal{A})}(L_{|x|}^S) = ckt_{\text{prSV}}(L_{|x|}^S)$. For n -length input $x1^y$ with $|x| = m$, Merlin sends an $(2y + m)^{k+1}$ -length input w for \mathcal{A} . Arthur produces the circuit $C = \mathcal{A}(w)$ to use as an oracle for M . Arthur then guesses random bits for the simulation of M and sends them to Merlin. Merlin in return sends the certificates that sets the flag bit of C to 1, on all the queries that M makes to C . Arthur uses these certificates, to compute the value bits of C , and thus successfully simulates M on x (using C as oracle).

Completeness follows easily. If $x \in L^S$, Merlin can send the input on which the algorithm \mathcal{A} outputs the correct SV circuit for L^S . If $x \notin L^S$, soundness follows from the fact that the algorithm \mathcal{A} always generates SV circuits, and thus the oracle used by M is consistent (to some language), and thus M rejects with probability at least $2/3$. \square

For each input length, assigning multiple input lengths corresponding to each possible advice, and making the input lengths with the correct advice as the promise input lengths, we get the following theorem. Also note that, for any class Γ , $\text{ip}(\Gamma \cap \text{co-}\Gamma)$ is a special case of the class $\text{pr}(\Gamma \cap \text{co-}\Gamma)$, which is a special case of $\text{pr}\Gamma \cap \text{prco-}\Gamma$.

Theorem 60. $\forall k \text{ ip}(\text{AM} \cap \text{Co-AM}) \not\subseteq \text{prSVNSIZE}(n^k)$.

Proof. In the proof of the Theorem 59 for $sc(n) = 1 + \log n$, if we use the machine M' from the Lemma 8.6.1, we can get a Co-AM/ $(1 + \log n)$ protocol for the language L^k that uses the same advice that the AM/ $1 + \log n$ protocol used. This protocol accepts if the input length n doesn't split into a valid m and y pair (using the first advice bit), and if the input is not in the $x1^y$ format. It both

these tests are passed by the input, then it accepts if M' accepts x (i.e. $x \notin L^S$). Only change is that Arthur simulates M' instead of M (using the same advice). Let's denote these $\text{AM}/(1 + \log n)$ and $\text{Co-AM}/(1 + \log n)$ protocols for L^k , by \mathcal{A} and \mathcal{A}' respectively.

Now, define a new modified language $T^k = \{g1^h \mid \mathcal{A} \text{ accepts } g \text{ on } h^{\text{th}} \text{ advice}\}$. Any input length $n + j$ in the range $[n + 1, 3n]$ is dedicated to the simulation of \mathcal{A} on the j^{th} advice (lexicographically j^{th} among all the $1 + \log n$ bits long advice strings). Now we create a promise problem prT^k using the language T^k , whose promise input lengths are the ones, that correspond to the correct advice for \mathcal{A} and \mathcal{A}' .

The protocols \mathcal{A} and \mathcal{A}' decide T^k and $\overline{T^k}$ correctly on the promise inputs of prT^k . \mathcal{A} and \mathcal{A}' also satisfy the semantic promises on these input lengths. Thus $prT^k \in \text{ip-}(\text{AM} \cap \text{Co-AM})$, and from the same arguments as in the proof of the Lemma 8.6.2, $prT^k \notin \text{prSVNSIZE}(n^{k-1})$. \square

Chapter 8 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Chapter 9

Conclusions and Open Problems

The main open problem is whether there are any connections between fast algorithms and non-uniform lower bounds possible within deterministic classes such as EXP. In almost all of the prior connections, non-uniformity is simulated with non-determinism, by having a non-deterministic machine guess the appropriate circuit. Can we substitute a recursive argument for non-determinism here? Our results show that, while still allowing non-determinism, the form of non-determinism can be weakened. In what other ways could we get such connections for smaller classes by restricting the use of non-determinism? The circuit model combines two features: time and non-uniformity. Can we get a fine-grained version of easy-witness lemma by distinguishing these two parameters?

Next obvious question in this line is whether we can get lower bounds for UEXP and related classes using our connections. Designing fast algorithms is one direct strategy. One other, seemingly easier strategy is to design tight hierarchy theorems for these semantic classes, possibly under the assumption that they have small circuits.

Our results also show that, if we are using unrestricted non-determinism to simulate non-uniformity, we can extract more out of it. That is, the guessed circuit is also allowed to use non-determinism that is promise-single-valued. In what other ways can we extend this allowance? Can we remove the promise condition? Specifically, can we prove NEXP easy-witness lemmas and Karp-Lipton theorems for circuit classes above $(\text{NP} \cap \text{Co-NP})/\text{poly}$?

We also show unconditional NEXP lower bounds where sub-polynomial and sub-linear amounts of promise-single-valued non-determinism is allowed. Can we increase the amount of non-determinism allowed, to polynomial or linear? Designing fast algorithms is one direct strategy. Can we do it without changing the satisfiability upper bounds? This would lead to super-linear and super-polynomial lower bounds against unrestricted Boolean circuits. Or can we get lower bounds against TC^0 circuits by simulating threshold gates, by the use of less expressive gates and limited non-determinism?

Chapter 9 contains material from “On Limiting & Limited Non-determinism in NEXP Lower Bounds”, by Anant Dhayal and Russell Impagliazzo, which is currently under review in ACM Transactions on Computation Theory (TOCT). The dissertation author was the primary investigator and author of this paper.

Bibliography

- [1] S. Aaronson and A. Wigderson. Algebrization: A new barrier in complexity theory. *ACM Trans. Comput. Theory*, 1(1):2:1–2:54, 2009.
- [2] M. Ajtai. Σ^1_1 -formulae on finite structures. *Ann. Pure Appl. Log.*, 24(1):1–48, 1983.
- [3] E. Allender. The complexity of sparse sets in P. In *Structure in Complexity Theory, Proceedings of the Conference hold at the University of California, Berkeley, California, USA, June 2-5, 1986*, pages 1–11, 1986.
- [4] E. Allender. When worlds collide: Derandomization, lower bounds, and kolmogorov complexity. In R. Hariharan, V. Vinay, and M. Mukund, editors, *FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science*, pages 1–15, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [5] N. Alon and R. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987.
- [6] D. Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1987.
- [7] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
- [8] B. Aydinlioglu and D. van Melkebeek. Nondeterministic circuit lower bounds from mildly derandomizing arthur-merlin games. *Computational Complexity*, 26(1):79–118, 2017.
- [9] L. Babai, L. Fortnow, N. Nisan, and A. Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3:307–318, 1993.
- [10] T. P. Baker, J. Gill, and R. Solovay. Relativizations of the P =? NP question. *SIAM J. Comput.*, 4(4):431–442, 1975.
- [11] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. P. Vadhan. Short pcps verifiable in polylogarithmic time. In *20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA*, pages 120–134. IEEE Computer Society, 2005.
- [12] N. H. Bshouty, R. Cleve, R. Gavaldà, S. Kannan, and C. Tamon. Oracles and queries that are sufficient for exact learning. *J. Comput. Syst. Sci.*, 52(3):421–433, 1996.

- [13] H. Buhrman, L. Fortnow, and T. Thierauf. Nonrelativizing separations. In *Proceedings of the 13th Annual IEEE Conference on Computational Complexity, Buffalo, New York, USA, June 15-18, 1998*, pages 8–12. IEEE Computer Society, 1998.
- [14] H. Buhrman and S. Homer. Superpolynomial circuits, almost sparse oracles and the exponential hierarchy. In R. K. Shyamasundar, editor, *Foundations of Software Technology and Theoretical Computer Science, 12th Conference, New Delhi, India, December 18-20, 1992, Proceedings*, volume 652 of *Lecture Notes in Computer Science*, pages 116–127. Springer, 1992.
- [15] J. Cai. SP_2 subseteq zpp^{NP} . In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 620–629. IEEE Computer Society, 2001.
- [16] J.-Y. Cai, V. T. Chakaravarthy, L. A. Hemaspaandra, and M. Ogihara. Competing provers yield improved karp–lipton collapse results. *Information and Computation*, 198(1):1 – 23, 2005.
- [17] C. Calabro, R. Impagliazzo, and R. Paturi. A duality between clause width and clause density for SAT. In *21st Annual IEEE Conference on Computational Complexity (CCC 2006), 16-20 July 2006, Prague, Czech Republic*, pages 252–260. IEEE Computer Society, 2006.
- [18] C. Calabro, R. Impagliazzo, and R. Paturi. The complexity of satisfiability of small depth circuits. In J. Chen and F. V. Fomin, editors, *Parameterized and Exact Computation, 4th International Workshop, IWPEC 2009, Copenhagen, Denmark, September 10-11, 2009, Revised Selected Papers*, volume 5917 of *Lecture Notes in Computer Science*, pages 75–85. Springer, 2009.
- [19] M. Carmosino, R. Impagliazzo, V. Kabanets, and A. Kolokolova. Tighter connections between derandomization and circuit lower bounds. In N. Garg, K. Jansen, A. Rao, and J. D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, volume 40 of *LIPICs*, pages 645–658. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.
- [20] M. L. Carmosino, J. Gao, R. Impagliazzo, I. Mihajlin, R. Paturi, and S. Schneider. Non-deterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In M. Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 261–270. ACM, 2016.
- [21] B. Chapman and R. Williams. The circuit-input game, natural proofs, and testing circuits with data. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15*, page 263–270, New York, NY, USA, 2015. Association for Computing Machinery.
- [22] L. Chen, R. Rothblum, R. Tell, and E. Yogev. On exponential-time hypotheses, derandomization, and circuit lower bounds. *Electron. Colloquium Comput. Complex.*, 26:169, 2019.

- [23] R. Chen, V. Kabanets, A. Kolokolova, R. Shaltiel, and D. Zuckerman. Mining circuit lower bound proofs for meta-algorithms. *Comput. Complex.*, 24(2):333–392, 2015.
- [24] T. Y. Chow. Almost-natural proofs. *Journal of Computer and System Sciences*, 77(4):728 – 737, 2011. JCSS IEEE AINA 2009.
- [25] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pages 151–158, 1971.
- [26] S. A. Cook. A hierarchy for nondeterministic time complexity. In P. C. Fischer, H. P. Zeiger, J. D. Ullman, and A. L. Rosenberg, editors, *Proceedings of the 4th Annual ACM Symposium on Theory of Computing, May 1-3, 1972, Denver, Colorado, USA*, pages 187–192. ACM, 1972.
- [27] H. Dell, V. Kabanets, D. van Melkebeek, and O. Watanabe. Is valiant-vazirani’s isolation probability improvable? *Computational Complexity*, 22(2):345–383, 2013.
- [28] A. Dhayal and R. Impagliazzo. UTIME easy-witness lemma & some consequences. *Electron. Colloquium Comput. Complex.*, 26:167, 2019.
- [29] J. Feigenbaum and L. Fortnow. Random-self-reducibility of complete sets. *SIAM J. Comput.*, 22(5):994–1005, 1993.
- [30] L. Fortnow and A. R. Klivans. NP with small advice. In *20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA*, pages 228–234. IEEE Computer Society, 2005.
- [31] L. Fortnow and A. R. Klivans. Efficient learning algorithms yield circuit lower bounds. *J. Comput. Syst. Sci.*, 75(1):27–36, 2009.
- [32] L. Fortnow, R. Lipton, D. van Melkebeek, and A. Viglas. Time-space lower bounds for satisfiability. *J. ACM*, 52(6):835–865, Nov. 2005.
- [33] L. Fortnow and R. Santhanam. Robust simulations and significant separations. In L. Aceto, M. Henzinger, and J. Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, volume 6755 of *Lecture Notes in Computer Science*, pages 569–580. Springer, 2011.
- [34] L. Fortnow, R. Santhanam, and L. Trevisan. Hierarchies for semantic classes. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, STOC '05*, pages 348–355, New York, NY, USA, 2005. ACM.
- [35] M. L. Furst, J. B. Saxe, and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. *Math. Syst. Theory*, 17(1):13–27, 1984.
- [36] R. C. Harkins and J. M. Hitchcock. Exact learning algorithms, betting games, and circuit lower bounds. *TOCT*, 5(4):18:1–18:11, 2013.

- [37] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965.
- [38] J. Håstad. Almost optimal lower bounds for small depth circuits. In J. Hartmanis, editor, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 6–20. ACM, 1986.
- [39] F. C. Hennie and R. E. Stearns. Two-tape simulation of multitape turing machines. *J. ACM*, 13(4):533–546, 1966.
- [40] T. Hertli. 3-sat faster and simpler - unique-sat bounds for ppsz hold in general. In *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS '11*, page 277–284, USA, 2011. IEEE Computer Society.
- [41] R. Impagliazzo, V. Kabanets, and I. Volkovich. The power of natural properties as oracles. In R. A. Servedio, editor, *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*, volume 102 of *LIPICs*, pages 7:1–7:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [42] R. Impagliazzo, V. Kabanets, and A. Wigderson. In search of an easy witness: exponential time vs. probabilistic polynomial time. *J. Comput. Syst. Sci.*, 65(4):672–694, 2002.
- [43] R. Impagliazzo, W. Matthews, and R. Paturi. A satisfiability algorithm for ac^0 . In Y. Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 961–972. SIAM, 2012.
- [44] R. Impagliazzo and R. Paturi. Complexity of k-sat. In *Proceedings of the 14th Annual IEEE Conference on Computational Complexity, Atlanta, Georgia, USA, May 4-6, 1999*, pages 237–240, 1999.
- [45] R. Impagliazzo, R. Paturi, and S. Schneider. A satisfiability algorithm for sparse depth two threshold circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 479–488. IEEE Computer Society, 2013.
- [46] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? In *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, pages 653–663, 1998.
- [47] R. Impagliazzo and G. Tardos. Decision versus search problems in super-polynomial time. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 222–227. IEEE Computer Society, 1989.
- [48] R. Impagliazzo and A. Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 220–229, 1997.

- [49] R. Impagliazzo and A. Wigderson. Randomness vs time: Derandomization under a uniform assumption. *J. Comput. Syst. Sci.*, 63(4):672–688, 2001.
- [50] K. Iwama and H. Morizumi. An explicit lower bound of $5n - o(n)$ for boolean circuits. In K. Diks and W. Rytter, editors, *Mathematical Foundations of Computer Science 2002, 27th International Symposium, MFCS 2002, Warsaw, Poland, August 26-30, 2002, Proceedings*, volume 2420 of *Lecture Notes in Computer Science*, pages 353–364. Springer, 2002.
- [51] H. Jahanjou, E. Miles, and E. Viola. Local reductions. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 749–760, 2015.
- [52] V. Kabanets. Easiness assumptions and hardness tests: Trading time for zero error. *J. Comput. Syst. Sci.*, 63(2):236–252, 2001.
- [53] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.
- [54] R. Kannan. Circuit-size lower bounds and non-reducibility to sparse sets. *Information and Control*, 55(1-3):40–56, 1982.
- [55] R. M. Karp and R. J. Lipton. Some connections between nonuniform and uniform complexity classes. In *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing, STOC '80*, pages 302–309, New York, NY, USA, 1980. ACM.
- [56] J. Köbler and O. Watanabe. New collapse consequences of np having small circuits. *SIAM Journal on Computing*, 28(1):311–324, 1998.
- [57] A. R. Klivans, P. Kothari, and I. C. Oliveira. Constructing hard functions using learning algorithms. In *Proceedings of the 28th Conference on Computational Complexity, CCC 2013, K.lo Alto, California, USA, 5-7 June, 2013*, pages 86–97, 2013.
- [58] A. R. Klivans and D. van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM J. Comput.*, 31(5):1501–1526, 2002.
- [59] M. Krause and S. Lucks. Pseudorandom functions in tc_0 and cryptographic limitations to proving lower bounds. *Comput. Complex.*, 10(4):297–313, May 2002.
- [60] O. Lachish and R. Raz. Explicit lower bound of $4.5n - o(n)$ for boolean circuits. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing, STOC '01*, page 399–408, New York, NY, USA, 2001. Association for Computing Machinery.
- [61] L. A. Levin. Universal sorting problems. *Problems of Information Transmission*, 9:265–266, 1973.
- [62] C. Lund, L. Fortnow, H. J. Karloff, and N. Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992.

- [63] K. Makino, S. Tamaki, and M. Yamamoto. Derandomizing the HSSW algorithm for 3-sat. *Algorithmica*, 67(2):112–124, 2013.
- [64] E. Miles and E. Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. In R. Safavi-Naini and R. Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, pages 68–85, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [65] B. Monien and E. Speckenmeyer. Solving satisfiability in less than 2^n steps. *Discret. Appl. Math.*, 10(3):287–295, 1985.
- [66] H. Morizumi. Lower bounds for the size of nondeterministic circuits. In D. Xu, D. Du, and D. Du, editors, *Computing and Combinatorics*, pages 289–296, Cham, 2015. Springer International Publishing.
- [67] R. A. Moser and D. Scheder. A full derandomization of schönig’s k-sat algorithm. In L. Fortnow and S. P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 245–252. ACM, 2011.
- [68] C. Murray and R. R. Williams. Circuit lower bounds for nondeterministic quasi-polytime: an easy witness lemma for NP and NQP. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 890–901, 2018.
- [69] M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. ACM*, 51(2):231–262, Mar. 2004.
- [70] N. Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica*, 11(1):63–70, 1991.
- [71] N. Nisan and A. Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994.
- [72] I. C. Oliveira. Algorithms versus circuit lower bounds. *CoRR*, abs/1309.0249, 2013.
- [73] I. C. Oliveira and R. Santhanam. Conspiracies between learning algorithms, circuit lower bounds, and pseudorandomness. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, pages 18:1–18:49, 2017.
- [74] R. Paturi, P. Pudlák, M. E. Saks, and F. Zane. An improved exponential-time algorithm for k-sat. *J. ACM*, 52(3):337–364, 2005.
- [75] R. Paturi, P. Pudlák, and F. Zane. Satisfiability coding lemma. *Chicago J. Theor. Comput. Sci.*, 1999, 1999.
- [76] A. A. Razborov. Lower bounds for the monotone complexity of some boolean functions. *Soviet Math. Dokl.*, 31:354–357, 1985.
- [77] A. A. Razborov. On the method of approximations. In D. S. Johnson, editor, *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 167–176. ACM, 1989.

- [78] A. A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1):24 – 35, 1997.
- [79] S. Rudich. Super-bits, demi-bits, and np/qpoly-natural proofs. 1269, 08 1997.
- [80] R. Santhanam. Circuit lower bounds for merlin–arthur classes. *SIAM J. Comput.*, 39(3):1038–1061, 2009.
- [81] R. Santhanam. Fighting peregbor: New and improved algorithms for formula and QBF satisfiability. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 183–192. IEEE Computer Society, 2010.
- [82] U. Schöning. A probabilistic algorithm for k-sat and constraint satisfaction problems. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*, pages 410–414. IEEE Computer Society, 1999.
- [83] R. Schuler. An algorithm for the satisfiability problem of formulas in conjunctive normal form. *J. Algorithms*, 54(1):40–44, 2005.
- [84] J. I. Seiferas, M. J. Fischer, and A. R. Meyer. Separating nondeterministic time complexity classes. *J. ACM*, 25(1):146–167, 1978.
- [85] K. Seto and S. Tamaki. A satisfiability algorithm and average-case hardness for formulas over the full binary basis. In *Proceedings of the 27th Conference on Computational Complexity, CCC 2012, Porto, Portugal, June 26-29, 2012*, pages 107–116. IEEE Computer Society, 2012.
- [86] R. Shaltiel and C. Umans. Pseudorandomness for approximate counting and sampling. *Computational Complexity*, 15(4):298–341, 2006.
- [87] R. Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In A. V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 77–82. ACM, 1987.
- [88] D. M. Stull. Some results on circuit lower bounds and derandomization of arthur-merlin problems. *CoRR*, abs/1701.04428, 2017.
- [89] M. Sudan, L. Trevisan, and S. P. Vadhan. Pseudorandom generators without the XOR lemma. *J. Comput. Syst. Sci.*, 62(2):236–266, 2001.
- [90] I. Tourlakis. Time–space tradeoffs for sat on nonuniform machines. *Journal of Computer and System Sciences*, 63(2):268 – 287, 2001.
- [91] G. S. Tseitin. *On the Complexity of Derivation in Propositional Calculus*, pages 466–483. Springer Berlin Heidelberg, Berlin, Heidelberg, 1983.
- [92] C. Umans. Pseudo-random generators for all hardnesses. *J. Comput. Syst. Sci.*, 67(2):419–440, 2003.

- [93] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984.
- [94] L. G. Valiant and V. V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
- [95] N. Vinodchandran. A note on the circuit complexity of pp. *Theoretical Computer Science*, 347(1):415 – 418, 2005.
- [96] N. V. Vinodchandran. $\text{Am}_{\exp}[\text{nsub}_e](\text{np}[\text{cap}]\text{conp})/\text{poly}$. *Inf. Process. Lett.*, 89(1):43–47, 2004.
- [97] N. Vyas and R. Williams. On super strong eth. In M. Janota and I. Lynce, editors, *Theory and Applications of Satisfiability Testing – SAT 2019*, pages 406–423, Cham, 2019. Springer International Publishing.
- [98] N. Vyas and R. R. Williams. Lower bounds against sparse symmetric functions of ACC circuits: Expanding the reach of #sat algorithms. In C. Paul and M. Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPICs*, pages 59:1–59:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [99] R. Williams. Improving exhaustive search implies superpolynomial lower bounds. *SIAM J. Comput.*, 42(3):1218–1244, 2013.
- [100] R. Williams. Nonuniform ACC circuit lower bounds. *J. ACM*, 61(1):2:1–2:32, 2014.
- [101] R. R. Williams. Natural proofs versus derandomization. *SIAM J. Comput.*, 45(2):497–529, 2016.
- [102] R. R. Williams. New algorithms and lower bounds for circuits with linear threshold gates. *Theory of Computing*, 14(1):1–25, 2018.
- [103] A. C. Yao. Separating the polynomial-time hierarchy by oracles (preliminary version). In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*, pages 1–10. IEEE Computer Society, 1985.
- [104] S. Zak. A turing machine time hierarchy. *Theoretical Computer Science*, 26(3):327 – 333, 1983.