Electronic Colloquium on Computational Complexity, Report No. 178 (2026)

Optimal Hitting Set Generators via A Potential-Descent Framework

Gonen Krak[†]

Abstract

We develop a new framework for constructing explicit hitting-set generators (HSGs) for classes of Boolean functions. The framework is inspired by the Ajtai–Wigderson approach based on iterated pseudorandom restrictions.

We use this framework to construct HSGs for read-once CNFs and read-once CNFs with parities. For formulas on n variables with density at least ε, our generators use $O(\log(n/\varepsilon))$ seed bits. This bound is optimal up to a constant factor. Applying the hitting reduction of Gopalan, Meka, Reingold, Trevisan, and Vadhan [GMR+12] gives the same seed bound for general ordered width-3 read-once branching programs. For this class, our construction is the first to achieve optimal dependence on both length and acceptance density. The result provides further evidence supporting the longstanding conjecture L = RL.

[†]Tel Aviv University. gonenok10@gmail.com

ISSN 1433-8092

Contents

1 Introduction

Pseudorandom generators are a central tool for reducing the randomness used by algorithms. A pseudorandom generator (PRG) replaces a long sequence of independent uniform bits by the output of a deterministic function of a short uniform seed. Its output should approximately preserve the acceptance probability of every function in a given class. A hitting-set generator (HSG) asks for less: whenever a function accepts at least an ε fraction of all inputs, some output of the generator must be accepted. The distinction matters when the goal is to find one accepting input.

The main contribution of this paper is a new *potential-descent framework* for constructing hitting-set generators. The framework is stated for an arbitrary class of Boolean functions and gives sufficient conditions for an HSG with a short seed. It separates the construction of the generator from the estimates needed for the particular function class. The analysis assigns variables in rounds and uses a potential function to bound the total randomness needed along an accepting sequence of assignments. We apply the framework first to read-once CNFs and then to read-once CNFs with parities, obtaining new HSGs with optimal seed length $O(\log(n/\varepsilon))$ for both classes. A known hitting reduction then gives a new optimal HSG for ordered width-3 read-once branching programs (ROBPs), with the same seed length up to a constant factor. These optimal HSG constructions are further main results of the paper.

1.1 Motivation and Related Work

1.1.1 Read-once branching programs

One motivation comes from space-bounded computation. A logspace-computable HSG with logarithmic seed length for polynomial-width ROBPs would imply $\mathrm{L} = \mathrm{RL}$, equating deterministic logarithmic space with randomized logarithmic space with one-sided error. A corresponding PRG would imply $\mathrm{L} = \mathrm{BPL}$, the analogous equality for two-sided error. Cheng and Hoza [CH20] showed that optimal logspace hitting sets for this model would in fact also imply $\mathrm{L} = \mathrm{BPL}$.

An ordered ROBP of length n and width w reads n input bits once, in a fixed order, and has at most w states at each step. We use δ for the additive error of a PRG and ε for the density threshold of an HSG. Nisan [Nis92] constructed a PRG for these programs with seed length

$$O\Big(\log n \cdot \log \frac{nw}{\delta}\Big) = O\big(\log n\big(\log n + \log w + \log(1/\delta)\big)\big)\,.$$

For polynomial width and constant or inverse-polynomial error, this gives $O(\log^2 n)$ seed bits. Improving this dependence on n is a central problem in space-bounded derandomization; the general goal is seed length $O(\log n + \log w + \log(1/\delta))$.

For hitting, one can improve the dependence on the density threshold. Hoza and Zuckerman [HZ20] constructed HSGs for general ordered ROBPs with, in particular, seed length

$$O(\log n \cdot \log(nw) + \log(1/\varepsilon))\,.$$

Thus, for polynomial width, their bound is $O(\log^2 n + \log(1/\varepsilon))$.

More efficient PRGs are known when the transition graphs have additional structure. A regular ROBP has indegree zero or two at every state outside the first layer. Braverman, Rao, Raz, and Yehudayoff [BRRY14] constructed PRGs for regular ROBPs of length n and width w, with error δ, using

$$O\big(\log n\big(\log w + \log\log n + \log(1/\delta)\big)\big)$$

seed bits.

A permutation ROBP has the additional property that, in each layer, the transitions labeled 0 and those labeled 1 each form a permutation of the state space. Cohen, Doron, and Goldgraber [CDG26] gave a new analysis of the Impagliazzo–Nisan–Wigderson (INW) generator for permutation ROBPs of length n and width w, obtaining error δ with seed length

$$O\big(\log n\big(\log w + \log(1/\delta)\big)\big)\,.$$

For width-2 ROBPs, an optimal PRG is already obtained from the small-bias sets of Naor and Naor [NN93, BDVY13]. Their construction uses $O(\log n + \log(1/\delta))$ seed bits to generate a δ-biased distribution.

Width-3 programs provide a small but expressive setting: they can combine parity functions, local constraints, and transitions that merge states. Šíma and Žák [SZ11] obtained explicit polynomial-size hitting sets for width-3 read-once branching programs at sufficiently large constant acceptance probability. Their full proof [SZ21] gives this guarantee for every fixed density threshold $\varepsilon > 5/6$, even when the variable order can depend on the input. Enumerating these sets gives HSGs with $O(\log n)$ seed length in this density range.

For arbitrary density thresholds, Gopalan, Meka, Reingold, Trevisan, and Vadhan (GMRTV) [GMR+12] constructed an HSG for ordered width-3 ROBPs with seed length

$$O\Big(\log(n/\varepsilon)\big(\log\log(n/\varepsilon)\big)^3\Big)\,.$$

Meka, Reingold, and Tal [MRT19] subsequently constructed a PRG for the same class with error δ and seed length

$$\widetilde{O}(\log(n/\delta)) + O(\log n \cdot \log(1/\delta))\,.$$

Here $\widetilde{O}(t)$ denotes $O(t\ \mathrm{poly}(\log t))$. Both works use pseudorandom restrictions. This approach also motivates our framework.

1.1.2 Pseudorandom restrictions

Ajtai and Wigderson [AW85] introduced an approach to derandomization based on pseudorandom restrictions. The idea is to assign some variables, preserve the acceptance probability of the function, and repeat until the remaining function is easy to fool. GMRTV developed this approach using *mild* restrictions, which assign only a controlled fraction of the remaining variables in each round [GMR+12]. They obtained PRGs for read-once CNFs and read-once CNFs with parities, with seed length $\widetilde{O}(\log(n/\delta))$ for error δ. These generators also give their width-3 HSG through the structural reduction.

Several later results developed this approach further. In addition to the width-3 PRG of Meka, Reingold, and Tal [MRT19], Doron, Hatami, and Hoza [DHH19] obtained $\widetilde{O}(\log(n/\delta))$

seed length for constant-depth read-once formulas over AND, OR, and NOT. Related work by Lee and Viola [LV20] studies read-once polynomials and products of functions on disjoint inputs. Lee [Lee19] proved Fourier bounds and improved PRGs for such products, obtaining $\widetilde{O}(m + \log(n/\delta))$ seed length for the AND or XOR of functions on disjoint blocks of at most m bits. These results show how the structure of disjoint constraints can be used within the iterated-restriction approach.

Doron, Hatami, and Hoza [DHH20] subsequently used iterated restrictions to construct an explicit PRG for read-once depth-2 formulas over AND, OR, and parity, with negations allowed at the inputs. For error δ, their seed length is

$$O\Big(\log n + \log(1/\delta)\big(\log\log(4/\delta)\big)^5\Big) .$$

This class includes read-once CNFs and read-once CNFs with parities. The bound is optimal for constant error and has near-optimal dependence on smaller errors.

In their closing discussion [DHH20, Section 6], they ask whether iterated restrictions with a growing number of rounds can give a PRG, or even an HSG, with fully optimal dependence on both the input length and the error or acceptance density. Our results answer the HSG part of this question affirmatively: the potential-descent framework gives HSGs for read-once CNFs and read-once CNFs with parities with seed length $O(\log(n/\varepsilon))$, using $\Theta(\log\log n)$ rounds.

1.2 Main results

Our application to ordered width-3 ROBPs gives the following seed-length bound.

Theorem 1.1 (Main theorem)**.** *For every $n \geq 2$ and every $\varepsilon \in (0, 1/2]$, there is an explicit generator*

$$G_{n,\varepsilon} : \{0,1\}^s \to \{0,1\}^n, \qquad s = O\Big(\log \frac{n}{\varepsilon}\Big) ,$$

such that every ordered width-3 ROBP P on n variables satisfying

$$\Pr_{x\sim\mathcal{U}_n}[P(x) = 1] \geq \varepsilon$$

is hit by the support of $G_{n,\varepsilon}$.

This is optimal up to an absolute constant factor.

The construction uses the structural reduction of GMRTV [GMR+12]. Their reduction turns the width-3 hitting problem into a hitting problem for read-once CNFs with parities. These are conjunctions of constraints on disjoint variable sets, where each constraint is either an OR of literals or an affine parity equation. We write $\mathrm{CNF}^{\oplus}$ for this class. The new hitting-set construction is the following.

Theorem 1.2 (HSG for read-once CNF with parities)**.** *For every $n \geq 2$ and every $\varepsilon \in (0, 1/2]$, there is an explicit HSG for read-once $\mathrm{CNF}^{\oplus}$ formulas on n variables of uniform density at least ε, with seed length*

$$O\Big(\log \frac{n}{\varepsilon}\Big) .$$

We first prove the same seed bound for ordinary read-once CNFs, whose clauses use disjoint variable sets.

Theorem 1.3 (Warm-up theorem)**.** *For every $n \geq 2$ and every $\varepsilon \in (0, 1/2]$, there is an explicit HSG for read-once CNF formulas on n variables of uniform density at least ε, with seed length*

$$O\left(\log \frac{n}{\varepsilon}\right).$$

The read-once CNF proof introduces the partition construction, the potential, and the choice of supported assignments. We then extend these arguments to parity constraints and apply the GMRTV reduction. The formula generators work for every choice of disjoint variable sets; the generator is not given the clauses or their variable sets.

1.3 The potential-descent framework at a high level

Our framework constructs an HSG by assigning variables in rounds. We choose an ordered partition $\mathcal{P} = (I_1, \ldots, I_J)$ of the coordinates and assign the variables in I_k in round k. For each function F and each sequence of earlier assignments, the proof specifies a *round obligation*: a condition on the next assignment. These obligations are chosen so that satisfying all of them gives an input accepted by F. The generator must be fixed independently of F, but its support must contain at least one such accepting sequence for every function in the class.

The construction addresses four questions.

1. How can a fixed source satisfy an obligation that depends on the function?

Suppose the current obligation can be satisfied by setting at most K_k variables to suitable values. We call this a *local sufficient condition*. For a read-once CNF, for example, we can choose one literal from each clause that must be satisfied in the current round. Setting these literals to true satisfies the obligation, so K_k is the number of clauses that need attention in that round.

If these variables were known, assigning them K_k independent uniform bits would give every possible pattern probability 2^{-K_k}. In particular, some assignment would satisfy the local condition.

The generator, however, is not given F. It therefore cannot determine which variables or values are needed: these can depend on the function, the partition, and the earlier assignments. We instead use a short-seeded source that assigns all $|I_k|$ variables and is approximately uniform on every set of at most K_k coordinates. The approximation must be accurate enough that every pattern still has positive probability. One fixed source then contains a suitable assignment regardless of which local condition arises.

2. How can the generator choose a source without knowing K_k?

We prepare a family of round sources $G_{k,r}$, indexed by a nonnegative integer r. The parameter r specifies how many coordinates the source must handle. The generator guesses the required parameters by enumerating *round profiles*

$$q = (q_1, \ldots, q_J).$$

For a particular profile, it uses G_{k,q_k} in round k.

Larger parameters require longer seeds. We therefore enumerate only profiles satisfying

$$\sum_{k=1}^{J} q_k \leq L,$$

where L is a fixed budget. Such profiles can be encoded using $O(L + J)$ bits. The proof must show that, for every function we want to hit, some accepting sequence in the round-source supports has

$$\sum_{k=1}^{J} K_k \leq L.$$

Its actual sequence of parameters will then appear among the profiles the generator enumerates.

3. How do we keep the parameters in later rounds small?

Meeting the current obligation is only part of the task. Different assignments satisfying it can leave very different obligations for later rounds. We need to find assignments that also keep the sum of the later parameters small.

To guide these choices, we define a nonnegative potential Φ_k. It includes the current parameter K_k and weights for obligations that remain in later rounds. Our goal is to find an assignment from the current source that satisfies the obligation and also gives

$$\Phi_{k+1} \leq \alpha(\Phi_k - K_k), \qquad \alpha = 1 + \frac{1}{20J}.$$

We prove this by bounding the average next potential over source outputs that satisfy the local sufficient condition. At least one of these outputs has next potential no larger than that average.

This requires more pseudorandomness than merely ensuring that the current condition can be satisfied. We choose a threshold t_k: the source must behave approximately uniformly on sets of up to $K_k + t_k$ coordinates. The first K_k coordinates handle the current condition; the additional coordinates let us estimate the terms of the future potential. Designing these terms and proving the conditional expectation bound are the main tasks for each function class.

With $\Phi_J = K_J$, the descent inequalities imply

$$\sum_{k=1}^{J} K_k = O(\Phi_1)$$

along one accepting sequence in the source supports. Thus a small initial potential guarantees an accepting sequence with a small total round parameter.

4. How do we find a partition with a small initial potential?

The initial potential depends on the partition. We construct a short-seeded distribution over partitions, fixed independently of the function, and prove that its average initial potential is small for every function we want to hit. Writing B for the desired bound, we seek

$$\mathbb{E}_{\mathcal{P}}[\Phi_1] \leq B.$$

The choice of B depends on the function class and the parameters of the construction.

The proof may first replace the function by a stronger one whose accepting inputs also satisfy the original function. The potential is then evaluated for that replacement.

The average bound guarantees that at least one partition has initial potential at most B. For this partition, the descent argument gives an accepting sequence in the round-source supports with

$$\sum_{k=1}^{J} K_k = O(B).$$

We choose the profile budget L large enough to include this sequence of parameters.

The generator enumerates the partitions in the distribution's support and all profiles satisfying $\sum_k q_k \leq L$. For each partition and profile, it combines outputs from the corresponding round sources. A seed therefore specifies three things: a partition, a profile, and a seed for each round source.

For every function we want to hit, the proof identifies one combination that produces an accepting input. The final seed length is determined by the bits needed to choose the partition, encode the profile, and generate the round assignments. Bounding these three contributions completes the construction.

Organization. Section 1.4 records the definitions and results used in the construction. Section 1.5 reviews the PRG iterated-restriction framework. Section 2 develops the general potential-descent framework and its seed-length calculation. Section 3 applies it to read-once CNF: first for a fixed partition, then by constructing the partition distribution. Section 4 adds parity constraints, explains which arguments carry over, and proves the additional estimates. Section 4.7 uses the GMRTV reduction to prove the width-3 HSG result. Section 5 discusses future research directions.

1.4 Definitions and results used

Notation. We write $[n] = \{1, \ldots, n\}$ and $\mathcal{U}_A$ for the uniform distribution on $\{0,1\}^A$; when $A = [n]$, we write $\mathcal{U}_n$. The *density* of a Boolean function F is $\Pr_{x \sim \mathcal{U}_n}[F(x) = 1]$. For a generator G, its support is the set of outputs obtained from all seeds. Logarithms are base two unless written as ln. An empty product is one. We call a generator *explicit* if its output can be computed in time polynomial in the output length and the inverse error or density parameter.

Pseudorandom generators and hitting-set generators. For a class $\mathcal{F}_n$ of Boolean functions on n variables, a generator $G : \{0,1\}^s \to \{0,1\}^n$ is a PRG with error δ if

$$|\Pr[F(G(\mathcal{U}_s)) = 1] - \Pr[F(\mathcal{U}_n) = 1]| \leq \delta \qquad (F \in \mathcal{F}_n).$$

It is an HSG with density threshold ε if

$$\Pr[F(\mathcal{U}_n) = 1] \geq \varepsilon \quad \Longrightarrow \quad \exists z \in \{0,1\}^s \;\; F(G(z)) = 1 \qquad (F \in \mathcal{F}_n).$$

In both cases, s is the seed length. A PRG with error smaller than ε is an HSG with threshold ε.

Ordered read-once branching programs. An ordered ROBP of length n and width w is a directed graph with layers $0, \ldots, n$, each containing at most w vertices. It has a designated start vertex in layer 0 and a set of accepting vertices in layer n. Each vertex in layer $i-1$ has two outgoing edges to layer i, labeled 0 and 1. On input x, the program follows the edge labeled x_i at step i and accepts if it reaches an accepting vertex. We use the standard order $x_1, \ldots, x_n$; another fixed known order can be handled by relabeling the coordinates.

Read-once CNFs and parity constraints. A read-once CNF is a conjunction $F = \bigwedge_j C_j$ in which each C_j is an OR of literals and no variable occurs in more than one literal of the formula. A read-once CNF with parities is a conjunction

$$F(x) = \bigwedge_j T_j(x_{B_j}),$$

where the sets $B_j \subseteq [n]$ are pairwise disjoint and each nonconstant T_j is either an OR of literals or an equation $\bigoplus_{i \in B_j} x_i = b_j$ with $b_j \in \{0,1\}$. The width of a term is $|B_j|$.

Small-bias distributions. A distribution M on $\{0,1\}^N$ is δ-*biased* if

$$\left|\mathbb{E}_{z \sim M}\left[(-1)^{\sum_{i \in S} z_i}\right]\right| \leq \delta \qquad (\varnothing \neq S \subseteq [N]).$$

Naor and Naor [NN93] give explicit such distributions samplable with $O(\log(N+1) + \log(1/\delta))$ uniform seed bits. We use small bias to choose the partition and to satisfy the final parity constraints.

Probably uniform distributions. A distribution X on $\{0,1\}^N$ is d-*wise* η-*probably uniform* if every event E depending on at most d coordinates satisfies

$$\Pr[X \in E] \geq (1-\eta)\Pr[\mathcal{U}_N \in E].$$

For $0 \leq d \leq N$ and $0 < \eta < 1$, Hoza and Lv [HL25, Theorem 3] give explicit distributions with this property using $O(d + \log(1/\eta) + \log\log(N+2))$ seed bits. For $d = 0$ the condition is automatic. These are the round sources used to meet local sufficient conditions and control the next potential.

The width-3 hitting reduction. We use the following structural theorem from GMRTV. Its one-sided inclusion is the reason a hitting-set construction for $\mathrm{CNF}^{\oplus}$ suffices.

Theorem 1.4 (GMRTV reduction [GMR+12, Theorem 6.2])**.** *There is an absolute constant C_{G} such that, for every ordered width-3 ROBP f on n variables with uniform acceptance probability $\varepsilon > 0$, there are a prefix length $k \in \{0, \dots, n\}$ and a read-once $\mathrm{CNF}^{\oplus}$ formula φ on the remaining $n-k$ variables such that, for every $x \in \{0,1\}^{n-k}$,*

$$\varphi(x) = 1 \quad \Longrightarrow \quad f(0^k x) = 1,$$

and

$$\mathbb{E}[\varphi] \geq \left(\frac{\varepsilon}{n}\right)^{C_{\mathrm{G}}}.$$

The formula in this theorem can depend on the branching program. The generator only needs to enumerate the prefix length and use a single HSG that works for every formula with the stated density. We carry this out in Section 4.7.

1.5 The PRG iterated-restriction framework

The iterated-restriction framework goes back to Ajtai and Wigderson [AW85]. The basic derandomization idea is to simplify a Boolean function through a sequence of partial assignments, while replacing truly random restrictions by restrictions drawn from small explicit pseudorandom families. GMRTV later developed a particularly useful version based on *iterated mild restrictions*, where each round fixes only a controlled fraction of the variables that are still unassigned [GMR+12].

A *restriction* is a string

$$\pi \in \{0, 1, *\}^n,$$

where $\pi_i \in \{0,1\}$ fixes x_i, while $\pi_i = *$ leaves x_i unassigned. Applying π to a Boolean function F produces the residual function $F \restriction \pi$ on the remaining variables. The purpose of a restriction is to simplify F without assigning the entire input at once.

The framework repeats this operation for several rounds. After rounds $1, \dots, k-1$, let $\pi^{(k-1)}$ be the cumulative restriction, let R_{k-1} be the set of variables still unassigned, and let

$$F_{k-1} := F \restriction \pi^{(k-1)}$$

be the current residual function. In round k, one would ideally sample a useful random restriction on R_{k-1}. For a pseudorandom generator, however, this restriction must be sampled from a small explicit family $\mathcal{D}_k(R_{k-1})$, so that a short seed y_k selects

$$\sigma^{(k)} \in \mathcal{D}_k(R_{k-1}).$$

The restriction $\sigma^{(k)}$ fixes some variables in R_{k-1} and leaves the others unassigned, producing

$$F_k = F_{k-1} \restriction \sigma^{(k)}.$$

The challenge is therefore not merely to show that a good random restriction exists, but to construct a small pseudorandom space whose behavior is sufficient for the argument.

For a PRG, the crucial requirement is that each pseudorandom restriction approximately preserves the acceptance probability of the residual function produced by the preceding rounds. Schematically, one seeks

$$\mathbb{E}_{\sigma^{(k)}\sim\mathcal{D}_k(R_{k-1})}\left[\mathbb{E}_x\left[(F_{k-1}\restriction\sigma^{(k)})(x)\right]\right]\approx\mathbb{E}_x[F_{k-1}(x)].$$

Importantly, this guarantee is conditional on the earlier rounds: once they determine a particular residual function F_{k-1}, the round-k family must still preserve the relevant quantity for that residual function. This gives the iterative evolution

$$F_0\longrightarrow F_1\longrightarrow\cdots\longrightarrow F_T.$$

The restriction process usually stops once F_T belongs to a simpler class for which a suitable PRG is already known. A complete seed therefore has the form

$$y=(y_1,\dots,y_T,z),$$

where y_k selects the pseudorandom restriction used in round k, and z generates values for the variables that remain after the final restriction. If round k loses at most δ_k in acceptance probability and the final generator has error δ_{final}, a standard hybrid argument gives total error at most

$$\delta_1+\cdots+\delta_T+\delta_{\text{final}}.$$

This is the framework to keep in mind when comparing PRGs with the one-sided HSG construction below. Both proceed by assigning variables through a sequence of pseudorandom rounds drawn from small explicit spaces, but a PRG must approximately preserve acceptance probability throughout the process, whereas an HSG only needs its support to contain one successful sequence of round assignments.

2 The potential-descent framework for hitting-set generators

Section 1.5 describes a PRG construction that repeatedly restricts a function while approximately preserving its acceptance probability. An HSG needs a different conclusion: one sequence of choices in the generator's support must end at an accepting input. This permits a corresponding one-sided framework in which a *potential* controls the total profile size of that sequence instead of tracking the acceptance probability of the residual function. We first present the framework for a Boolean function $F:\{0,1\}^n\to\{0,1\}$, and then specify the objects used for read-once CNF. The framework gives sufficient conditions for a short-seed HSG; obtaining those conditions is the class-specific part of the proof.

2.1 Rounds, histories, and obligations

As in the PRG framework, we assign variables gradually. Here we separate choosing *when* each variable is assigned from choosing its value: fix an ordered partition $\mathcal{P}=(I_1,\dots,I_J)$ of $[n]$, and assign all coordinates in I_k in round k. A round assignment is $y_k\in\{0,1\}^{I_k}$, and $y_{<k} := (y_1, \dots, y_{k-1})$ is the history before round k, with $y_{<1} = \varnothing$. The remaining coordinate set is $R_{k-1} := \bigcup_{h=k}^{J} I_h$, and the residual function is $F \restriction y_{<k}$ on R_{k-1}. A full path $y = (y_1, \dots, y_J)$ denotes the assignment obtained by placing each y_k on its coordinate set I_k.

Instead of asking each round to preserve the residual acceptance probability, we specify a *round obligation*: a condition that the next assignment must satisfy. Write $\mathcal{S}_k^{F,\mathcal{P}}(y_{<k}) \subseteq \{0,1\}^{I_k}$ for the assignments meeting this obligation. A history is valid if every preceding assignment met its obligation; the empty history is valid. The obligations must be chosen so that a valid history always has a valid next assignment, and meeting all obligations implies acceptance:

$$\mathcal{S}_k^{F,\mathcal{P}}(y_{<k}) \neq \varnothing \quad \text{at every valid history,} \qquad \big[\, y_k \in \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}) \text{ for all } k \,\big] \Longrightarrow F(y) = 1.$$

A full sequence meeting every obligation is a *valid path.* These requirements are part of the design, not consequences of merely dividing the variables into rounds.

The HSG will realize such a path through one source for each round. At a high level, let

$$G_k : \{0,1\}^{s_k} \longrightarrow \{0,1\}^{I_k}$$

be the source used in round k; each of its outputs assigns all $|I_k|$ variables of that round. We do not require every output of G_k to satisfy every obligation that might arise. Instead, our goal is to find one valid path $y = (y_1, \dots, y_J)$ such that, for every k,

$$y_k \in \operatorname{supp}(G_k) \cap \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}).$$

Equivalently, every history $y_{<k}$ on this path is produced by the supports of the preceding round sources. The round sources themselves must be fixed without knowing which formula-dependent history will ultimately be the successful one. The next question is therefore what structure of the valid sets $\mathcal{S}_k^{F,\mathcal{P}}(y_{<k})$ lets small, fixed sources contain such a path.

2.2 Local conditions and short round sources

The round-by-round formulation above replaces the single global requirement $F(y) = 1$ by a sequence of requirements

$$y_k \in \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}), \qquad k \in [J].$$

This raises a new question. The set $\mathcal{S}_k^{F,\mathcal{P}}(y_{<k})$ depends on the unknown function and on the history reached in the preceding rounds, whereas the round source used by the HSG must be fixed before either is known. Why should one fixed small source contain a point of every valid set that may arise?

For an arbitrary nonempty subset of $\{0,1\}^{I_k}$ this is impossible. The valid set could be a singleton, and a source that had to hit every singleton would have to contain the entire cube. We therefore need the round obligations to have additional structure. In nonfinal rounds, we use the following property: membership in the full valid set can be guaranteed by satisfying a much smaller *local* condition.

Formally, for every nonfinal round $k < J$ and every valid history, we assign a nonnegative integer $K_k^{F,\mathcal{P}}(y_{<k})$ and identify a nonempty event $\mathcal{T}_k^{F,\mathcal{P}}(y_{<k})$ that depends on at most $K_k^{F,\mathcal{P}}(y_{<k})$ coordinates and satisfies

$$\mathcal{T}_k^{F,\mathcal{P}}(y_{<k}) \subseteq \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}).$$

Thus the source does not need to recognize or hit the full set $\mathcal{S}_k^{F,\mathcal{P}}(y_{<k})$ directly. It is enough that its support contain one assignment satisfying the local sufficient condition $\mathcal{T}_k^{F,\mathcal{P}}(y_{<k})$.

For example, consider the read-once CNF instance developed below. Suppose that, at a particular history $y_{<k}$, exactly K still-unsatisfied clauses must be completed in round k. The valid set $\mathcal{S}_k^{F,\mathcal{P}}(y_{<k})$ consists of all assignments to I_k that satisfy all K clauses. This set may depend on many coordinates, since a clause can contain several variables of I_k. However, we may choose one literal from each current clause and require only these K literals to take their satisfying values. Because the CNF is read-once, the chosen literals belong to K distinct variables. Let $\mathcal{T}_k^{F,\mathcal{P}}(y_{<k})$ be the event that all K chosen literals are satisfied. Then

$$\mathcal{T}_k^{F,\mathcal{P}}(y_{<k}) \subseteq \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}),$$

and $\mathcal{T}_k^{F,\mathcal{P}}(y_{<k})$ depends on exactly K coordinates. In this example we take $K_k^{F,\mathcal{P}}(y_{<k}) = K$.

We call $K_k^{F,\mathcal{P}}(y_{<k})$ the *round parameter*. When F and $\mathcal{P}$ are fixed, we suppress the superscripts. In a nonfinal round, it bounds the number of coordinates used by the sufficient condition $\mathcal{T}_k$, even though the round assigns all $|I_k|$ coordinates. We also assign a nonnegative integer parameter K_J to each valid history before the final round; its role is explained below. The sequence of round parameters along a path y is its *round profile*, and its *total profile size* is

$$\sum_{k=1}^{J} K_k^{F,\mathcal{P}}(y_{<k}).$$

The seed-length calculation below will use this sum to bound the total number of random bits used by the round sources.

The source must be formula-oblivious, so it cannot know which coordinates the local certificate will use or which values they will require. It must work simultaneously for every such choice. This is exactly the role of *probable uniformity*: a distribution X on $\{0,1\}^{I_k}$ is d-wise η-probably uniform if every event E depending on at most d coordinates satisfies

$$\Pr[X \in E] \geq (1-\eta)\Pr[U \in E], \qquad U \sim \mathcal{U}_{I_k}, \quad 0 < \eta < 1.$$

A nonempty event depending on at most K Boolean coordinates has uniform probability at least 2^{-K}. Hence, at a history with $K = K_k(y_{<k})$, a K-wise η-probably uniform source satisfies

$$\Pr[X \in \mathcal{S}_k(y_{<k})] \geq \Pr[X \in \mathcal{T}_k(y_{<k})] \geq (1-\eta)2^{-K} > 0.$$

Therefore its support contains a valid round assignment, even though the source does not know the obligation. The key gain is that one short seed can determine the complete string in $\{0,1\}^{I_k}$ while its seed length is governed by the round parameter K, rather than by $|I_k|$.

The final round may use a different sufficient condition. There is no future potential to preserve: its source only needs to contain an assignment meeting the final obligation. Thus K_J measures the parameter needed by that source and need not bound the number of coordinates in a local event. The RO-CNF construction can use the same local approach in the final round; the construction with parities will use a separate small-bias argument.

2.3 The potential and one-round descent

The round parameter $K_k(y_{<k})$ concerns the current obligation. To bound the total profile size, we must also control the parameters needed in later rounds. Two assignments meeting the current obligation can lead to different round parameters later. We therefore need to keep track of how the current assignment affects their sum.

For this purpose we introduce a nonnegative, history-dependent potential

$$\Phi_k^{F,\mathcal{P}}(y_{<k}).$$

Its intended role is to measure the current round parameter together with an estimate of the sum of the parameters needed in later rounds. We require

$$\Phi_k(y_{<k}) \geq K_k(y_{<k}), \qquad \Phi_J(y_{<J}) = K_J(y_{<J}),$$

and write $\Phi_1(F,\mathcal{P})$ for the initial value before any assignments are made.

The goal in a nonfinal round is to find a valid next assignment for which the potential decreases by approximately the current round parameter. Setting

$$\alpha := 1 + \frac{1}{20J},$$

we seek

$$\Phi_{k+1}(y_{\leq k}) \leq \alpha\big(\Phi_k(y_{<k}) - K_k(y_{<k})\big).$$

With descent in every nonfinal round and a supported valid assignment in the final round, the inequalities telescope and give

$$\sum_{k=1}^{J} K_k^{F,\mathcal{P}}(y_{<k}) = O(\Phi_1(F,\mathcal{P}))\,.$$

However, the assignment achieving this decrease must come from the *actual short round source*. The source must therefore do more than realize the K_k-local sufficient condition for the current obligation: it must also provide enough pseudorandomness to control how the current assignment changes the future part of the potential. We choose a threshold t_k for this purpose.

Fix nonnegative integers $t_1, \ldots, t_J$, with $t_J = 0$. For every round k and integer source parameter $r \geq 0$, let

$$G_{k,r} : \{0,1\}^{s_{k,r}} \to \{0,1\}^{I_k}$$

be a fixed round source. For each nonfinal round $k < J$, take its output under a uniform seed to be $(r+t_k)$-wise $\eta_{k,r}$-probably uniform, where

$$\eta_{k,r} := \frac{2^{-(r+t_k)}}{100J}.$$

Hoza and Lv [HL25, Theorem 3] give such sources with

$$s_{k,r} = O(r + t_k + \log(J+1) + \log\log(n+2))\,.$$

The source can be fixed on n coordinates and projected onto I_k; apart from this projection, it depends only on n, k, r, t_k, J, and not on F or the history reached so far.[1]

For the final source family, the class-specific proof must establish

$$\operatorname{supp}\bigl(G_{J,K_J(y_{<J})}\bigr) \cap \mathcal{S}_J(y_{<J}) \neq \varnothing \qquad \text{at every valid history } y_{<J}. \tag{1}$$

This family is also fixed independently of F and the history, apart from projection onto I_J. It may use the same probable-uniformity construction with $t_J = 0$, or a different construction suited to the final obligations. The composition theorem below allows an additional $O(\log(n+2))$ seed bits in this one round.

Thus the potential and the round source are designed together. The round parameter K_k in a nonfinal round bounds the locality of the sufficient condition for the current obligation. The threshold t_k bounds the additional coordinates used to analyze the future potential.

Finding a good assignment inside the round-source support. Fix a nonfinal round $k < J$ and a valid history $y_{<k}$. Our goal is to find a round assignment y_k satisfying two properties simultaneously: it must belong to the support of the round source and meet the current obligation, and it must satisfy the potential-descent inequality. We now derive a useful condition stated directly in terms of the round-source distribution that is sufficient for proving this potential-descent inequality:

$$\Phi_{k+1}(y_{\le k}) \le \alpha\bigl(\Phi_k(y_{<k}) - K_k(y_{<k})\bigr).$$

The round source itself gives a natural way to prove that such an assignment exists. Set

$$K := K_k(y_{<k}),$$

let Z_k be a uniform seed for $G_{k,K}$, and write

$$X_k := G_{k,K}(Z_k).$$

Thus X_k is a random assignment drawn from the actual support available to the HSG in round k.

Recall that $\mathcal{T}_k(y_{<k}) \subseteq \mathcal{S}_k(y_{<k})$ is a local sufficient condition for meeting the current obligation. Therefore, every source output satisfying

$$X_k \in \mathcal{T}_k(y_{<k})$$

is already a valid round assignment. By probable uniformity,

$$\Pr[X_k \in \mathcal{T}_k(y_{<k})] > 0,$$

so there is at least one such output in the source support.

We now need to ensure that among these valid source outputs, one also leaves a sufficiently small future potential. It is enough to prove that their average next potential satisfies

$$\mathbb{E}[\Phi_{k+1}(y_{<k}, X_k) \mid X_k \in \mathcal{T}_k(y_{<k})] \le \alpha\bigl(\Phi_k(y_{<k}) - K\bigr).$$

[1] If $r + t_k$ exceeds $|I_k|$, the probable-uniformity requirement concerns all subsets of the available coordinates.

Indeed, the conditional expectation is an average over a nonempty collection of actual outputs of $G_{k,K}$, all of which meet the current obligation. Hence at least one of these outputs, say y_k, has next potential no larger than the average. For this y_k,

$$y_k \in \operatorname{supp}(G_{k,K}) \cap \mathcal{S}_k(y_{<k})$$

and

$$\Phi_{k+1}(y_{\le k}) \le \alpha\big(\Phi_k(y_{<k}) - K\big).$$

Therefore the remaining task in designing the potential is clear: its one-round change must be simple enough that we can bound this conditional expectation using the local pseudorandomness of $G_{k,K}$. The conditioning is only part of the existence proof; the HSG itself does not need to recognize $\mathcal{T}_k$ or sample from a conditional distribution.

Composing the rounds. Starting with the empty history, choose an extension satisfying the two conditions above in every nonfinal round, and a supported valid assignment in the final round using (1). Inductively, if the earlier assignments lie in their round-source supports, adjoining one more supported output preserves that property. The resulting full path is valid and hence accepted by F.

The bound on the total profile size follows by composing the potential inequalities along this path. Abbreviate $K_k = K_k(y_{<k})$ and $\Phi_k = \Phi_k(y_{<k})$. The descent inequality gives

$$\frac{K_k}{\alpha^{k-1}} \le \frac{\Phi_k}{\alpha^{k-1}} - \frac{\Phi_{k+1}}{\alpha^k} \qquad (k < J).$$

Summing over the nonfinal rounds and using $\Phi_J = K_J$ yields

$$\sum_{k=1}^{J} \frac{K_k}{\alpha^{k-1}} \le \Phi_1, \qquad \sum_{k=1}^{J} K_k \le \alpha^{J-1}\Phi_1 \le e^{1/20}\Phi_1.$$

In Section 1.5, one-step preservation errors add to bound the final fooling error. Here the rescaled potential inequalities telescope to bound the total profile size of one supported accepting path. Neither statement says that every seed produces a favorable path.

2.4 Conditions for a function class

We can now summarize the potential-descent framework as a concrete recipe for constructing an HSG. Let $\mathcal{F}_n$ be a family of Boolean functions $F : \{0,1\}^n \to \{0,1\}$. For each ordered partition $\mathcal{P} = (I_1, \ldots, I_J)$ and each valid history $y_{<k}$, the class-specific construction must specify a nonempty round obligation

$$\mathcal{S}_k^{F,\mathcal{P}}(y_{<k}) \subseteq \{0,1\}^{I_k}$$

such that satisfying the obligation in every round implies $F(y) = 1$, and a nonnegative integer round parameter $K_k^{F,\mathcal{P}}(y_{<k})$. For each nonfinal round $k < J$, it must also specify a nonempty local sufficient event

$$\mathcal{T}_k^{F,\mathcal{P}}(y_{<k}) \subseteq \mathcal{S}_k^{F,\mathcal{P}}(y_{<k})$$

depending on at most $K_k^{F,\mathcal{P}}(y_{<k})$ coordinates. In the final round, it must instead verify the support condition (1) for the chosen source family and parameter K_J. The total profile size of a path is

$$\sum_{k=1}^{J} K_k^{F,\mathcal{P}}(y_{<k}).$$

Finally, the class-specific analysis must define a nonnegative potential $\Phi_k^{F,\mathcal{P}}(y_{<k})$ with $\Phi_k \geq K_k$ and $\Phi_J = K_J$ for which the round-source argument above gives supported descent in every nonfinal round. The purpose of these definitions is to prove two separate theorems: one for every fixed partition, and one selecting a useful partition from a small formula-oblivious family.

First theorem: a supported path with bounded profile size. Fix the number of rounds J, thresholds $t_1, \ldots, t_J$ with $t_J = 0$, and formula-oblivious round sources $G_{k,r}$, using the convention that each $G_{k,r}$ is obtained by projecting a fixed n-coordinate source onto the current part I_k. The first class-specific theorem must show that for every $F \in \mathcal{F}_n$ and every ordered partition $\mathcal{P}$, there is a valid path in these source supports whose total profile size is at most a constant times the initial potential. Its quantifier order is

$$\forall F \in \mathcal{F}_n \; \forall \mathcal{P} \; \exists q \; \exists y.$$

The successful profile and path may depend on F and $\mathcal{P}$, while the source family is fixed independently of them.

Theorem 2.1 (Abstract fixed-partition supported-path theorem)**.** *For the chosen function family $\mathcal{F}_n$, suppose that the obligations, round parameters, source families, and potential can be chosen so that there is an absolute constant C_{path} with the following property. For every $F \in \mathcal{F}_n$ and every ordered partition $\mathcal{P} = (I_1, \ldots, I_J)$, there exist a profile $q = (q_1, \ldots, q_J) \in \mathbb{Z}_{\geq 0}^J$ and an assignment $y = (y_1, \ldots, y_J)$, with $y_k \in \{0,1\}^{I_k}$, such that for every $k \in [J]$,*

$$q_k = K_k^{F,\mathcal{P}}(y_{<k}), \qquad y_k \in \operatorname{supp}(G_{k,q_k}) \cap \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}),$$

and

$$\sum_{k=1}^{J} q_k = \sum_{k=1}^{J} K_k^{F,\mathcal{P}}(y_{<k}) \leq C_{\text{path}} \, \Phi_1(F, \mathcal{P}).$$

In particular, $F(y) = 1$.

The discussion above gives the standard way to prove this theorem for a particular class: in each nonfinal round, construct $\mathcal{T}_k$ and Φ_k so that, for the actual round source $X_k = G_{k,K_k}(Z_k)$, conditioning on $X_k \in \mathcal{T}_k$ has positive probability and the conditional expected next potential satisfies the one-round descent inequality. Verify the final support condition separately. The supported choices then compose, and the rescaled potential inequalities telescope using $\Phi_J = K_J$.

Second theorem: a small partition family containing a good partition. The fixed-partition theorem does not guarantee that $\Phi_1(F, \mathcal{P})$ is small. The second class-specific theorem must therefore construct, before the unknown function is known, a small explicit distribution over partitions such that every sufficiently dense function has a useful partition in its support. We allow a one-sided preprocessing step: the theorem may replace F by a function $\widehat{F} \in \mathcal{F}_n$ satisfying $\widehat{F}(x) = 1 \Rightarrow F(x) = 1$. In applications where no preprocessing is needed, simply take $\widehat{F} = F$.

Theorem 2.2 (Abstract partition-selection theorem)**.** *For every n and $\varepsilon \in (0, 1]$, suppose there exist a number of rounds J, nonnegative thresholds $t_1, \ldots, t_J$ with $t_J = 0$, a bound $B = B(n, \varepsilon)$, and an explicit distribution $\mathfrak{P}_{n,\varepsilon}$ over ordered partitions of $[n]$, samplable using s_{part} bits, such that these objects depend only on n and ε and the following holds. For every $F \in \mathcal{F}_n$ with*

$$\Pr_{x \sim \mathcal{U}_n}[F(x) = 1] \geq \varepsilon,$$

there exist a function $\widehat{F} \in \mathcal{F}_n$ and a partition $\mathcal{P} \in \mathrm{supp}(\mathfrak{P}_{n,\varepsilon})$ such that

$$\widehat{F}(x) = 1 \Longrightarrow F(x) = 1 \quad \textit{for every } x \in \{0,1\}^n, \qquad \Phi_1(\widehat{F}, \mathcal{P}) \leq B(n, \varepsilon).$$

The important quantifier order is that J, the thresholds, and the entire partition family are fixed first, using only n and ε; after an eligible F is given, the proof may identify one useful $\widehat{F}$ and one useful partition in that fixed support. A convenient sufficient way to obtain the last inequality is an average bound on Φ_1 over $\mathfrak{P}_{n,\varepsilon}$. For an optimal construction, the targets are

$$B(n, \varepsilon) = O\left(1 + \log \frac{1}{\varepsilon}\right), \qquad \sum_{k=1}^{J} t_k = O(\log n), \qquad J = O(\log \log n), \qquad s_{\mathrm{part}} = O(\log n).$$

2.5 From the two theorems to one HSG

Once the two class-specific statements above are available, the rest of the construction is generic. The following theorem records the precise seed and output maps and proves the resulting hitting guarantee.

Theorem 2.3 (HSG consequence of the two framework theorems)**.** *Assume Theorems 2.1 and 2.2 hold for $\mathcal{F}_n$. Assume also that for every round k and source parameter r the fixed round source has seed length*

$$s_{k,r} = O(r + t_k + \log(J + 1) + \log \log(n + 2) + \mathbf{1}[k = J] \log(n + 2)) .$$

The last term permits an additional $O(\log(n + 2))$ seed bits in the final round, including a final source with seed length $O(r + \log(n + 2))$. Then there is an explicit HSG $H_{n,\varepsilon}$ for all $F \in \mathcal{F}_n$ of density at least ε with seed length

$$s_{\mathrm{part}} + O\left(B + J + \sum_{k=1}^{J} t_k + J \log(J + 1) + J \log \log(n + 2) + \log(n + 2)\right).$$

In particular, if

$$B = O\left(1 + \log\frac{1}{\varepsilon}\right), \qquad \sum_k t_k = O(\log n), \qquad J = O(\log\log n), \qquad s_{\mathrm{part}} = O(\log n),$$

then $H_{n,\varepsilon}$ *has seed length* $O(\log(n/\varepsilon))$.

Proof. Let C_{path} be the constant from Theorem 2.1, set the *profile budget*

$$L := \lceil C_{\mathrm{path}} B \rceil,$$

and enumerate the round profiles

$$\mathcal{Q}_L := \left\{ q = (q_1, \dots, q_J) \in \mathbb{Z}_{\geq 0}^J : \sum_{k=1}^{J} q_k \leq L \right\}.$$

Let

$$P : \{0,1\}^{s_{\mathrm{part}}} \to \mathrm{Part}_J([n])$$

be the sampler for $\mathfrak{P}_{n,\varepsilon}$. We first describe the HSG by a structured seed

$$\xi = (u, q, z_1, \dots, z_J),$$

where $u \in \{0,1\}^{s_{\mathrm{part}}}$, $q \in \mathcal{Q}_L$, and $z_k \in \{0,1\}^{s_{k,q_k}}$. If $P(u) = \mathcal{P} = (I_1, \dots, I_J)$, then $H_{n,\varepsilon}(\xi)$ is the unique assignment $y \in \{0,1\}^n$ satisfying

$$y|_{I_k} = G_{k,q_k}(z_k) \qquad \text{for every } k \in [J].$$

Therefore a seed specifies three things: an ordered partition, a round profile, and one seed for the corresponding round source in each round.

We next bound the number of bits needed to encode this structured seed. By stars and bars,

$$|\mathcal{Q}_L| = \binom{L+J}{J} \leq 2^{L+J},$$

so the profile uses $O(B + J)$ bits. For every fixed $q \in \mathcal{Q}_L$,

$$\begin{aligned}
\sum_{k=1}^{J} s_{k,q_k} &= O\left(\sum_{k=1}^{J} q_k + \sum_{k=1}^{J} t_k + J\log(J+1) + J\log\log(n+2) + \log(n+2)\right) \\
&= O\left(B + \sum_{k=1}^{J} t_k + J\log(J+1) + J\log\log(n+2) + \log(n+2)\right).
\end{aligned}$$

The additional $\log(n+2)$ term appears only once, from the final source. Adding the partition seed and the profile encoding gives the claimed seed length. A fixed encoding of q, together with padding of the *combined* round-seed string to the largest total length over $q \in \mathcal{Q}_L$, turns this structured description into an ordinary fixed-length binary seed. Under the quantitative bounds in the theorem, the total length is $O(\log(n/\varepsilon))$.

It remains to prove hitting. Fix $F \in \mathcal{F}_n$ with density at least ε. By Theorem 2.2, there are $\widehat{F} \in \mathcal{F}_n$ and $\mathcal{P} \in \operatorname{supp}(\mathfrak{P}_{n,\varepsilon})$ such that $\widehat{F}(x) = 1 \Rightarrow F(x) = 1$ for every x and $\Phi_1(\widehat{F}, \mathcal{P}) \le B$. Applying Theorem 2.1 to $\widehat{F}$ and this partition gives a valid path $y = (y_1, \dots, y_J)$ and a profile q satisfying

$$\sum_{k=1}^{J} q_k \le C_{\text{path}} \Phi_1(\widehat{F}, \mathcal{P}) \le L, \qquad y_k \in \operatorname{supp}(G_{k,q_k}) \quad \text{for every } k.$$

Hence $q \in \mathcal{Q}_L$. The useful partition has some partition seed u, and for every k there is a round seed z_k with $G_{k,q_k}(z_k) = y_k$. The structured seed $\xi = (u, q, z_1, \dots, z_J)$ is therefore among the seeds enumerated by the fixed generator and satisfies $H_{n,\varepsilon}(\xi) = y$. The supported-path theorem gives $\widehat{F}(y) = 1$, and the one-sided implication then gives $F(y) = 1$. Therefore

$$\exists\, H_{n,\varepsilon} \quad \forall F \in \mathcal{F}_n \text{ of density at least } \varepsilon \quad \exists \xi : \quad F(H_{n,\varepsilon}(\xi)) = 1,$$

which is exactly the HSG guarantee. □

In this form, the class-specific work is concentrated in the two theorem statements above. The profile enumeration, the composition of the round sources, and the final seed-length calculation are generic.

3 Hitting read-once CNFs

We now prove Theorem 1.3 using the framework from Section 2. We first state the two properties needed for this class, then prove them and combine them into the generator.

3.1 Round obligations and the two main theorems

We now specialize the abstract framework of Section 2 to a read-once CNF

$$F = C_1 \wedge \cdots \wedge C_m$$

on $[n]$. For a fixed partition $\mathcal{P} = (I_1, \dots, I_J)$, let $\ell_j = \ell_j^{\mathcal{P}}$ be the last round containing a variable of C_j, called its *completion round*. A clause is *alive* at a history if the earlier assignments have not satisfied it. The round-k obligation is to satisfy every alive clause completed in that round. In the notation of the framework, this means

$$\mathcal{S}_k^{F,\mathcal{P}}(y_{<k}) := \{a \in \{0,1\}^{I_k} : a \text{ satisfies every alive } C_j \text{ with } \ell_j = k\},$$
$$K_k^{F,\mathcal{P}}(y_{<k}) := \#\{j : \ell_j = k \text{ and } C_j \text{ is alive after } y_{<k}\}.$$

We choose one literal with a variable in I_k from each of these clauses and require its satisfying value. The resulting event $\mathcal{T}_k(y_{<k})$ uses exactly $K_k(y_{<k})$ distinct coordinates, since the clauses are read-once, and is contained in $\mathcal{S}_k(y_{<k})$. Therefore the valid set is nonempty. Every valid path satisfies F, because each clause is either satisfied earlier or satisfied when its completion round is reached. Each clause that survives until its completion round contributes one to the total profile size.

The truncated potential. Fix nonnegative integer thresholds $t_1, \dots, t_J$, with $t_J = 0$. For each clause C_j, let B_j be its variable set and put $D_{j,h} := |B_j \cap I_h|$. We use the potential

$$\begin{aligned}\Phi_k(y_{<k}) &:= K_k(y_{<k}) \\ &\quad + \sum_{\substack{j:\ C_j \text{ is alive after } y_{<k} \\ \ell_j > k}} \prod_{h=k}^{\ell_j - 1} 2^{-\min(D_{j,h}, t_h)}.\end{aligned} \tag{2}$$

The first term counts the current clauses. Each later clause has a weight that accounts for its possible survival through the intervening rounds, using at most t_h variables in round h. We explain these factors in Section 3.2 and prove that they give the required conditional descent. In the final round the sum is empty, so $\Phi_J(y_{<J}) = K_J(y_{<J})$.

At the empty history every clause is alive. Using the convention that an empty product equals one, the initial potential is

$$\Phi_1(F, \mathcal{P}) = \sum_{j=1}^{m} \prod_{h<\ell_j} 2^{-\min(D_{j,h}, t_h)}.$$

In particular, a clause completing in round 1 contributes one. The potential depends on F, $\mathcal{P}$, and the fixed thresholds; we suppress these dependencies in $\Phi_k(y_{<k})$ and the threshold dependence in $\Phi_1(F, \mathcal{P})$.

The two remaining tasks are to find a supported valid path whose total profile size is bounded by this initial potential, and to construct a small explicit partition family containing a partition with small initial potential. These are the contents of Theorems 3.1 and 3.2 below. The second theorem also includes a width-shortening step in the analysis; every satisfying assignment of the shortened formula still satisfies the original formula.

Theorem 3.1 (Fixed-partition supported-path theorem)**.** *Let F be a read-once CNF on $[n]$, let $\mathcal{P} = (I_1, \dots, I_J)$ be an ordered partition of $[n]$, and let $t_1, \dots, t_J$ be nonnegative integers, with $t_J = 0$. Use the potential in (2) with these thresholds. For every $k \in [J]$ and integer $r \geq 0$, let*

$$G_{k,r} : \{0,1\}^{s_{k,r}} \to \{0,1\}^{I_k}$$

be an explicit $(r + t_k)$-wise $\eta_{k,r}$-probably uniform source, where

$$\eta_{k,r} \leq \frac{2^{-(r+t_k)}}{100J} \qquad \text{and} \qquad s_{k,r} = O(r + t_k + \log J + \log\log n)\,.$$

Apart from projection onto I_k, the family $\{G_{k,r}\}$ depends only on n, k, r, t_k and the partition size J, and is independent of both F and the particular partition $\mathcal{P}$.

Then there exist a profile $q = (q_1, \dots, q_J) \in \mathbb{Z}_{\geq 0}^J$ and an assignment $y = (y_1, \dots, y_J)$, where $y_k \in \{0,1\}^{I_k}$, such that for every $k \in [J]$,

$$q_k = K_k^{F,\mathcal{P}}(y_{<k}), \qquad y_k \in \operatorname{supp}(G_{k,q_k}) \cap \mathcal{S}_k^{F,\mathcal{P}}(y_{<k}),$$

and

$$\sum_{k=1}^{J} q_k = \sum_{k=1}^{J} K_k^{F,\mathcal{P}}(y_{<k}) \leq C\, \Phi_1(F, \mathcal{P})$$

for an absolute constant C. In particular, $F(y) = 1$. Moreover, every history $y_{<k} = (y_1, \dots, y_{k-1})$ lies in

$$\mathrm{supp}(G_{1,q_1}) \times \cdots \times \mathrm{supp}(G_{k-1,q_{k-1}}),$$

so the entire valid path is contained in the support of the concrete formula-oblivious round sources indexed by the profile q.

Theorem 3.2 (Partition-selection theorem)**.** *Let $n \geq 1$ and $\varepsilon \in (0, 1]$, and set*

$$\Lambda := 1 + \log(1/\varepsilon) \qquad \text{and} \qquad W := \lceil \log(8n) \rceil.$$

There exist $J = \Theta(\log \log n)$, thresholds $t_1, \dots, t_J \geq 0$ with $t_J = 0$ and

$$\sum_{k=1}^{J} t_k = O(\log n),$$

and an explicit distribution $\mathfrak{P}$ over ordered partitions $\mathcal{P} = (I_1, \dots, I_J)$ of $[n]$, samplable using $O(\log n)$ random bits, such that the following holds.

For every read-once CNF $F = C_1 \wedge \cdots \wedge C_m$ on $[n]$ of density at least ε, first shorten every clause of width greater than W by retaining any W of its literals, and let $F^{\leq W}$ denote the resulting read-once CNF. Then

$$F^{\leq W}(x) = 1 \quad \Longrightarrow \quad F(x) = 1$$

for every $x \in \{0, 1\}^n$, and the initial value of the potential in (2), with the thresholds above, satisfies

$$\mathbb{E}_{\mathcal{P} \sim \mathfrak{P}} \left[\Phi_1(F^{\leq W}, \mathcal{P}) \right] \leq C\Lambda$$

for an absolute constant C. Consequently, for every such F there exists $\mathcal{P} \in \mathrm{supp}(\mathfrak{P})$ such that

$$\Phi_1(F^{\leq W}, \mathcal{P}) \leq C\Lambda.$$

The parameters J, $t_1, \dots, t_J$, and the distribution $\mathfrak{P}$ depend only on n and ε, and are independent of F.

3.2 Proof of the fixed-partition supported-path theorem

We first prove Theorem 3.1 by proving conditional descent for the truncated potential defined in Section 3.1. In Section 3.3, we will then construct the ordered-partition family and prove Theorem 3.2.

Fix a read-once CNF

$$F = C_1 \wedge \cdots \wedge C_m,$$

an ordered partition

$$\mathcal{P} = (I_1, \dots, I_J),$$

and thresholds $t_1, \dots, t_J \geq 0$, with $t_J = 0$. The round sources $G_{k,r}$ are those appearing in Theorem 3.1. Throughout the proof, $F, \mathcal{P}$, and the thresholds are fixed, and we suppress them from history-dependent notation when no ambiguity arises.

We use the round obligations, local events, and round parameters defined in Section 3.1. Thus ℓ_j is the completion round of C_j; $K_k(y_{<k})$ is the number of clauses that are still alive and complete in round k; and $\mathcal{T}_k(y_{<k}) \subseteq \mathcal{S}_k(y_{<k})$ is obtained by choosing one literal from each such clause and requiring its satisfying value. Since the clauses are read-once, $\mathcal{T}_k(y_{<k})$ depends on exactly $K_k(y_{<k})$ coordinates.

The abstract framework of Section 2 tells us what remains to prove. For the potential in (2), we must show that, at every valid history and every nonfinal round k, if

$$K := K_k(y_{<k}), \qquad X_k := G_{k,K}(Z_k)$$

for a uniform seed Z_k, then

$$\Pr[X_k \in \mathcal{T}_k(y_{<k})] > 0$$

and

$$\mathbb{E}[\Phi_{k+1}(y_{<k}, X_k) \mid X_k \in \mathcal{T}_k(y_{<k})] \leq \alpha\big(\Phi_k(y_{<k}) - K\big), \qquad \alpha := 1 + \frac{1}{20J}.$$

Once this is established, the supported-choice and telescoping argument from Section 2 yields the required path. We now verify precisely this condition for read-once CNF.

Why we truncate the clause counts. Recall that $D_{j,h} = |B_j \cap I_h|$ counts the variables of C_j assigned in round h. Suppose C_j is alive before a round $h < \ell_j$. It survives that round only if all of its $D_{j,h}$ variables in I_h receive their forbidden values. The exact survival event may depend on more than t_h coordinates, while the round source at a history with round parameter K_h gives us pseudorandomness only on $K_h + t_h$ coordinates. We therefore retain only

$$\min(D_{j,h}, t_h)$$

of these variables. Requiring those selected variables to receive their forbidden values gives a local event containing the true survival event and having uniform probability

$$2^{-\min(D_{j,h}, t_h)}.$$

These are the factors used in (2). They are uniform probabilities of local events containing the true survival events. We will use them to bound conditional survival under the round sources and prove the required one-round descent inequality.

The one-round estimate. We first state the estimate without reference to clauses. This will let us use it again when parity constraints are added.

Lemma 3.3 (Conditioning on fixed coordinates)**.** *Let $K, t \geq 0$ and $J \geq 1$ be integers. Let X be a distribution on $\{0,1\}^I$ that is $(K+t)$-wise η-probably uniform, where*

$$0 < \eta \leq \frac{2^{-(K+t)}}{100J}.$$

Suppose T fixes the values of K distinct coordinates, and E depends on at most t other coordinates. For uniform $U \in \{0,1\}^I$,

$$\Pr[X \in T] \geq (1-\eta)2^{-K} > 0, \qquad \Pr[X \in E \mid X \in T] \leq \left(1 + \frac{1}{20J}\right) \Pr[U \in E].$$

Proof. Since T fixes K coordinates, $\Pr[U \in T] = 2^{-K}$, and probable uniformity gives the stated lower bound for $\Pr[X \in T]$. If E is empty, the conditional bound is immediate. Otherwise put $p := \Pr[U \in E]$. A nonempty event on at most t bits has $p \geq 2^{-t}$. The events use disjoint coordinates, so

$$\Pr[U \in T \cap E] = 2^{-K} p.$$

The joint event depends on at most $K + t$ coordinates. Applying probable uniformity to its complement gives

$$\Pr[X \in T \cap E] \leq 2^{-K} p + \eta.$$

Dividing by the lower bound for $\Pr[X \in T]$, we obtain

$$\begin{aligned}\Pr[X \in E \mid X \in T] &\leq \frac{p + \eta 2^K}{1 - \eta} \\ &\leq p \frac{1 + 1/(100J)}{1 - 1/(100J)} = p\left(1 + \frac{2}{100J - 1}\right) \\ &\leq p\left(1 + \frac{1}{20J}\right).\end{aligned}$$

The second line uses $\eta 2^K \leq 2^{-t}/(100J) \leq p/(100J)$ and $\eta \leq 1/(100J)$; the last line follows from $J \geq 1$. $\square$

Now fix a valid history $y_{<k}$ with $k < J$, and put $K := K_k(y_{<k})$ and $X_k := G_{k,K}(Z_k)$ for a uniform seed Z_k. For an alive clause C_j with $\ell_j > k$, set $d_{j,k} := \min(D_{j,k}, t_k)$. Choose $d_{j,k}$ of its variables in I_k, and let $E'_{j,k}$ require their forbidden values. True survival of C_j implies $E'_{j,k}$, whose uniform probability is $2^{-d_{j,k}}$; when $d_{j,k} = 0$, it is the whole space.

The event $\mathcal{T}_k(y_{<k})$ fixes K coordinates from clauses completing now. These coordinates are disjoint from those of $E'_{j,k}$, since F is read-once. Applying Lemma 3.3 with $t = t_k$ gives

$$\Pr[X_k \in \mathcal{T}_k(y_{<k})] \geq (1 - \eta_{k,K}) 2^{-K} > 0, \qquad \Pr[X_k \in E'_{j,k} \mid X_k \in \mathcal{T}_k(y_{<k})] \leq \alpha 2^{-d_{j,k}}.$$

This is the only probabilistic estimate needed in the descent proof.

Verifying the potential-descent condition. For the same future clause C_j, define

$$c_j := \prod_{h=k+1}^{\ell_j - 1} 2^{-\min(D_{j,h}, t_h)}.$$

Before round k, the contribution of this clause to $\Phi_k(y_{<k}) - K$ is $2^{-d_{j,k}} c_j$. After the round, it contributes c_j if it remains alive and 0 if it becomes satisfied. This also includes the case $\ell_j = k + 1$, for which the empty product gives $c_j = 1$, so a surviving clause appears as one unit of K_{k+1}.

Since true survival implies $E'_{j,k}$,

$$\begin{aligned}&\mathbb{E}[\text{contribution of } C_j \text{ to } \Phi_{k+1} \mid X_k \in \mathcal{T}_k(y_{<k})] \\ &\qquad \leq c_j \Pr[E'_{j,k} \mid X_k \in \mathcal{T}_k(y_{<k})] \leq \alpha\, 2^{-d_{j,k}} c_j.\end{aligned}$$

On $\mathcal{T}_k(y_{<k})$, every alive clause completing in round k is satisfied, so such clauses contribute nothing to the next potential. Summing the estimate above over all currently alive clauses with $\ell_j > k$, and using only linearity of expectation, gives

$$\begin{aligned}
&\mathbb{E}[\Phi_{k+1}(y_{<k}, X_k) \mid X_k \in \mathcal{T}_k(y_{<k})] \\
&\le \alpha \sum_{\substack{j:\ C_j \text{ is alive after } y_{<k} \\ \ell_j > k}} 2^{-\min(D_{j,k},t_k)} \prod_{h=k+1}^{\ell_j - 1} 2^{-\min(D_{j,h},t_h)} \\
&= \alpha\big(\Phi_k(y_{<k}) - K_k(y_{<k})\big).
\end{aligned}$$

This is exactly the round-source-induced condition identified in Section 2. In particular, we did not require independence among the future clauses under $G_{k,K}$; each clause was analyzed separately using at most $K + t_k$ coordinates, and the resulting expectations were then summed.

Constructing the supported path. We now apply the consequence proved in Section 2. Start with the empty history. At every nonfinal round k, set

$$q_k := K_k(y_{<k}).$$

The conditioning event for G_{k,q_k} has positive probability, and the conditional expected potential satisfies the required descent inequality. Hence there exists

$$y_k \in \operatorname{supp}(G_{k,q_k}) \cap \mathcal{S}_k(y_{<k})$$

such that

$$\Phi_{k+1}(y_{\le k}) \le \alpha\big(\Phi_k(y_{<k}) - q_k\big).$$

The new history remains valid, so the argument continues.

In the final round, set

$$q_J := K_J(y_{<J}).$$

Here $t_J = 0$ and no future potential remains. The local event $\mathcal{T}_J(y_{<J})$ depends on q_J coordinates, and probable uniformity gives

$$\Pr[G_{J,q_J}(Z_J) \in \mathcal{T}_J(y_{<J})] \ge (1 - \eta_{J,q_J})2^{-q_J} > 0.$$

Therefore we may choose

$$y_J \in \operatorname{supp}(G_{J,q_J}) \cap \mathcal{S}_J(y_{<J}).$$

The resulting path $y = (y_1, \ldots, y_J)$ is valid in every round. By the definition of the read-once CNF obligations in Section 3.1, every clause is satisfied either before or at its completion round, and hence

$$F(y) = 1.$$

Moreover,

$$q_k = K_k^{F,\mathcal{P}}(y_{<k})$$

and every y_k lies in the support of G_{k,q_k}, as required.

Bounding the total profile size. It remains only to bound the total profile size. Along the selected path abbreviate

$$K_k := K_k^{F,\mathcal{P}}(y_{<k}), \qquad \Phi_k := \Phi_k^{F,\mathcal{P}}(y_{<k}).$$

For $k < J$,

$$\Phi_{k+1} \leq \alpha(\Phi_k - K_k),$$

and therefore

$$\frac{K_k}{\alpha^{k-1}} \leq \frac{\Phi_k}{\alpha^{k-1}} - \frac{\Phi_{k+1}}{\alpha^k}.$$

Summing over $k = 1, \dots, J-1$ and using $\Phi_J = K_J$ gives

$$\sum_{k=1}^{J} \frac{K_k}{\alpha^{k-1}} \leq \Phi_1(F, \mathcal{P}).$$

Since $K_k \geq 0$,

$$\begin{aligned}\sum_{k=1}^{J} K_k &\leq \alpha^{J-1} \sum_{k=1}^{J} \frac{K_k}{\alpha^{k-1}} \\ &\leq \alpha^{J-1} \Phi_1(F, \mathcal{P}) \\ &\leq e^{1/20} \Phi_1(F, \mathcal{P}),\end{aligned}$$

where

$$\alpha^{J-1} \leq \left(1 + \frac{1}{20J}\right)^J \leq e^{1/20}.$$

Finally, because $q_k = K_k^{F,\mathcal{P}}(y_{<k})$, we obtain

$$\sum_{k=1}^{J} q_k = \sum_{k=1}^{J} K_k \leq e^{1/20} \Phi_1(F, \mathcal{P}).$$

Together with the validity and source-support properties proved above, this proves Theorem 3.1. □

3.3 Proof of the partition-selection theorem

Theorem 3.1 shows that we can choose an output from each round source so that the resulting assignment satisfies the formula. The sum of the round parameters for these choices is at most a constant times the initial potential. Our task now is to make this potential small by choosing a suitable partition. We will construct a small family of partitions and show that, for every formula with enough satisfying inputs, at least one partition in this family works.

Fix a read-once CNF F on n variables with $\Pr[F(\mathcal{U}_n) = 1] \geq \varepsilon$, and set

$$\Lambda := 1 + \log(1/\varepsilon), \qquad W := \lceil \log_2(8n) \rceil.$$

We will first work with clauses that have at most W variables. For each wider clause, keep any W literals and delete the others. Call the new formula $F^{\leq W}$. Deleting literals from an

OR makes it harder to satisfy, so every assignment satisfying $F^{\le W}$ also satisfies F. We use this shorter formula only in the proof. The generator does not need to know which literals we kept.

We will construct a distribution $\mathfrak{P}$ on ordered partitions, using $O(\log n)$ seed bits, such that

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\le W},\mathcal{P})] = O(\Lambda). \tag{3}$$

If this average is $O(\Lambda)$, at least one partition that the generator can choose has potential $O(\Lambda)$. This is the conclusion needed in Theorem 3.2. The distribution and the choices made for the rounds will depend only on n, so the same generator will work for every F. We first explain what bound we need for each clause and why it is enough. We then show how to choose the partitions to get that bound.

The contribution of one clause. Write $F^{\le W} = C_1 \wedge \cdots \wedge C_m$. For each clause C_j, let $B_j \subseteq [n]$ be the set of indices of the variables appearing in it, and let $w_j := |B_j| \le W$. For example, if $C_j = x_2 \vee \neg x_5 \vee x_9$, then $B_j = \{2,5,9\}$ and $w_j = 3$. The set B_j records which input variables the clause uses.

Fix an ordered partition $\mathcal{P} = (I_1,\dots,I_J)$, where round h assigns the variables whose indices are in I_h. Write $\tau_i = h$ when $i \in I_h$, and let $D_{j,h} := |B_j \cap I_h|$ count how many variables of C_j are assigned in round h. Its completion round $\ell_j := \max_{i\in B_j} \tau_i$ is the last round assigning any of its variables. Before the first round, no variables have been assigned and every clause is still alive. The definition of the potential in Section 3.1 therefore gives

$$\Phi_1(F^{\le W},\mathcal{P}) = K_1 + \sum_{j:\,\ell_j>1} \prod_{h=1}^{\ell_j-1} 2^{-\min(D_{j,h},t_h)}.$$

Here $K_1 = \#\{j : \ell_j = 1\}$: each clause completed in round 1 contributes one. Since an empty product is one, these clauses can be included in the same sum as the other clauses. Thus

$$\Phi_1(F^{\le W},\mathcal{P}) = \sum_{j=1}^{m} \Psi(B_j,\mathcal{P}), \qquad \Psi(B_j,\mathcal{P}) := \prod_{h<\ell_j} 2^{-\min(D_{j,h},t_h)}. \tag{4}$$

We can therefore bound the potential by bounding one clause's contribution at a time and adding the results. The formula for Ψ uses only the rounds in which the clause's variables are assigned. The signs of its literals do not affect this formula.

The assumption that F accepts at least an ε fraction of inputs tells us how small we need to make each clause's average contribution. A clause of width w is false on a uniform input with probability 2^{-w}. Let w_j^{old} be the width of the jth clause before shortening. No variable appears in two different clauses, so the events that the clauses are satisfied are independent. It follows that

$$\Pr[F(\mathcal{U}_n)=1] = \prod_j (1-2^{-w_j^{\mathrm{old}}}) \ge \varepsilon.$$

Taking natural logarithms and using $-\ln(1-u) \ge u$ gives

$$\sum_j 2^{-w_j^{\mathrm{old}}} \le -\ln\Pr[F(\mathcal{U}_n)=1] \le \ln(1/\varepsilon).$$

Shortening a clause increases its chance of being false by at most 2^{-W}. There are at most n nonconstant clauses, and $2^W \geq 8n$. Thus, after shortening, we still have

$$Z_{\mathrm{OR}}(F^{\leq W}) := \sum_j 2^{-w_j} \leq \ln(1/\varepsilon) + n2^{-W} \leq \ln(1/\varepsilon) + \tfrac{1}{8}. \tag{5}$$

So the sum of the clauses' failure probabilities is small. We want each clause's average contribution to the potential to be at most a constant times its failure probability. In symbols, it is enough to prove

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Psi(B,\mathcal{P})] \leq C2^{-|B|} \qquad (1 \leq |B| \leq W), \tag{6}$$

with the same constant C for every such set B. We could then add the bounds for the clauses. Since the average of a sum is the sum of the averages, equation (4) would give

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\leq W},\mathcal{P})] \leq C\sum_j 2^{-w_j} \leq C\left(\ln(1/\varepsilon) + \tfrac{1}{8}\right) = O(\Lambda).$$

This is the bound we wanted in (3). Notice why the factor 2^{-w} matters: a formula can have many wide clauses and still have many satisfying inputs. A constant bound for each clause would grow with the number of clauses. The factors 2^{-w} let us use the smaller sum in (5).

We now focus on one clause C. To shorten the notation, write B for its variable set, w for its width, D_h for the number of its variables assigned in round h, and ℓ for its completion round. Put $\rho := D_\ell$. Thus ρ variables are assigned in the completion round, and the other $w-\rho$ variables are assigned earlier. Recall that in Theorem 3.1, the clause contributes one to K_ℓ if it is still unsatisfied when its completion round begins.

First consider the untruncated contribution and give the earlier variables independent uniform values. The clause is still unsatisfied exactly when all $w-\rho$ earlier literals are false. Each is false with probability $1/2$, independently, so

$$\Psi_{\mathrm{untr}}(B,\mathcal{P}) = \prod_{h<\ell} 2^{-D_h} = 2^{-(w-\rho)} = 2^{-w}2^{\rho}.$$

The factor 2^{-w} depends only on the number of variables in the clause, not on the chosen partition $\mathcal{P}$. Since we fix the clause before choosing $\mathcal{P}$, we treat 2^{-w} as a constant when taking the expectation over $\mathcal{P}$. We can therefore take it outside the expectation. Before truncation, the bound we need in (6) is equivalent to

$$\mathbb{E}_{\mathcal{P}}[\Psi_{\mathrm{untr}}(B,\mathcal{P})] \leq C2^{-w} \iff 2^{-w}\mathbb{E}_{\mathcal{P}}[2^{\rho(B,\mathcal{P})}] \leq C2^{-w} \iff \mathbb{E}_{\mathcal{P}}[2^{\rho(B,\mathcal{P})}] \leq C.$$

Our first goal is therefore to bound the average of 2^ρ by a constant.

We next account for truncation. The potential uses at most t_h variables of the clause in round h, as required by the round-source argument in Section 3.2. It therefore counts only $\min(D_h, t_h)$ of the D_h variables assigned in that round. The number left out is $(D_h - t_h)^+$, where $x^+ := \max\{x, 0\}$. Each variable left out removes a factor $1/2$ from the product. Hence

$$\Psi(B,\mathcal{P}) = 2^{-w}R(B,\mathcal{P}), \qquad R(B,\mathcal{P}) := 2^{\rho}\prod_{h<\ell} 2^{(D_h-t_h)^+}. \tag{7}$$

Because $(D_h - t_h)^+ \le D_h$ and $\rho + \sum_{h<\ell} D_h = w$, we have $0 \le R(B,\mathcal{P}) \le 2^w$. Equation (7) shows that we need to bound the average of R by a constant. This will give (6) with truncation included.

Our goal is to construct a distribution $\mathfrak{P}$ over partitions, using $O(\log n)$ random bits, such that

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[R(B,\mathcal{P})] \le C \qquad \text{for every fixed } B \subseteq [n], \quad 1 \le |B| \le W.$$

The distribution must depend only on n. We will analyze its effect on one fixed set B at a time.

We build and analyze the distribution in three steps.

1. **Choose independent round numbers and first ignore truncation.**

   We begin with a distribution ν on the round numbers $[J]$. Choose

   $$\tau_1, \ldots, \tau_n \text{ independently from } \nu, \qquad I_k := \{i : \tau_i = k\}.$$

   These choices define a random partition $\mathcal{P} = (I_1, \ldots, I_J)$. Write $\mathbb{E}_{\text{ind}}$ for expectation under these independent choices.

   We will find a distribution ν, with $J = O(\log W)$, for which

   $$\mathbb{E}_{\text{ind}}[2^{\rho(B,\mathcal{P})}] \le C_0$$

   for an absolute constant C_0. This bounds the clause's contribution before accounting for truncation. Independence lets us calculate the contribution of each possible completion round and then bound their sum.

2. **Choose the truncation thresholds and bound their effect.**

   The actual potential counts at most t_h variables in each earlier round h. We therefore need to control the additional factor

   $$\prod_{h<\ell} 2^{(D_h - t_h)^+}$$

   appearing in R.

   We will choose fixed thresholds with $t_J = 0$ and $\sum_h t_h = O(W)$, and prove

   $$\mathbb{E}_{\text{ind}}[R(B,\mathcal{P})] \le 2\,\mathbb{E}_{\text{ind}}[2^{\rho(B,\mathcal{P})}] \le 2C_0.$$

   Thus we obtain a constant bound for the potential's full contribution, including truncation, while keeping the total threshold budget small.

3. **Replace the independent choices with a short seed.**

   Our chosen distribution ν will have a sampling function

   $$f : \{0,1\}^b \longrightarrow [J], \qquad b = O(\log W),$$

such that $f(\mathcal{U}_b)$ has distribution ν. Applying f to an independent block of b uniform bits for each variable would use nb independent uniform bits.

Instead, we generate these nb bits using a small-bias source with an $O(\log n)$-bit seed. Divide its output into n consecutive blocks $z^{(1)}, \ldots, z^{(n)} \in \{0,1\}^b$, one for each variable, and set

$$\tau_i := f(z^{(i)}), \qquad I_k := \{i : f(z^{(i)}) = k\}.$$

Every variable receives exactly one round number, so these sets form an ordered partition $\mathcal{P} = (I_1, \ldots, I_J)$. Choosing the seed uniformly makes this partition random; we call its distribution $\mathfrak{P}$. We will prove that

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[R(B,\mathcal{P})] \le \mathbb{E}_{\mathrm{ind}}[R(B,\mathcal{P})] + O(1) = O(1).$$

To prove this comparison, we consider the possible lists of round numbers for the variables in B. We keep polynomially many lists that account for most of the probability under independent choices. We then define an upper bound for R: it equals R on the lists we kept and equals $2^{|B|}$ on all other lists. This changes the independent expectation by only a small amount. Small bias lets us compare the expectation of this upper bound under the two sources, which gives the required bound for R.

The same distribution and thresholds work for every fixed B. Applying the resulting bound to each clause and adding gives

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\le W}, \mathcal{P})] \le C \sum_j 2^{-|B_j|} \le C\left(\ln(1/\varepsilon) + \tfrac{1}{8}\right).$$

Thus the bound for one set gives the required bound for the whole formula.

Lemma 3.4 (Partition bound for one set)**.** *There are $J = \Theta(\log W)$ rounds, integer thresholds $t_h \ge 0$ with $t_J = 0$ and $\sum_h t_h = O(W)$, and an explicit distribution $\mathfrak{P}$ over ordered partitions of $[n]$, which can be sampled using $O(\log n)$ bits, such that every nonempty $B \subseteq [n]$ with $|B| \le W$ satisfies*

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[R(B,\mathcal{P})] \le C, \qquad \textit{and hence} \qquad \mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Psi(B,\mathcal{P})] \le C 2^{-|B|},$$

for an absolute constant C. All parameters and the distribution are fixed using only n.

To prove Lemma 3.4, we first state and prove the untruncated bound, then choose a distribution that satisfies its assumptions. We then introduce truncation and use a short seed, stating and proving a lemma for each step.

How to choose a round for each variable. For now, choose the round number of each variable independently from the same distribution ν on $[J]$. Define

$$p_k := \Pr_{\tau\sim\nu}[\tau = k], \qquad \theta_k := \Pr_{\tau\sim\nu}[\tau > k], \qquad A_k := 1 + p_k - \theta_k.$$

Here p_k is the chance that a variable is assigned in round k, and θ_k is the chance that it is assigned later. The expression A_k will appear in the calculation in Step 1. We have

$\theta_{k-1} = p_k + \theta_k$, because a variable that remains after round $k-1$ is assigned either in round k or later. Also, $\theta_0 = 1$ and $\theta_J = 0$. We first prove the bound we need under a few conditions on ν. After proving the lemma, we will show how to choose a distribution that satisfies all of these conditions.

Lemma 3.5 (Step 1: first ignore truncation)**.** *Choose the round number of each variable independently from a distribution ν on $[J]$. Suppose that its probabilities satisfy*

$$\begin{aligned} p_k \le \theta_k/2 \quad (k < J), \qquad \theta_{k+1} \le \tfrac{3}{4}\theta_k \quad (1 \le k < J), \\ A_k \ge \tfrac{1}{2} \quad (k \in [J]), \qquad p_J \le 1/W, \qquad J = O(\log W). \end{aligned} \tag{8}$$

Then, for every nonempty set $B \subseteq [n]$ with $|B| \le W$, the number $\rho(B, \mathcal{P})$ of its variables assigned in its completion round satisfies

$$\mathbb{E}_{\text{ind}}[2^{\rho(B,\mathcal{P})}] \le 4 + e.$$

Here $\mathbb{E}_{\text{ind}}$ is the expectation over the independent choices of round numbers.

Proof. The first condition bounds the chance of assigning a variable in a nonfinal round. Indeed, given that a variable is still unassigned before round k, its chance of being assigned in that round is $p_k/(p_k + \theta_k) \le 1/3$, because $\theta_k \ge 2p_k$. The second condition says that the chance of remaining unassigned drops by at least a quarter in each round. These two conditions will bound the contributions from nonfinal rounds. The condition $p_J \le 1/W$ will handle the final round. We will use $A_k \ge 1/2$ and $J = O(\log W)$ later, when we introduce truncation.

Fix B of size $w \le W$. We will calculate $\mathbb{E}_{\text{ind}}[2^\rho]$ by separating the possible completion rounds. For each round k, put

$$M_k := \mathbb{E}_{\text{ind}}[2^\rho \mathbf{1}[\ell = k]].$$

Since the clause completes in exactly one round,

$$\mathbb{E}_{\text{ind}}[2^\rho] = \sum_{k=1}^{J} M_k.$$

Here M_k is the contribution to the average from partitions that complete the clause in round k.

To calculate M_k, fix a round k. For each variable $i \in B$, define

$$Y_i := \begin{cases} 2, & \tau_i = k, \\ 1, & \tau_i < k, \\ 0, & \tau_i > k. \end{cases}$$

If any variable is assigned after round k, the product $\prod_{i \in B} Y_i$ is zero. Otherwise, each of the D_k variables assigned in round k contributes a factor of two, and every earlier variable contributes a factor of one. Thus

$$\prod_{i \in B} Y_i = 2^{D_k} \mathbf{1}[\ell \le k].$$

We want to count only partitions with $\ell = k$. The product also counts partitions with $\ell < k$, but on those partitions $D_k = 0$, so its value is exactly one. Subtracting their indicator therefore gives

$$2^\rho \mathbf{1}[\ell = k] = \prod_{i \in B} Y_i - \mathbf{1}[\ell < k].$$

The round choices are independent, so the random variables Y_i are independent. Moreover,

$$\mathbb{E}_{\mathrm{ind}}[Y_i] = 2p_k + \Pr[\tau_i < k] = 1 + p_k - \theta_k = A_k.$$

Also, $\ell < k$ exactly when all w variables are assigned before round k, which has probability $(1 - p_k - \theta_k)^w$. Taking expectations in the preceding identity yields

$$\begin{aligned} M_k &= \mathbb{E}_{\mathrm{ind}}\left[\prod_{i \in B} Y_i\right] - \Pr_{\mathrm{ind}}[\ell < k] \\ &= \prod_{i \in B} \mathbb{E}_{\mathrm{ind}}[Y_i] - (1 - p_k - \theta_k)^w \\ &= A_k^w - (1 - p_k - \theta_k)^w. \end{aligned} \tag{9}$$

For each nonfinal round k, we will show that

$$M_k \le \Pr_{\mathrm{ind}}[k \le \ell \le k + 3].$$

In words, round k's contribution is at most the probability that the clause completes in one of four consecutive rounds. Since θ_m decreases by a factor of at most $3/4$ per round,

$$\theta_{k+3} \le (3/4)^3 \theta_k \le \theta_k / 2,$$

where we set $\theta_m = 0$ for $m \ge J$. Together with $p_k \le \theta_k/2$, this gives

$$A_k = 1 + p_k - \theta_k \le 1 - \theta_k/2 \le 1 - \theta_{k+3}.$$

Using $1 - p_k - \theta_k = 1 - \theta_{k-1}$ in (9), we therefore obtain

$$\begin{aligned} M_k &= A_k^w - (1 - \theta_{k-1})^w \\ &\le (1 - \theta_{k+3})^w - (1 - \theta_{k-1})^w \\ &= \Pr_{\mathrm{ind}}[k \le \ell \le k + 3]. \end{aligned}$$

The last equality holds because $(1 - \theta_m)^w$ is the probability that all clause variables have been assigned by round m. Subtracting the probability of completion before round k leaves exactly the probability of completion between rounds k and $k + 3$.

We now sum over the nonfinal rounds. For any fixed partition $\mathcal{P}$, let $\ell = \ell(B, \mathcal{P})$. The condition $k \le \ell \le k + 3$ is equivalent to $\ell - 3 \le k \le \ell$, so it holds for at most four integer values of k. Therefore, for every $\mathcal{P}$,

$$\sum_{k=1}^{J-1} \mathbf{1}[k \le \ell(B, \mathcal{P}) \le k + 3] \le 4.$$

Taking expectations and using linearity of expectation gives

$$\begin{aligned}\sum_{k=1}^{J-1} M_k &\le \sum_{k=1}^{J-1} \Pr_{\text{ind}}[k \le \ell(B,\mathcal{P}) \le k+3] \\ &= \mathbb{E}_{\text{ind}}\left[\sum_{k=1}^{J-1} \mathbf{1}[k \le \ell(B,\mathcal{P}) \le k+3]\right] \\ &\le 4.\end{aligned}$$

For the final round, equation (9) instead gives

$$M_J = (1+p_J)^w - (1-p_J)^w \le (1+p_J)^w \le e^{wp_J} \le e.$$

The first inequality drops the subtracted nonnegative term. The next uses $1+x \le e^x$, and the last uses $w \le W$ and $p_J \le 1/W$. Adding the bounds for the nonfinal and final rounds gives

$$\mathbb{E}_{\text{ind}}[2^\rho] = \sum_{k=1}^{J} M_k \le 4+e. \tag{10}$$

□

Choosing the round probabilities. We now construct a distribution satisfying the conditions of Lemma 3.5. The main requirement is that each nonfinal round assigns a variable with conditional probability at most $1/3$, given that it has not yet been assigned. At the same time, the probability of remaining unassigned must decrease quickly, so that only $O(\log W)$ rounds are needed.

A natural starting point is to assign half of the remaining probability in each round. However, this would give conditional assignment probability $1/2$, exceeding the required $1/3$. We therefore *split each halving step into two rounds.*

To see how this works, suppose a variable remains unassigned with probability Q before a pair of rounds. Assign probability $Q/4$ to the first round and another $Q/4$ to the second round. The remaining probability after the pair is $Q/2$. Within the pair:

- In the first round, the conditional assignment probability is $(Q/4)/Q = 1/4$.
- In the second round, it is $(Q/4)/(3Q/4) = 1/3$.

Thus neither round assigns too much probability, while every pair of rounds reduces the remaining probability by half. Starting from $Q = 1$, the round probabilities are

$$\frac{1}{4}, \frac{1}{4}, \quad \frac{1}{8}, \frac{1}{8}, \quad \frac{1}{16}, \frac{1}{16}, \quad \ldots$$

Once the remaining probability is at most $1/W$, we assign all of it to the final round. We call this the *split-dyadic distribution.*

Formally, set

$$r_\star := \lceil \log_2 W \rceil, \qquad J := 2r_\star + 1,$$

and define the distribution ν by

$$p_{2r-1} = p_{2r} := 2^{-(r+1)} \quad (1 \le r \le r_\star), \qquad p_J := 2^{-r_\star}.$$

These probabilities sum to one: each pair uses half of the probability remaining before it, and the final round receives everything left over.

We next verify the conditions of Lemma 3.5. Before pair r, the remaining probability is $2^{-(r-1)}$. After its first and second rounds, respectively, it is

$$\theta_{2r-1} = 3 \cdot 2^{-(r+1)}, \qquad \theta_{2r} = 2^{-r}.$$

Therefore,

$$\frac{p_{2r-1}}{\theta_{2r-1}} = \frac{1}{3}, \qquad \frac{p_{2r}}{\theta_{2r}} = \frac{1}{2},$$

so $p_k \le \theta_k/2$ in every nonfinal round.

The remaining probability decreases by a factor $2/3$ from the first round of a pair to the second, and by a factor $3/4$ from there to the first round of the next pair. At the final round it becomes zero. Hence $\theta_{k+1} \le (3/4)\theta_k$. Also,

$$A_{2r-1} = 1 - 2^{-r} \ge \frac{1}{2}, \qquad A_{2r} = 1 - 2^{-(r+1)} \ge \frac{3}{4}, \qquad A_J = 1 + p_J \ge 1.$$

Finally,

$$p_J = 2^{-r_\star} \le \frac{1}{W}, \qquad J = \Theta(\log W).$$

Thus all the required conditions hold.

The split-dyadic distribution can be sampled using $r_\star + 1$ independent uniform bits. Define the sampling function

$$f : \{0,1\}^{r_\star+1} \longrightarrow [J]$$

as follows. Write its input as $(u_1, \ldots, u_{r_\star}, v)$. If the first zero among the u-bits occurs at position r, output $2r-1$ when $v = 0$, and $2r$ when $v = 1$. If all the u-bits are one, output J.

When the input bits are independent and uniform, the first zero occurs at r with probability 2^{-r}, and v divides this probability equally between the two rounds. All the u-bits are one with probability $2^{-r_\star}$, giving the stated probability for the final round. Thus $f(\mathcal{U}_{r_\star+1})$ has distribution ν. Applying f to an independent block of uniform bits for each variable gives independent round numbers with this distribution. We will use the same sampling function in Step 3.

The first two properties below will be used when we replace the independent bits by a small-bias source. The third will be used for parity constraints in Section 4.

Lemma 3.6 (Properties of the split-dyadic construction)**.** *Apply the sampling function f to each of n blocks of $r_\star + 1$ input bits, and put $N := n(r_\star + 1)$. For a fixed set $B \subseteq [n]$, let $\sigma(z) := (\tau_i(z))_{i \in B}$ be its list of round numbers on input $z \in \{0,1\}^N$.*

(a) For every list $\sigma \in [J]^B$, the event $\{\sigma(z) = \sigma\}$ fixes some specified input bits and leaves the others free. The absolute values of its indicator's Fourier coefficients sum to one. In particular, for any δ-biased distribution M on the input bits,

$$\left| \Pr_M[\sigma(z) = \sigma] - \Pr_{\text{ind}}[\sigma(z) = \sigma] \right| \le \delta.$$

(b) The round probabilities satisfy

$$\sum_{k=1}^{J} \sqrt{p_k} < 2 + \sqrt{2} < 4.$$

(c) For every variable i and nonfinal round k, the event $\{\tau_i > k\}$ fixes some input bits, or is a disjoint union of two events of this form. Consequently,

$$\sum_{S \subseteq [N]} \left| \widehat{\mathbf{1}[\tau_i > k]}(S) \right| \leq 2.$$

Here the subscript ind *means that all input bits are independent and uniform.*

Proof. For part (a), fix a set $B \subseteq [n]$ and a list of round numbers $\sigma = (\sigma_i)_{i \in B}$. We first describe exactly which input bits the event

$$E_\sigma := \{z : \sigma(z) = \sigma\}$$

fixes. We will then expand its indicator as a sum of parity functions.

For each variable i, write its block of input bits as

$$(u_{i,1}, \ldots, u_{i,r_\star}, v_i).$$

By the definition of the sampling function:

- If $\sigma_i = 2r - 1$ or $\sigma_i = 2r$, then $\tau_i = \sigma_i$ holds exactly when

  $$u_{i,1} = \cdots = u_{i,r-1} = 1, \qquad u_{i,r} = 0,$$

  and $v_i = 0$ or $v_i = 1$, respectively. The bits $u_{i,r+1}, \ldots, u_{i,r_\star}$ are unrestricted. When $r = 1$, there are no bits before $u_{i,r}$ to fix.

- If $\sigma_i = J$, then $\tau_i = \sigma_i$ holds exactly when

  $$u_{i,1} = \cdots = u_{i,r_\star} = 1,$$

  with v_i unrestricted.

Different variables use disjoint blocks of input bits. Therefore, all the requirements for the list σ can be combined into one set of fixed bit values. More precisely, there exist a set of input positions $S_{\text{fix}} \subseteq [N]$ and values $a_j \in \{0, 1\}$, for $j \in S_{\text{fix}}$, such that

$$E_\sigma = \{z \in \{0,1\}^N : z_j = a_j \text{ for every } j \in S_{\text{fix}}\}.$$

All positions outside S_{fix} are unrestricted.

Put $d := |S_{\text{fix}}|$, and define the parity functions

$$\chi_S(z) := (-1)^{\sum_{j \in S} z_j}, \qquad S \subseteq [N].$$

For one fixed bit, we have

$$\mathbf{1}[z_j = a_j] = \frac{1 + (-1)^{a_j} \chi_{\{j\}}(z)}{2}.$$

Indeed, the right-hand side is one when $z_j = a_j$ and zero otherwise. Multiplying these identities over the fixed positions gives

$$\begin{aligned}
\mathbf{1}[E_\sigma](z) &= 2^{-d} \prod_{j \in S_{\mathrm{fix}}} \bigl(1 + (-1)^{a_j} \chi_{\{j\}}(z)\bigr) \\
&= 2^{-d} \sum_{S \subseteq S_{\mathrm{fix}}} (-1)^{\sum_{j \in S} a_j} \chi_S(z).
\end{aligned}$$

In the expansion, S records the positions where we choose the parity term rather than the constant term. Since $\prod_{j \in S} \chi_{\{j\}} = \chi_S$, this is the Fourier expansion of the indicator.

There are 2^d coefficients in this expansion, each of absolute value 2^{-d}; every other Fourier coefficient is zero. Hence

$$\sum_{S \subseteq [N]} \left| \widehat{\mathbf{1}[E_\sigma]}(S) \right| = 2^d \cdot 2^{-d} = 1.$$

Finally, suppose M is a δ-biased distribution. For every nonempty S,

$$|\mathbb{E}_M[\chi_S]| \le \delta, \qquad \mathbb{E}_{\mathrm{ind}}[\chi_S] = 0.$$

The empty-set term is the constant function $\chi_\varnothing = 1$, so its expectation is identical under both distributions. Taking expectations in the Fourier expansion and subtracting therefore gives

$$\Pr_M[E_\sigma] - \Pr_{\mathrm{ind}}[E_\sigma] = 2^{-d} \sum_{\varnothing \ne S \subseteq S_{\mathrm{fix}}} (-1)^{\sum_{j \in S} a_j} \mathbb{E}_M[\chi_S].$$

By the triangle inequality and the small-bias assumption,

$$\begin{aligned}
\left| \Pr_M[E_\sigma] - \Pr_{\mathrm{ind}}[E_\sigma] \right| &\le 2^{-d} \sum_{\varnothing \ne S \subseteq S_{\mathrm{fix}}} |\mathbb{E}_M[\chi_S]| \\
&\le 2^{-d}(2^d - 1)\delta \\
&\le \delta.
\end{aligned}$$

This proves both claims in part (a).

For part (b), the split-dyadic probabilities give

$$\sum_{k=1}^{J} \sqrt{p_k} = \sum_{r=1}^{r_\star} 2^{(1-r)/2} + 2^{-r_\star/2} < \sum_{r=1}^{\infty} 2^{(1-r)/2} = 2 + \sqrt{2} < 4.$$

The strict inequality holds because the part of the infinite sum after $r_\star$ starts with $2^{-r_\star/2}$ and has further positive terms.

For part (c), the sampling function gives, for $1 \le r \le r_\star$,

$$\begin{aligned}
\{\tau_i > 2r\} &= \{u_{i,1} = \cdots = u_{i,r} = 1\}, \\
\{\tau_i > 2r - 1\} &= \{u_{i,1} = \cdots = u_{i,r} = 1\} \\
&\quad \cup \{u_{i,1} = \cdots = u_{i,r-1} = 1,\ u_{i,r} = 0,\ v_i = 1\}.
\end{aligned}$$

The two events in the union are disjoint because they give different values to $u_{i,r}$. Each event fixes some input bits, so its indicator's Fourier coefficients have absolute values summing to one, by the calculation in part (a). The triangle inequality bounds the sum for the union by two. $\square$

Choosing the truncation thresholds. Set

$$t_h := \lceil 4Wp_h + 3\log_2(J+1)\rceil \quad (h < J), \qquad t_J := 0. \tag{11}$$

For a set of at most W variables, the average number assigned in round h is at most Wp_h. We choose the truncation threshold to be four times this upper bound, plus a term that depends on the number of rounds. Step 2 will show why this choice keeps the total increase caused by truncation below a factor of two. Since $\sum_h p_h = 1$,

$$\sum_h t_h \le 4W + (J-1)\big(3\log_2(J+1) + 1\big) = O(W + J\log(J+1)) = O(W),$$

where the last equality uses $J = O(\log W)$.

Lemma 3.7 (Step 2: introduce truncation)**.** *Let ν be a distribution on the round numbers $[J]$, and define*

$$p_k := \Pr_{\tau\sim\nu}[\tau = k], \qquad \theta_k := \Pr_{\tau\sim\nu}[\tau > k], \qquad A_k := 1 + p_k - \theta_k.$$

Suppose that

$$\begin{gathered} p_k \le \theta_k/2 \quad (k < J), \qquad \theta_{k+1} \le \tfrac{3}{4}\theta_k \quad (1 \le k < J), \\ A_k \ge \tfrac{1}{2} \quad (k \in [J]), \qquad p_J \le 1/W, \qquad J = O(\log W). \end{gathered}$$

Set

$$t_h := \lceil 4Wp_h + 3\log_2(J+1)\rceil \quad (h < J), \qquad t_J := 0.$$

Choose $\tau_1, \dots, \tau_n$ independently from ν, and let $\mathcal{P} = (I_1, \dots, I_J)$ be the partition defined by $I_h := \{i : \tau_i = h\}$.

For every nonempty set $B \subseteq [n]$ of size $w \le W$, define

$$D_h := |B \cap I_h|, \qquad \ell := \max_{i\in B} \tau_i, \qquad \rho(B, \mathcal{P}) := D_\ell,$$

Thus ℓ is the last round assigning a variable in B, and $\rho(B, \mathcal{P})$ counts the variables of B assigned in that round. Put

$$R(B, \mathcal{P}) := 2^{\rho(B,\mathcal{P})} \prod_{h<\ell} 2^{(D_h - t_h)^+}, \qquad x^+ := \max\{x, 0\}.$$

Then

$$\mathbb{E}_{\text{ind}}[R(B, \mathcal{P})] \le 2\mathbb{E}_{\text{ind}}[2^{\rho(B,\mathcal{P})}] \le 2(4+e),$$

where $\mathbb{E}_{\text{ind}}$ is the expectation over the independent round choices.

Proof. Fix a nonempty set $B \subseteq [n]$ of size $w \leq W$. We will condition on both the completion round and the variables assigned in that round. This makes 2^ρ constant, leaving only the factor introduced by truncation to bound.

For a round k and a nonempty set $C \subseteq B$, let

$$E_{k,C} := \{\tau_i = k \text{ for } i \in C, \quad \tau_i < k \text{ for } i \in B \setminus C\}.$$

Thus $E_{k,C}$ is the event that B completes in round k, with exactly the variables in C assigned in that round. On this event,

$$\rho = |C|, \qquad R(B, \mathcal{P}) = 2^{|C|} X_k, \qquad X_k := \prod_{h<k} 2^{(D_h - t_h)^+}.$$

We will show that, for every such event of positive probability,

$$\mathbb{E}_{\text{ind}}[X_k \mid E_{k,C}] \leq 2.$$

First consider the two earliest completion rounds. If $k = 1$, the product defining X_k is empty, so $X_k = 1$. If $k = 2$, the only earlier round is round 1. Since $\theta_1 = 1 - p_1$, the assumption $A_1 \geq \frac{1}{2}$ gives

$$A_1 = 1 + p_1 - \theta_1 = 2p_1 \geq \frac{1}{2}, \qquad \text{hence} \qquad p_1 \geq \frac{1}{4}.$$

Therefore,

$$t_1 \geq 4Wp_1 \geq W \geq D_1,$$

so $X_k = 1$ in this case as well.

Now suppose $k \geq 3$, and put $m := |B \setminus C|$. Conditioned on $E_{k,C}$, the labels of these m earlier variables remain independent: the conditioning imposes a separate requirement $\tau_i < k$ on each of them. Their common distribution is $\nu(\cdot \mid \tau < k)$, with probabilities

$$p'_h := \Pr[\tau = h \mid \tau < k] = \frac{p_h}{1 - \theta_{k-1}}, \qquad h < k.$$

We first bound how much this conditioning can increase each round probability. From $p_1 \geq \frac{1}{4}$ and the assumed decrease of the tail probabilities,

$$\theta_1 \leq \frac{3}{4}, \qquad \theta_2 \leq \frac{3}{4}\theta_1 \leq \frac{9}{16}.$$

The tail probabilities are nonincreasing, so $k \geq 3$ implies

$$1 - \theta_{k-1} \geq 1 - \theta_2 \geq \frac{7}{16}.$$

Consequently,

$$p'_h \leq \frac{16}{7} p_h \qquad (h < k).$$

We now bound the expected truncation factor under these conditional choices. For every $D, t \geq 0$,

$$2^{(D-t)^+} = \max\{1, 2^{D-t}\} \leq 1 + 2^{D-t}.$$

Applying this inequality in each earlier round and expanding the product gives

$$X_k \le \sum_{S\subseteq\{1,\dots,k-1\}} 2^{-\sum_{h\in S} t_h}\, 2^{\sum_{h\in S} D_h}.$$

Here S records the rounds where we choose the term $2^{D_h - t_h}$ in the expansion.

For a fixed S, the sum $\sum_{h\in S} D_h$ counts the earlier variables whose labels belong to S. Each such variable contributes a factor of two, and every other earlier variable contributes one. By their conditional independence,

$$\begin{aligned}\mathbb{E}_{\text{ind}}\left[2^{\sum_{h\in S} D_h} \mid E_{k,C}\right] &= \left(1 + \sum_{h\in S} p'_h\right)^m \\ &\le \exp\left(m \sum_{h\in S} p'_h\right) \\ &\le \exp\left(\frac{16}{7} W \sum_{h\in S} p_h\right).\end{aligned}$$

The first inequality uses $1 + x \le e^x$; the second uses $m \le W$ and $p'_h \le \frac{16}{7} p_h$.

Substituting this bound into the expansion, we obtain

$$\begin{aligned}\mathbb{E}_{\text{ind}}[X_k \mid E_{k,C}] &\le \sum_{S\subseteq\{1,\dots,k-1\}} \prod_{h\in S} \left(2^{-t_h} e^{(16/7)W p_h}\right) \\ &= \prod_{h<k} \left(1 + 2^{-t_h} e^{(16/7)W p_h}\right).\end{aligned}$$

The last equality follows by expanding the product: each round is either included in S or omitted.

The thresholds

$$t_h = \lceil 4W p_h + 3\log_2(J+1) \rceil \qquad (h < J)$$

make every extra term small. Indeed,

$$2^{-t_h} e^{(16/7)W p_h} \le (J+1)^{-3} e^{-(4\ln 2 - 16/7)W p_h} \le (J+1)^{-3},$$

because $4 \ln 2 > 16/7$. Using $1 + x \le e^x$ once more,

$$\mathbb{E}_{\text{ind}}[X_k \mid E_{k,C}] \le \exp\left(\frac{k-1}{(J+1)^3}\right) \le \exp\left(\frac{J}{(J+1)^3}\right) < 2.$$

This proves the required conditional bound for every completion round.

Finally, the events $E_{k,C}$ partition all choices of labels for B. On each event, $2^\rho = 2^{|C|}$, so

$$\begin{aligned}\mathbb{E}_{\text{ind}}[R(B,\mathcal{P})] &= \sum_{k,C} \Pr[E_{k,C}]\, 2^{|C|}\, \mathbb{E}_{\text{ind}}[X_k \mid E_{k,C}] \\ &\le 2 \sum_{k,C} \Pr[E_{k,C}]\, 2^{|C|} \\ &= 2\mathbb{E}_{\text{ind}}[2^\rho] \le 2(4+e).\end{aligned} \tag{12}$$

The sums run over events of positive probability, and the last inequality is the previously proved untruncated bound. □

How to use a short seed. We have bounded the average of R when all input bits are independent and uniform. We now replace them with a δ-biased source M, using the same sampling function f to choose the rounds. For each fixed set B of size $w \leq W$, we want to show that this replacement increases the average of R by at most a constant.

The value of R depends only on the list of rounds assigned to the variables in B. Small bias changes the probability of any one list by at most δ, as proved for the split-dyadic construction. But there are J^w possible lists, so adding an error bound for every list would introduce a factor of J^w. We will use a smaller collection instead. The proof has three steps.

1. **Choose lists that account for most of the probability.** Keep each list whose probability under independent choices is at least 2^{-8w}. There are at most 2^{8w} such lists. We will show that all other lists together have probability at most 2^{-2w}.

2. **Replace R by an upper bound.** Define $\widetilde{R}$ to equal R on the lists we kept and 2^w on all other lists. Since $R \leq 2^w$, this gives $R \leq \widetilde{R}$ for every input. Under independent choices, the replacement increases the expectation by at most $2^w \cdot 2^{-2w} = 2^{-w}$.

3. **Compare the average of this upper bound under the two sources.** Outside the chosen lists, $\widetilde{R}$ has the constant value 2^w. We can therefore write it using only this constant and the indicators of the lists we kept. There are at most 2^{8w} indicators, each multiplied by a number of absolute value at most 2^w. Small bias changes each indicator's expectation by at most δ, so it changes the expectation of $\widetilde{R}$ by at most $\delta 2^{9w}$.

These bounds fit together as follows:

$$\mathbb{E}_M R \leq \mathbb{E}_M \widetilde{R} \leq \mathbb{E}_{\mathrm{ind}} \widetilde{R} + \delta 2^{9w} \leq \mathbb{E}_{\mathrm{ind}} R + 2^{-w} + \delta 2^{9w}.$$

Since $w \leq W = O(\log n)$, we can make the last error term small with $\delta = (16n)^{-12}$. An explicit small-bias source with this choice of δ needs only $O(\log n)$ seed bits.

Lemma 3.8 (Step 3: choose all rounds with a short seed)**.** *Let $N := n(r_\star + 1)$, and let M be a δ-biased distribution on N bits. Apply the sampling function f to each block of $r_\star + 1$ bits to obtain the round numbers and hence a partition $\mathcal{P}$. Use the truncation thresholds*

$$t_h := \lceil 4Wp_h + 3\log_2(J+1) \rceil \quad (h < J), \qquad t_J := 0.$$

For every nonempty set $B \subseteq [n]$ of size $w \leq W$,

$$\mathbb{E}_M[R(B, \mathcal{P})] \leq \mathbb{E}_{\mathrm{ind}}[R(B, \mathcal{P})] + 2^{-w} + \delta 2^{9w}.$$

In particular, for $\delta = (16n)^{-12}$, this expectation is at most 16, and an explicit such distribution M can be sampled using $O(\log n)$ bits.

Proof. Fix B of size $w \leq W$. For each list $\sigma \in [J]^B$, write R_σ for the value of R on that list. Thus $0 \leq R_\sigma \leq 2^w$. Under independent bits, the list has probability $p_\sigma = \prod_{i \in B} p_{\sigma_i}$. Let

$$\Sigma^* := \{\sigma : p_\sigma \geq 2^{-8w}\}.$$

Since the list probabilities sum to one, $|\Sigma^*| \le 2^{8w}$.

We first bound the probability of all other lists. If $\sigma \notin \Sigma^*$, then $\sqrt{p_\sigma} < 2^{-4w}$, and hence $p_\sigma \le 2^{-4w}\sqrt{p_\sigma}$. Therefore,

$$\begin{aligned}\Pr_{\text{ind}}[\sigma \notin \Sigma^*] &\le 2^{-4w} \sum_{\sigma \in [J]^B} \sqrt{p_\sigma} \\ &= 2^{-4w} \left(\sum_{k=1}^{J} \sqrt{p_k} \right)^w \\ &\le 2^{-2w}.\end{aligned} \tag{13}$$

The equality uses independence to write p_σ as a product, then sums over each variable's round choice. The last inequality uses $\sum_k \sqrt{p_k} < 4$, proved in Lemma 3.6(b).

For input bits z, let $\sigma(z)$ be the list they produce, and define

$$\widetilde{R}(z) := \begin{cases} R_{\sigma(z)}, & \sigma(z) \in \Sigma^*, \\ 2^w, & \sigma(z) \notin \Sigma^*. \end{cases}$$

This function equals R on the lists we kept and bounds it from above everywhere else. The probability bound just proved gives

$$0 \le \mathbb{E}_{\text{ind}}\widetilde{R} - \mathbb{E}_{\text{ind}}R \le 2^w \Pr_{\text{ind}}[\sigma \notin \Sigma^*] \le 2^{-w}.$$

Indeed, the two functions agree on Σ^*, and their difference on any other list is at most 2^w.

To compare expectations under the two sources, write

$$\widetilde{R}(z) = 2^w - \sum_{\sigma \in \Sigma^*} (2^w - R_\sigma)\mathbf{1}[\sigma(z) = \sigma].$$

If the list is in Σ^*, exactly one indicator is one and the expression equals R_σ. Otherwise all indicators are zero and the expression equals 2^w. The constant has the same expectation under both sources. Each indicator's expectation changes by at most δ, by Lemma 3.6(a). Taking the difference of the two expectations and using the triangle inequality therefore gives

$$\begin{aligned}\left|\mathbb{E}_M\widetilde{R} - \mathbb{E}_{\text{ind}}\widetilde{R}\right| &\le \sum_{\sigma \in \Sigma^*} (2^w - R_\sigma) \left|\Pr_M[\sigma(z) = \sigma] - p_\sigma\right| \\ &\le \delta \sum_{\sigma \in \Sigma^*} (2^w - R_\sigma) \\ &\le \delta 2^w |\Sigma^*| \le \delta 2^{9w}.\end{aligned}$$

The last line uses $0 \le R_\sigma \le 2^w$ and $|\Sigma^*| \le 2^{8w}$. Since $R \le \widetilde{R}$ for every input, we conclude that

$$\begin{aligned}\mathbb{E}_M R &\le \mathbb{E}_M \widetilde{R} \\ &\le \mathbb{E}_{\text{ind}}\widetilde{R} + \delta 2^{9w} \\ &\le \mathbb{E}_{\text{ind}}R + 2^{-w} + \delta 2^{9w}.\end{aligned}$$

Choose $\delta := (16n)^{-12}$. Since $w \le W$ and $2^W < 16n$,

$$\delta 2^{9w} \le (16n)^{-3} < 1.$$

The previously proved bound $\mathbb{E}_{\text{ind}} R \le 2(4+e)$ now gives $\mathbb{E}_M R \le 2(4+e)+2 < 16$. Naor and Naor [NN93] give an explicit δ-biased distribution on N bits with seed length

$$O(\log N + \log(1/\delta)) = O(\log n),$$

because $N = n(r_\star + 1)$ and $r_\star = O(\log\log n)$. □

Proof of Theorem 3.2. Recall that

$$W = \lceil \log_2(8n) \rceil, \qquad \Lambda = 1 + \log(1/\varepsilon).$$

We must construct a distribution $\mathfrak{P}$ over ordered partitions into $J = \Theta(\log\log n)$ rounds, together with thresholds satisfying

$$t_J = 0, \qquad \sum_{h=1}^{J} t_h = O(\log n).$$

The distribution must be sampled using $O(\log n)$ random bits and must satisfy

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\le W}, \mathcal{P})] \le C\Lambda$$

for every read-once CNF F of density at least ε. We will choose all construction parameters using only n.

Set

$$r_\star := \lceil \log_2 W \rceil, \qquad b := r_\star + 1, \qquad J := 2r_\star + 1.$$

Use the split-dyadic distribution ν on round numbers, whose probabilities are

$$p_{2r-1} = p_{2r} := 2^{-(r+1)} \quad (1 \le r \le r_\star), \qquad p_J := 2^{-r_\star}.$$

Choose the truncation thresholds

$$t_h := \lceil 4Wp_h + 3\log_2(J+1) \rceil \quad (h < J), \qquad t_J := 0.$$

Let $f : \{0,1\}^b \to [J]$ be the sampling function defined above. Explicitly, on input $(u_1, \ldots, u_{r_\star}, v)$, if the first zero among the u-bits occurs at position r, then f outputs $2r - 1 + v$. If all the u-bits are one, it outputs J. Thus $f(\mathcal{U}_b)$ has distribution ν.

To choose all round numbers with a short seed, set

$$N := nb, \qquad \delta := (16n)^{-12}.$$

Use the explicit δ-biased source from Lemma 3.8,

$$G_{\text{sb}} : \{0,1\}^s \longrightarrow \{0,1\}^N, \qquad s = O\big(\log N + \log(1/\delta)\big) = O(\log n).$$

For a uniformly chosen seed, write its output as n consecutive blocks

$$z^{(1)}, \dots, z^{(n)} \in \{0,1\}^b,$$

and define

$$\tau_i := f(z^{(i)}), \qquad I_h := \{i \in [n] : \tau_i = h\}.$$

Every variable receives exactly one round number, so $\mathcal{P} = (I_1, \dots, I_J)$ is an ordered partition of $[n]$. Let $\mathfrak{P}$ be the distribution of this partition.

This construction has the required parameter bounds. Indeed,

$$J = \Theta(\log W) = \Theta(\log\log n),$$

and, since $\sum_h p_h = 1$,

$$\begin{aligned}\sum_{h=1}^{J} t_h &\le 4W + (J-1)\big(3\log_2(J+1)+1\big) \\ &= O(W) = O(\log n).\end{aligned}$$

The equality $O(W + J\log(J+1)) = O(W)$ uses $J = O(\log W)$. The partition is sampled using the $s = O(\log n)$ seed bits of G_{sb}.

We next bound the contribution of one clause. Fix a nonempty set $B \subseteq [n]$ of size $w \le W$, and write

$$D_h := |B \cap I_h|, \qquad \ell := \max_{i \in B} \tau_i, \qquad \rho := D_\ell.$$

Recall that

$$R(B,\mathcal{P}) = 2^\rho \prod_{h<\ell} 2^{(D_h - t_h)^+}, \qquad \Psi(B,\mathcal{P}) = 2^{-w} R(B,\mathcal{P}).$$

The split-dyadic probabilities satisfy the assumptions of Lemma 3.5, so independent round choices give

$$\mathbb{E}_{\mathrm{ind}}[2^\rho] \le 4 + e.$$

Lemma 3.7 bounds the additional factor caused by truncation:

$$\mathbb{E}_{\mathrm{ind}}[R(B,\mathcal{P})] \le 2\mathbb{E}_{\mathrm{ind}}[2^\rho] \le 2(4+e).$$

Finally, Lemma 3.8 compares our small-bias construction with these independent choices:

$$\begin{aligned}\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[R(B,\mathcal{P})] &\le \mathbb{E}_{\mathrm{ind}}[R(B,\mathcal{P})] + 2^{-w} + \delta 2^{9w} \\ &\le 2(4+e) + 2^{-w} + (16n)^{-3} \\ &< 16.\end{aligned}$$

Here $\delta 2^{9w} \le (16n)^{-3}$ follows from $w \le W$ and $2^W < 16n$.

Since B is fixed when taking the expectation, 2^{-w} is constant. Therefore,

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Psi(B,\mathcal{P})] = 2^{-w}\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[R(B,\mathcal{P})] \le 16\, 2^{-w}.$$

Together with the construction and parameter bounds above, this also proves Lemma 3.4.

Now let F be any read-once CNF of density at least ε. Shorten every clause wider than W by keeping any W of its literals, obtaining $F^{\le W}$. Removing literals from an OR can only make it harder to satisfy. Hence

$$F^{\le W}(x) = 1 \quad \Longrightarrow \quad F(x) = 1.$$

Write u_j for the original clause widths and $w_j = \min\{u_j, W\}$ for their widths after shortening. Because the clauses use disjoint variable sets,

$$\Pr[F(\mathcal{U}_n) = 1] = \prod_j (1 - 2^{-u_j}) \ge \varepsilon.$$

Using $-\ln(1-x) \ge x$, we obtain

$$\sum_j 2^{-u_j} \le -\ln \Pr[F(\mathcal{U}_n) = 1] \le \ln(1/\varepsilon).$$

Shortening increases each clause's failure probability by at most 2^{-W}. There are at most n nonconstant clauses, and $2^W \ge 8n$, so

$$\sum_j 2^{-w_j} \le \sum_j 2^{-u_j} + n2^{-W} \le \ln(1/\varepsilon) + \frac{1}{8}.$$

Let B_j be the variable set of the j-th shortened clause. The initial potential is the sum of the clause contributions:

$$\Phi_1(F^{\le W}, \mathcal{P}) = \sum_j \Psi(B_j, \mathcal{P}).$$

Applying the bound for each fixed set B_j and using linearity of expectation gives

$$\begin{aligned}
\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\le W}, \mathcal{P})] &= \sum_j \mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Psi(B_j, \mathcal{P})] \\
&\le 16 \sum_j 2^{-w_j} \\
&\le 16 \left(\ln(1/\varepsilon) + \frac{1}{8} \right) \\
&\le C\Lambda
\end{aligned}$$

for an absolute constant C.

At least one partition in $\operatorname{supp}(\mathfrak{P})$ has initial potential no larger than this average. All parameters and the partition distribution depend only on n, and hence are independent of F. This proves Theorem 3.2. $\square$

3.4 Completing the read-once CNF construction

Proof of Theorem 1.3. Apply the generator construction from Section 2.5. Theorem 3.1 supplies its fixed-partition guarantee, and Theorem 3.2 supplies a partition with initial potential $O(1+\log(1/\varepsilon))$. The latter theorem permits shortening the clauses because every satisfying assignment of the shortened formula also satisfies the original formula.

The partition seed has length $O(\log n)$, the number of rounds is $J = O(\log\log n)$, and $\sum_k t_k = O(\log n)$. Choose the profile budget $L = O(1+\log(1/\varepsilon))$ with a sufficiently large absolute constant. The profile encoding uses $O(L+J)$ bits, and the combined round seeds have length

$$O\left(L + \sum_k t_k + J\log(J+1) + J\log\log(n+2)\right) = O(\log n + \log(1/\varepsilon)).$$

Including the partition seed gives the same bound. The supported path associated with the good partition and profile is an output of this fixed generator and satisfies the formula. This proves the hitting guarantee. □

4 Hitting Read-once $\mathrm{CNF}^{\oplus}$ and the width-3 reduction

We now prove Theorem 1.2. We use the same partition distribution and the same nonfinal-round sources as for read-once CNF. To extend the potential to parity constraints, we must account for two differences: whether a constraint can be satisfied before its completion round, and how many variables we fix when it completes.

Fix an ordered partition $\mathcal{P} = (I_1, \ldots, I_J)$. For a constraint on the variable set B_j, let ℓ_j be its completion round, and write

$$D_{j,h} := |B_j \cap I_h|, \qquad \rho_j := D_{j,\ell_j}.$$

Thus ρ_j is the number of its variables assigned in its completion round.

An OR clause can be satisfied by any one of its literals. Consequently, it may disappear before its completion round. If it is still unsatisfied when that round arrives, fixing one remaining literal to its satisfying value is enough to satisfy the clause. This is why an alive OR clause contributes 1 to the current-round parameter. Its earlier contribution to the potential also accounts for the possibility that earlier assignments satisfy it. Under independent uniform assignments, the probability that all its literals in round h are false is $2^{-D_{j,h}}$. As in the RO-CNF proof, we use the truncated factor $2^{-\min(D_{j,h},t_h)}$. The clause's contribution to the initial potential is therefore

$$\Psi(B_j, \mathcal{P}) = \prod_{h<\ell_j} 2^{-\min(D_{j,h},t_h)}.$$

A parity constraint behaves differently. Assigning some of its variables changes the required parity of the remaining variables, but it does not satisfy the constraint while any variable remains unassigned. Thus every parity constraint must be handled in its completion round, and its contribution has no survival factors. If it completes in a nonfinal round, we choose one satisfying assignment to its ρ_j remaining variables and require those variables to take those values. This gives a sufficient event fixing exactly ρ_j bits, to which we can apply the same round-source argument as before. Accordingly, the parity constraint has weight ρ_j

throughout the rounds before completion, and contributes that weight to the current-round parameter when it completes.

The final round allows a smaller weight. There are no future obligations to preserve, so we use a small-bias source to satisfy the remaining parity equations directly. For each remaining OR clause, we again require one literal to be true; this is also an affine equation, on one variable. The sufficient final-round event therefore consists of one affine equation per remaining constraint. As we prove below, small bias handles this event with a parameter depending on the number of equations. We therefore give a parity constraint completing in the final round weight 1.

These choices give the following contributions to the initial potential:

$$\begin{cases} \prod_{h<\ell_j} 2^{-\min(D_{j,h},t_h)}, & \text{for an OR clause,} \\ \rho_j, & \text{for a parity constraint with } \ell_j < J, \\ 1, & \text{for a parity constraint with } \ell_j = J. \end{cases}$$

The OR contribution is exactly the one already analyzed. The new partition estimate must bound the expected parity contribution, including for constraints whose width is larger than W.

The proof has four parts.

1. Define the round parameters and add the future parity weights to the existing potential. The OR contribution is unchanged.

2. Prove a supported-path theorem for every fixed partition. Lemma 3.3 supplies the same conditional survival bound used for RO-CNF; the additional parity weights are fixed once the partition is chosen. The final round requires a separate small-bias argument.

3. Use the partition distribution from Section 3.3. Lemma 3.8 already bounds the OR contributions. For each parity constraint, we construct one function that bounds its contribution from above, bound the expectation under independent bits by a fourth moment, and compare expectations using small bias. This gives a constant expected weight, with no width restriction.

4. Combine these bounds using the profile enumeration from Section 2.5, and verify the total seed length.

4.1 Density, shortening, and round obligations

A formula of positive density has no identically false term, and any identically true terms can be removed. We may therefore write

$$F(x) = \bigwedge_{j\in\mathcal{O}} C_j(x_{B_j}) \wedge \bigwedge_{j\in\mathcal{A}} \mathbf{1}\left[\bigoplus_{i\in B_j} x_i = b_j\right],$$

where the nonempty sets B_j are pairwise disjoint. The index sets $\mathcal{O}$ and $\mathcal{A}$ identify the OR clauses and parity constraints, respectively. Write $m_\oplus := |\mathcal{A}|$. Throughout this section, set

$$\Lambda := 1 + \log_2(1/\varepsilon), \qquad W := \lceil \log_2(8n) \rceil.$$

As before, $F^{\leq W}$ is obtained by keeping any W literals in each OR clause wider than W. Every parity constraint keeps all its variables. Shortening a parity constraint in this way would not preserve the implication to the original constraint.

Lemma 4.1 (Density bounds and shortening)**.** *Define* $Z_{\mathrm{OR}}(F) := \sum_{j\in\mathcal{O}} 2^{-|B_j|}$*. If* F *has density at least* ε*, then*

$$Z_{\mathrm{OR}}(F) + m_\oplus \leq \log_2(1/\varepsilon).$$

Moreover, $F^{\leq W}(x) = 1$ *implies* $F(x) = 1$*, its parity constraints are unchanged, and*

$$Z_{\mathrm{OR}}(F^{\leq W}) + m_\oplus \leq \log_2(1/\varepsilon) + \tfrac{1}{8} \leq \Lambda.$$

Proof. The density calculation from the read-once CNF proof gains one factor of 1/2 per parity constraint. Disjointness gives

$$\Pr[F(\mathcal{U}_n) = 1] = 2^{-m_\oplus} \prod_{j\in\mathcal{O}} (1 - 2^{-|B_j|}) \geq \varepsilon.$$

Taking binary logarithms yields

$$\log_2(1/\varepsilon) \geq m_\oplus + \sum_{j\in\mathcal{O}} -\log_2(1 - 2^{-|B_j|}) \geq m_\oplus + Z_{\mathrm{OR}}(F),$$

where the last inequality uses $-\log_2(1-u) \geq u$. As in the RO-CNF proof, deleting literals makes an OR harder to satisfy and increases the total OR sum by at most $n2^{-W} \leq 1/8$. No parity constraint is changed. Adding this increase proves the bound for $F^{\leq W}$. □

Fix an ordered partition $\mathcal{P} = (I_1, \ldots, I_J)$, with $J \geq 2$. For each term, use the same completion data as before:

$$D_{j,h} := |B_j \cap I_h|, \qquad \ell_j := \max_{i\in B_j} \tau_i, \qquad \rho_j := D_{j,\ell_j}.$$

An OR clause is alive before round k if none of its earlier literals has been satisfied. A parity constraint remains an equation on its unassigned variables; assigning earlier variables only changes its right-hand side.

At a history $y_{<k}$, the round obligation consists of the remaining OR of each alive clause with $\ell_j = k$, and the remaining parity equation of each parity constraint with $\ell_j = k$. Write $\mathcal{S}_k(y_{<k})$ for the set of assignments to I_k meeting this obligation. The constraints use disjoint, nonempty variable sets, so $\mathcal{S}_k(y_{<k})$ is nonempty. As in the RO-CNF proof, a history is valid if it has met every earlier obligation, and a full valid path satisfies F.

For each parity constraint $j \in \mathcal{A}$, define its *weight* by

$$c_j := \begin{cases} \rho_j, & \ell_j < J, \\ 1, & \ell_j = J. \end{cases}$$

This weight depends only on the partition. For every round k, define the round parameter by

$$K_k(y_{<k}) := \#\{j \in \mathcal{O} : \ell_j = k,\ C_j \text{ is alive}\} + \sum_{\substack{j \in \mathcal{A} \\ \ell_j = k}} c_j.$$

Thus the final round parameter K_J counts all remaining constraints, one per OR clause and one per parity equation.

For $k < J$, obtain a sufficient event $\mathcal{T}_k(y_{<k}) \subseteq \mathcal{S}_k(y_{<k})$ as follows. Select one remaining literal from each alive OR clause completing now and fix its satisfying value. For each parity constraint completing now, choose one satisfying assignment to all of its ρ_j remaining variables and fix those values. These coordinates are distinct, and $c_j = \rho_j$ for parities in a nonfinal round. Hence $\mathcal{T}_k$ fixes exactly K_k bits and has uniform probability 2^{-K_k}. These choices are used only in the proof; the round sources do not need to know them. In the final round, the sufficient event will impose K_J affine equations. Each parity constraint contributes one equation, which may involve many variables. The small-bias argument below handles these equations directly.

4.2 Supported assignments for a fixed partition

The potential. Keep the truncation thresholds $t_1, \dots, t_J$, with $t_J = 0$. For every $k \in [J]$, define

$$\Phi_k(y_{<k}) := K_k(y_{<k}) + \sum_{\substack{j \in \mathcal{O}:\ \ell_j > k \\ C_j \text{ is alive at } y_{<k}}} \prod_{h=k}^{\ell_j - 1} 2^{-\min(D_{j,h}, t_h)} + \sum_{\substack{j \in \mathcal{A} \\ \ell_j > k}} c_j.$$

The OR sum is exactly the one in (2). The new sum records all future parity weights. They have no survival factors, because an earlier assignment cannot remove a parity constraint whose completion round has not arrived. When a parity completes, its weight moves from this sum into K_k. Both future sums are empty in the final round, so $\Phi_J = K_J$.

At the empty history, every OR clause is alive. Therefore

$$\Phi_1(F, \mathcal{P}) = \sum_{j \in \mathcal{O}} \Psi(B_j, \mathcal{P}) + \sum_{j \in \mathcal{A}} c_j, \qquad \Psi(B_j, \mathcal{P}) = \prod_{h < \ell_j} 2^{-\min(D_{j,h}, t_h)}. \tag{14}$$

The contribution of each OR clause is unchanged. In particular, with no parity constraints, these round parameters and potentials are exactly the ones used for RO-CNF.

The round sources. For $k < J$ and each integer $r \geq 0$, use the same type of source as in the RO-CNF proof: $G_{k,r}$ is a source on n bits that is $(r + t_k)$-wise $\eta_{k,r}$-probably uniform, where

$$\eta_{k,r} := \frac{2^{-(r+t_k)}}{100J}.$$

The construction of Hoza and Lv [HL25] has seed length

$$s_{k,r} = O\big(r + t_k + \log(J+1) + \log\log(n+2)\big).$$

If $r + t_k > n$, use the uniform n-bit source, which satisfies the same guarantee and seed bound.

For the final round, let $G_{J,r}$ be an explicit small-bias source on n bits with bias

$$\beta_r := 2^{-(r+3)}.$$

By Naor and Naor [NN93], its seed length is

$$s_{J,r} = O(\log n + r).$$

In round k, use the projection $G_{k,r}|_{I_k}$. Projection preserves probable uniformity and small bias: every event or parity function on I_k is also one on the full set of coordinates. For each k and source parameter r, the source is fixed independently of F, the history, and the density parameter.

Lemma 4.2 (Nonfinal-round descent). *Fix a valid history $y_{<k}$ with $k < J$, set $K := K_k(y_{<k})$, and let X be the output of $G_{k,K}|_{I_k}$ on a uniform seed. With $\alpha := 1 + 1/(20J)$,*

$$\Pr[X \in \mathcal{T}_k] \geq (1 - \eta_{k,K})2^{-K} > 0, \qquad \mathbb{E}[\Phi_{k+1}(y_{<k}, X) \mid X \in \mathcal{T}_k] \leq \alpha\big(\Phi_k(y_{<k}) - K\big).$$

Consequently, some $y_k \in \operatorname{supp}(G_{k,K}|_{I_k}) \cap \mathcal{S}_k(y_{<k})$ satisfies the descent inequality without the expectation.

Proof. The event $\mathcal{T}_k$ fixes exactly K coordinates. For every alive OR clause C_j completing later, the event $E'_{j,k}$ requiring forbidden values on $d := \min(D_{j,k}, t_k)$ selected variables uses coordinates disjoint from $\mathcal{T}_k$. Lemma 3.3 therefore gives the positive probability of $\mathcal{T}_k$ and

$$\Pr[X \in E'_{j,k} \mid X \in \mathcal{T}_k] \leq \alpha 2^{-d}.$$

True survival implies $E'_{j,k}$. Multiplying this bound by the clause's remaining potential factors gives the same contribution estimate as in Section 3.2, including clauses completing in the next round.

Every future parity contributes the same fixed weight c_j to $\Phi_k - K$ and to Φ_{k+1}; if it completes next, this weight enters K_{k+1}. On $\mathcal{T}_k$, all terms completing now are satisfied. Adding the OR estimates and these deterministic parity contributions, using $\alpha \geq 1$, proves the descent inequality. Since the conditioning event has positive probability, one supported assignment has next potential no larger than this conditional average. □

Lemma 4.3 (Completing the final round). *Fix a valid history $y_{<J}$, put $K := K_J(y_{<J})$, and let $r \geq K$. If X is the output of $G_{J,r}|_{I_J}$ on a uniform seed, then*

$$\Pr[X \in \mathcal{S}_J(y_{<J})] \geq \tfrac{7}{8}\, 2^{-K} > 0.$$

Proof. Suppose the final obligation contains a OR clauses and m parity constraints. By definition, $K = a + m$. Select one remaining literal from each OR clause and require it to be true. Keep each remaining parity equation. Let $\mathcal{T}_J$ be the set of assignments satisfying these K requirements. Then $\mathcal{T}_J \subseteq \mathcal{S}_J$.

Each requirement is an affine equation over $\mathbb{F}_2$, and their variable sets are disjoint and nonempty. Thus they are independent under a uniform assignment, and

$$\Pr[\mathcal{U}_{I_J} \in \mathcal{T}_J] = 2^{-K}.$$

To compare this probability with the one under X, write the equations as $\bigoplus_{i \in V_s} x_i = d_s$, for $s \in [K]$. Their common indicator is

$$\mathbf{1}[\mathcal{T}_J](x) = 2^{-K} \prod_{s=1}^{K} \big(1 + (-1)^{d_s} \chi_{V_s}(x)\big).$$

Here $\chi_V(x) = (-1)^{\sum_{i \in V} x_i}$. Expanding the product gives 2^K distinct parity functions, each with coefficient of absolute value 2^{-K}, because the sets V_s are disjoint and nonempty. Therefore, writing $\|\widehat{h}\|_1 := \sum_S |\widehat{h}(S)|$,

$$\|\widehat{\mathbf{1}[\mathcal{T}_J]}\|_1 = 1.$$

A β_r-biased source changes the expectation of any function h from its uniform expectation by at most $\beta_r \|\widehat{h}\|_1$: expand h in parity functions, whose nonconstant expectations have absolute value at most β_r. Applying this to $\mathbf{1}[\mathcal{T}_J]$ gives

$$\begin{aligned}\Pr[X \in \mathcal{S}_J] &\geq \Pr[X \in \mathcal{T}_J] \\ &\geq 2^{-K} - \beta_r \geq 2^{-K} - 2^{-(K+3)} = \tfrac{7}{8} 2^{-K} > 0.\end{aligned}$$

Here $\beta_r = 2^{-(r+3)} \leq 2^{-(K+3)}$, since $r \geq K$. The empty system, when $K = 0$, also satisfies these bounds. □

Theorem 4.4 (Supported path for a fixed partition)**.** *Let F be a read-once CNF with parity constraints on disjoint nonempty variable sets. For every ordered partition and the thresholds above, there exist a valid path $y = (y_1, \dots, y_J)$ and a profile $q \in \mathbb{Z}_{\geq 0}^J$ such that*

$$q_k = K_k(y_{<k}), \qquad y_k \in \operatorname{supp}(G_{k,q_k}|_{I_k}) \cap \mathcal{S}_k(y_{<k}) \quad (k \in [J]),$$

and

$$\sum_{k=1}^{J} q_k \leq e^{1/20} \Phi_1(F, \mathcal{P}).$$

In particular, $F(y) = 1$. No density assumption on F is needed.

Proof. Starting with the empty history, set $q_k := K_k(y_{<k})$ and use Lemma 4.2 to choose a supported assignment in each nonfinal round. In the final round, set $q_J := K_J(y_{<J})$ and apply Lemma 4.3. Every assignment meets its obligation, so the resulting path is valid and satisfies F.

With $\alpha = 1 + 1/(20J)$, the potential satisfies $\Phi_{k+1} \leq \alpha(\Phi_k - K_k)$ for $k < J$ and $\Phi_J = K_J$. The telescoping calculation from Section 2.3 therefore applies without change:

$$\sum_{k=1}^{J} q_k = \sum_{k=1}^{J} K_k \leq \alpha^{J-1} \Phi_1(F, \mathcal{P}) \leq e^{1/20} \Phi_1(F, \mathcal{P}).$$

All choices belong to the required source supports by the two lemmas. □

4.3 The same partition distribution

Use the partition distribution $\mathfrak{P}$ from Section 3.3, with the same parameters:

$$r_\star := \lceil \log_2 W \rceil, \quad b := r_\star + 1, \quad J := 2r_\star + 1, \quad N := nb, \quad \delta := (16n)^{-12},$$

$$p_{2r-1} = p_{2r} := 2^{-(r+1)} \ (1 \le r \le r_\star), \qquad p_J := 2^{-r_\star},$$
$$t_k := \lceil 4Wp_k + 3\log_2(J+1) \rceil \ \ (k < J), \qquad t_J := 0.$$

The input bits are drawn from the δ-biased source M already constructed there. Apply the same sampling function f to each of its n blocks of b bits to obtain the round numbers and the partition. Thus $\mathfrak{P}$ uses $O(\log n)$ seed bits, depends only on n, and has $J = \Theta(\log\log n)$ and $\sum_k t_k = O(\log n)$.

For every fixed set B with $1 \le |B| \le W$, Lemma 3.8 gives

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Psi(B,\mathcal{P})] \le 16\, 2^{-|B|}. \tag{15}$$

This already bounds each OR contribution. It remains to bound the weight of a parity constraint of arbitrary width under this same source.

4.4 The contribution of a parity constraint

For a parity constraint on B, the contribution before the final round is $\rho(B,\mathcal{P})\mathbf{1}[\ell(B,\mathcal{P}) < J]$. The RO-CNF comparison for sets of at most W variables does not apply, because B can be much larger. We need an estimate that holds for every width.

The idea is to examine the number of variables left after each round. When many are left on average, completion in that round requires this number to fall far below its mean. A fourth power will give an upper bound on the completion event. When few are left on average, we can simply count all variables assigned in that round.

We combine these bounds into one function H that is at least the nonfinal contribution for every input. We then bound its expectation under independent uniform bits and show that small bias changes this expectation by little. As in the RO-CNF proof, the comparison is made for the function that bounds the contribution from above. A parity completing in the final round has weight one, which we add after proving the next lemma.

Lemma 4.5 (Parity completion under small bias). *Let M be a δ-biased distribution on the N input bits, and let $\mathcal{P}$ be the partition obtained by applying the split-dyadic sampling function to its blocks. For every nonempty $B \subseteq [n]$, with no restriction on its size,*

$$\mathbb{E}_M[\rho(B,\mathcal{P})\mathbf{1}[\ell(B,\mathcal{P}) < J]] \le 16 + 81Jn^5\delta.$$

Proof. Fix B and write $w := |B|$, $D_k := |B \cap I_k|$, $\ell := \max_{i\in B} \tau_i$, and $\rho := D_\ell$. If $w = 1$, the contribution is at most one. Assume $w \ge 2$. Let U denote independent uniform bits used to sample the partition, so the round numbers are independent under U. Write

$$p_k := \Pr_U[\tau_i = k], \qquad \theta_k := \Pr_U[\tau_i > k], \qquad N_k := \#\{i \in B : \tau_i > k\}.$$

These are the same split-dyadic probabilities as before. Only the completion round contributes to the identity

$$\rho \mathbf{1}[\ell < J] = \sum_{k<J} D_k \mathbf{1}[N_k = 0] :$$

before completion $N_k > 0$, and after completion $D_k = 0$.

Step 1: choose an upper bound for every input. Put $x_k := w\theta_k$ and $\mu_k := (w-1)\theta_k$. The first is the uniform mean of N_k. The second is its mean after conditioning on one specified variable being assigned in round k, since only the other $w-1$ variables can then be assigned later. Define

$$H_k := \begin{cases} D_k, & x_k \le 4, \\ D_k \left(\dfrac{N_k - \mu_k}{\mu_k} \right)^4, & x_k > 4, \end{cases} \qquad H := \sum_{k<J} H_k.$$

In the second case, $\mu_k \ge x_k/2 > 2$, so the denominator is positive. When $N_k = 0$, the fourth-power factor is one; at every other input it is nonnegative. In both cases we therefore have

$$D_k \mathbf{1}[N_k = 0] \le H_k, \qquad \rho \mathbf{1}[\ell < J] \le H.$$

Step 2: bound the expectation under independent bits. For $x_k \le 4$, the inequality $p_k \le \theta_k/2$ gives

$$\mathbb{E}_U[H_k] = w p_k \le x_k/2.$$

For $x_k > 4$, let $Z_k \sim \mathrm{Bin}(w-1, \theta_k)$, with mean μ_k. Expanding D_k as a sum of indicators gives

$$\begin{aligned} \mathbb{E}_U[D_k(N_k - \mu_k)^4] &= \sum_{i \in B} p_k \, \mathbb{E}_U[(N_k - \mu_k)^4 \mid \tau_i = k] \\ &= w p_k \, \mathbb{E}[(Z_k - \mu_k)^4]. \end{aligned}$$

Indeed, conditioned on $\tau_i = k$ under U, the later-round indicators of the other $w-1$ variables remain independent Bernoulli variables with mean θ_k. This is the only conditioning used in this calculation.

To bound the fourth moment, write $Z_k - \mu_k$ as a sum of independent centered Bernoulli variables. A term in the fourth-power expansion has zero expectation if some index appears only once. The remaining terms have either one index appearing four times or two indices appearing twice each. The second and fourth centered moments of a Bernoulli variable of mean θ_k are at most θ_k, so

$$\mathbb{E}[(Z_k - \mu_k)^4] \le (w-1)\theta_k + 6\binom{w-1}{2}\theta_k^2 \le \mu_k + 3\mu_k^2.$$

Consequently,

$$\mathbb{E}_U[H_k] \le \frac{w p_k (\mu_k + 3\mu_k^2)}{\mu_k^4} \le \frac{4 w p_k}{\mu_k^2} \le \frac{8}{x_k}.$$

The second inequality uses $\mu_k > 2$, and the last uses $w p_k \le x_k/2$ and $\mu_k \ge x_k/2$. The two estimates $x_k/2$ and $8/x_k$ agree at $x_k = 4$, which explains the cutoff.

Now $x_{k+1} \le (3/4)x_k$ between consecutive nonfinal rounds. The rounds with $x_k \le 4$ form a final interval. Starting at its first round, the bound $x_k/2$ is at most $2(3/4)^a$ after a further rounds. The rounds with $x_k > 4$ form an initial interval. Going backwards from its last round, the bound $8/x_k$ is also at most $2(3/4)^a$ after a steps. Each interval therefore contributes at most $2\sum_{a\geq 0}(3/4)^a = 8$; an empty interval contributes zero. Thus

$$\mathbb{E}_U[H] \leq 16.$$

Step 3: compare expectations using small bias. We bound the Fourier norm of H. Each event $\{\tau_i = k\}$ has indicator norm one, and each event $\{\tau_i > k\}$ has indicator norm at most two, by Lemma 3.6. Since these indicators sum to D_k and N_k, respectively,

$$\|\widehat{D_k}\|_1 \leq w, \qquad \|\widehat{N_k - \mu_k}\|_1 \leq 2w + \mu_k \leq 3w.$$

We use the general inequality $\|\widehat{gh}\|_1 \leq \|\widehat{g}\|_1 \|\widehat{h}\|_1$, which follows by multiplying the Fourier expansions and applying the triangle inequality. It holds even when g and h use the same input bits. For $x_k > 4$, this gives

$$\|\widehat{H_k}\|_1 \leq \frac{w(3w)^4}{\mu_k^4} \leq 81n^5,$$

using $\mu_k > 2$ and $w \leq n$. For $x_k \leq 4$, we have $H_k = D_k$, so the same bound holds. Summing over rounds yields

$$\|\widehat{H}\|_1 \leq 81Jn^5.$$

Finally, apply the small-bias comparison from Lemma 4.3 to H on the N sampling bits:

$$\begin{aligned}
\mathbb{E}_M[\rho\mathbf{1}[\ell < J]] &\leq \mathbb{E}_M[H] \\
&\leq \mathbb{E}_U[H] + \delta\|\widehat{H}\|_1 \\
&\leq 16 + 81Jn^5\delta.
\end{aligned}$$

All independence and conditioning arguments were under U; the last comparison is the only step that uses the distribution M. □

Including the final round. The full parity weight is $c_j = \rho_j\mathbf{1}[\ell_j < J] + \mathbf{1}[\ell_j = J]$. For the chosen bias $\delta = (16n)^{-12}$, the error in the lemma is less than one. Indeed, $J \leq 2W - 1 \leq 4n$ for $n \geq 2$, so

$$81Jn^5\delta \leq \frac{324}{16^{12}n^6} < 1.$$

A parity completing in the final round contributes at most one. Therefore every parity constraint, of any width, satisfies

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[c_j] \leq 17 + 81Jn^5\delta \leq 18.$$

4.5 Partition selection for the full formula

Theorem 4.6 (Partition selection with parities)**.** *Use the same partition distribution $\mathfrak{P}$ and thresholds as in the RO-CNF construction. They depend only on n, satisfy $t_J = 0$ and*

$\sum_k t_k = O(\log n)$*, and the partition is sampled with* $O(\log n)$ *bits. For every read-once CNF with parity constraints* F *of density at least* ε*,*

$$F^{\le W}(x) = 1 \implies F(x) = 1, \qquad \mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\le W}, \mathcal{P})] \le 18\Lambda.$$

Consequently, some partition in $\operatorname{supp}(\mathfrak{P})$ *has initial potential at most* 18Λ*.*

Proof. The parameter and sampling bounds follow from the RO-CNF construction, and Lemma 4.1 gives the implication $F^{\le W}(x) = 1 \Rightarrow F(x) = 1$. Put $Z := Z_{\mathrm{OR}}(F^{\le W})$ and $m := m_\oplus(F)$. The parity constraints are unchanged by shortening, and the same lemma gives

$$Z + m \le \log_2(1/\varepsilon) + \tfrac{1}{8} \le \Lambda.$$

Each OR clause now has width at most W, so its expected contribution is at most $16\,2^{-|B_j|}$ by (15). Each parity constraint has expected weight at most 18, by the preceding calculation. Summing these bounds in the initial potential (14) gives

$$\mathbb{E}_{\mathcal{P}\sim\mathfrak{P}}[\Phi_1(F^{\le W}, \mathcal{P})] \le 16Z + 18m \le 18(Z + m) \le 18\Lambda.$$

At least one partition in the support has potential no larger than this average. □

4.6 The generator and its seed length

We now apply the abstract composition theorem. Its final-round seed bound allows the small-bias source used here, so the same profile enumeration and seed encoding apply.

Proof of Theorem 1.2. If $\varepsilon < 2^{-n}$, outputting the n seed bits suffices, with $n \le \log_2(1/\varepsilon)$. Assume $2^{-n} \le \varepsilon \le 1/2$. Use the obligations, round parameters, potential, and sources constructed above. Theorem 4.4 gives the supported-path property with $C_{\mathrm{path}} = e^{1/20}$. Theorem 4.6 gives the partition-selection property with

$$\widehat{F} = F^{\le W}, \qquad B(n, \varepsilon) = 18\Lambda.$$

In particular, $\widehat{F}(x) = 1$ implies $F(x) = 1$, and some partition in the sampler's support has $\Phi_1(\widehat{F}, \mathcal{P}) \le 18\Lambda$. These are precisely the two class-specific statements required by Theorem 2.3.

The nonfinal sources satisfy its usual seed bound, and the final source has seed length $O(r + \log n)$, which satisfies its extended bound. The partition sampler and all sources are explicit, and

$$s_{\mathrm{part}} = O(\log n), \qquad \sum_k t_k = O(\log n), \qquad J = O(\log\log n).$$

Thus Theorem 2.3 gives an explicit HSG with seed length $O(\Lambda + \log n) = O(\log(n/\varepsilon))$. Concretely, its construction enumerates the partition seeds, all profiles of total size at most $L = \lceil 18e^{1/20}\Lambda \rceil$, and all corresponding round-source seeds. The supported path guaranteed above is among these choices. □

4.7 The width-3 hitting-set generator

Proof of Theorem 1.1. The GMRTV reduction stated in Section 1.4 now gives the result. Use density parameter $\eta := (\varepsilon/n)^{C_G}$ and enumerate every prefix length $k \in \{0, \ldots, n\}$. Output 0^k followed by the CNF-with-parities generator on the last $n-k$ variables; when $n-k \leq 1$, include every assignment to those variables. For one of these prefixes, Theorem 1.4 supplies a read-once formula of density at least η whose satisfying assignments, after adding the prefix, are accepted by the program. A seed hitting that formula therefore hits the original width-3 program. Choosing the prefix uses $O(\log n)$ bits, and

$$O(\log n) + O(\log(n/\eta)) = O(\log(n/\varepsilon)).$$

Thus the full width-3 HSG has the claimed seed length. □

5 Future research directions

The potential-descent framework separates an HSG construction into two tasks: proving descent within fixed round-source supports, and constructing a distribution over partitions with a small expected initial potential. We ask whether this approach can give improved HSGs for other function classes.

5.1 Applications to other function classes

A natural direction is to apply the framework to classes for which iterated pseudorandom restrictions have already led to strong PRG constructions. Examples include:

- **Constant-depth read-once formulas.**
  Doron, Hatami, and Hoza [DHH19] constructed PRGs for read-once formulas over AND, OR, and NOT using iterated restrictions. Extending our analysis from read-once CNFs to depth-three read-once formulas would be a natural first step toward this broader class.

- **ANDs and XORs of functions on disjoint blocks.**
  These classes allow more general functions in place of the OR and parity constraints studied here. Lee [Lee19] obtained nearly optimal PRGs for such functions when each block has bounded size. Their disjoint variable sets make them a natural setting in which to seek extensions of our potential function.

- **Constant-width monotone ROBPs.**
  Doron, Meka, Reingold, Tal, and Vadhan [DMR+21] used iterated restrictions to reduce the width of these programs and eventually obtain a function depending on few variables. It would be interesting to determine whether this simplification can instead be used to bound the total round parameters along an accepting sequence.

For each class, the main question is how to turn simplification under restrictions into suitable round obligations and a potential function. The descent argument must ensure that an assignment in the current source's support both satisfies the current obligation and leaves a small potential for later rounds. We must also bound the expected initial potential using the function's acceptance density. Existing restriction arguments provide a starting point, but these additional properties require separate proofs. The goal is to obtain simpler constructions or improved seed-length bounds by using the freedom to preserve one accepting sequence.

5.2 Beyond width three

Another direction is to extend the construction to general ordered ROBPs of larger constant width, beginning with width four. Our width-three result uses the hitting reduction of Gopalan, Meka, Reingold, Trevisan, and Vadhan [GMR+12] to read-once CNFs with parities. An extension could seek a corresponding reduction for larger widths, or apply the potential-descent framework directly to branching programs.

For a direct approach, one challenge is to define obligations that account for how assignments at different positions affect the remaining accepting paths. The potential would need to control these interactions while keeping the total round parameters small. Unlike the formula proof, such an argument would have to work without a given decomposition into constraints on disjoint variable sets.

The long-term goal is an explicit HSG with seed length

$$O_w\left(\log \frac{n}{\varepsilon}\right)$$

for every fixed width w, where the implicit constant may depend on w. Establishing suitable obligations and potential descent even for a restricted class of width-four programs would be a useful first step.

References

[AW85] Miklós Ajtai and Avi Wigderson. Deterministic Simulation of Probabilistic Constant Depth Circuits. In *26th Annual Symposium on Foundations of Computer Science (FOCS 1985)*, pages 11–19, 1985.

[BDVY13] Andrej Bogdanov, Zeev Dvir, Elad Verbin, and Amir Yehudayoff. Pseudorandomness for Width-2 Branching Programs. *Theory of Computing*, 9(7):283–293, 2013.

[BRRY14] Mark Braverman, Anup Rao, Ran Raz, and Amir Yehudayoff. Pseudorandom Generators for Regular Branching Programs. *SIAM Journal on Computing*, 43(3):973–986, 2014.

[CH20] Kuan Cheng and William M. Hoza. Hitting Sets Give Two-Sided Derandomization of Small Space. In *35th Computational Complexity Conference (CCC 2020)*, LIPIcs 169, 10:1–10:25, 2020.

[CDG26] Gil Cohen, Dean Doron, and Noam Goldgraber. A Forward–Backward Weight Analysis of INW for Permutation Branching Programs. arXiv:2607.18168, 2026.

[DHH19] Dean Doron, Pooya Hatami, and William M. Hoza. Near-Optimal Pseudorandom Generators for Constant-Depth Read-Once Formulas. In *34th Computational Complexity Conference (CCC 2019)*, LIPIcs 137, 16:1–16:34, 2019.

[DHH20] Dean Doron, Pooya Hatami, and William M. Hoza. Log-Seed Pseudorandom Generators via Iterated Restrictions. In *35th Computational Complexity Conference (CCC 2020)*, LIPIcs 169, 6:1–6:36, 2020.

[DMR+21] Dean Doron, Raghu Meka, Omer Reingold, Avishay Tal, and Salil Vadhan. Pseudorandom Generators for Read-Once Monotone Branching Programs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, LIPIcs 207, 58:1–58:21, 2021.

[GMR+12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Better Pseudorandom Generators from Milder Pseudorandom Restrictions. In *53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012)*, pages 120–129, 2012. Full version: ECCC TR12-123 / arXiv:1210.0049.

[HL25] William M. Hoza and Zelin Lv. Fooling Near-Maximal Decision Trees. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025)*, LIPIcs 353, 35:1–35:24, 2025.

[HZ20] William M. Hoza and David Zuckerman. Simple Optimal Hitting Sets for Small-Success RL. *SIAM Journal on Computing*, 49(4):811–820, 2020.

[Lee19] Chin Ho Lee. Fourier Bounds and Pseudorandom Generators for Product Tests. In *34th Computational Complexity Conference (CCC 2019)*, LIPIcs 137, 7:1–7:25, 2019.

[LV20] Chin Ho Lee and Emanuele Viola. More on Bounded Independence Plus Noise: Pseudorandom Generators for Read-Once Polynomials. *Theory of Computing*, 16(7):1–50, 2020.

[MRT19] Raghu Meka, Omer Reingold, and Avishay Tal. Pseudorandom Generators for Width-3 Branching Programs. In *51st Annual ACM Symposium on Theory of Computing (STOC 2019)*, pages 626–637, 2019.

[NN93] Joseph Naor and Moni Naor. Small-Bias Probability Spaces: Efficient Constructions and Applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.

[Nis92] Noam Nisan. Pseudorandom Generators for Space-Bounded Computation. *Combinatorica*, 12(4):449–461, 1992.

[SZ11] Jiří Šíma and Stanislav Žák. Almost k-Wise Independent Sets Establish Hitting Sets for Width-3 1-Branching Programs. In *Computer Science – Theory and Applications (CSR 2011)*, LNCS 6651, pages 120–133, 2011.

[SZ21] Jiří Šíma and Stanislav Žák. A Polynomial-Time Construction of a Hitting Set for Read-Once Branching Programs of Width 3. *Fundamenta Informaticae*, 184(4):307–354, 2021. Full version: arXiv:2101.01151v3 (2022).

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il