

Constant-Probability Witness Isolation Implies $\text{NP} \subseteq \text{P/poly}$

Sebastian Ben Daniel

Ben-Gurion University of the Negev

23 September 2026

Exposition revised 9 October 2026

Abstract

A pruning procedure maps a Boolean circuit to a circuit on the same variables that accepts only satisfying assignments of the original. Valiant and Vazirani give a pruning procedure that leaves exactly one satisfying assignment of every satisfiable circuit with probability $\Omega(1/n)$, where n is the number of variables. Dell, Kabanets, van Melkebeek, and Watanabe showed that success $2/3 + 1/\text{poly}$ would imply $\text{NP} \subseteq \text{P/poly}$. We show that every fixed positive success probability implies the same collapse. For $0 < \epsilon \leq 1$ it suffices to have success at least ϵ on circuits whose satisfying set is a nonempty affine subspace with at most $2^{\lfloor 2/\epsilon \rfloor}$ points. We also derive the collapse from success $10/\log L$ on affine inputs with at most $L^{1/3}$ satisfying assignments, where L is the description length, and from success $3/5$ plus an inverse-polynomial margin on inputs with one or two satisfying assignments. The procedure may be nonuniform and may read the entire circuit description. No cryptographic assumption is used.

The proof combines a pool of circuits, each with zero or one satisfying assignment, into one circuit whose satisfying assignments are tagged by points of $\mathbb{F}_2^d$. Each pool member owns an affine region of tag space. If that member is the only unsatisfiable one, the satisfying set shrinks to its region and can be listed using the other members' witnesses, so isolation can be tested on the region without a witness for the candidate. Regions chosen uniformly within each dimension $0, \dots, d-1$, with equal weight per dimension, have the property that every set of tags meets at most a $2/d$ fraction of them in exactly one point. A standard domination argument then yields polynomial advice. We also show that this counting method cannot go below $\Omega(1/\log L)$.

Keywords: witness isolation; affine spaces; tagged gadgets; nonuniform domination; selectors.

1 Introduction

A *pruning procedure* takes a Boolean circuit H and returns a circuit on the same variables that accepts only satisfying assignments of H . It *isolates* H when the returned circuit has exactly one satisfying assignment. The procedure only has to return the circuit; it does not have to find the surviving assignment or know whether it succeeded.

Valiant and Vazirani [1] give a randomized pruning procedure that isolates every satisfiable circuit with probability $\Omega(1/n)$, where n is the number of variables. Their procedure adds random linear constraints over $\mathbb{F}_2$ to the circuit, with the number of constraints chosen at random, so that with noticeable probability exactly one satisfying assignment survives.

The consequence is a randomized reduction from SAT to USAT, the promise problem of telling circuits with exactly one satisfying assignment apart from unsatisfiable ones. This is the sense of their title, “NP is as easy as detecting unique solutions.” Instances with a unique solution are a natural hope for algorithm design, since there is no ambiguity about which solution to find, and the theorem shows that this hope gives nothing against NP-hardness, at least under randomized reductions. Isolation has since become a standard tool. Toda’s theorem that the polynomial hierarchy lies in $P^{\#P}$ starts from a version of the reduction showing $NP \subseteq BPP^{\oplus P}$ [3]. A related weighting technique, the isolation lemma of Mulmuley, Vazirani, and Vazirani, places perfect matching in randomized NC [4] and underlies the nonuniform collapse $NL/poly = UL/poly$ of Reinhardt and Allender [5]. In exponential-time complexity, Calabro, Impagliazzo, Kabanets, and Paturi proved an isolation lemma for k -CNFs, showing that the Exponential Time Hypothesis already implies exponential hardness for instances of k -SAT with a unique solution [6].

In all of these applications the success probability is the price of isolation. A procedure that succeeds with probability p must be repeated about $1/p$ times before one of the outputs is uniquely satisfiable, and nobody can tell which one. A constant success probability would make isolation nearly free. Dell, Kabanets, van Melkebeek, and Watanabe (DKMW) [2] asked whether this probability can be improved, and showed that success $2/3 + 1/\text{poly}(L)$, where L is the description length, would imply $NP \subseteq P/poly$. Their argument needs the guarantee only on circuits with one or two satisfying assignments. The result leaves a success probability in the range between $\Theta(1/n)$ and $2/3$ open. The question we study is how low the success probability can be while still forcing the collapse.

Our answer is that every fixed positive probability suffices. This holds even when the guarantee is required only on circuits whose satisfying set is a nonempty affine subspace of $\mathbb{F}_2^n$ of size bounded by a constant that depends on the probability. The affine structure is hidden: the procedure sees only the circuit description, with no origin, basis, or dimension. The procedure may use the syntax of the circuit in any way it likes, and it may be nonuniform. All conclusions are implications of the form “such a procedure exists $\Rightarrow NP \subseteq P/poly$ ”; none of them is an unconditional separation.

1.1 Results

Theorem 1.1 (Every positive constant). *Fix $0 < \epsilon \leq 1$ and put $d = \lfloor 2/\epsilon \rfloor + 1$. Suppose a randomized nonuniform polynomial-size pruning transformation isolates, with probability at least ϵ , every circuit whose satisfying set is a nonempty affine subspace with at most 2^{d-1} points, at every sufficiently large description length. Then $NP \subseteq P/poly$.*

Corollary 1.2. *For each fixed $0 < \epsilon \leq 1$, an efficient nonuniform pruning transformation with affine-input isolation guarantee ϵ exists if and only if $NP \subseteq P/poly$.*

One direction is Theorem 1.1. The other is easy: if $NP \subseteq P/poly$, polynomial-size circuits for SAT find a satisfying assignment by prefix search, and the circuit accepting only that assignment isolates with probability one (Corollary 5.3).

For example, $\epsilon = 1/2$ gives $d = 5$, so success $1/2$ on affine inputs with at most 16 satisfying assignments already implies the collapse. The witness bound depends only on ϵ . These parameters are convenient, not optimal. Note also that bounded witness counts do not make the theorem trivial: if the number of witnesses is known to be at most 2^{d-1} , hashing already isolates with some probability of order $1/d$; the theorem needs the specific guarantee $\epsilon > 2/d$.

Theorem 1.3 (An inverse-logarithmic threshold). *Suppose a randomized nonuniform polynomial-size pruning transformation isolates, with probability at least $10/\log L$, every circuit of description*

length L whose satisfying set is a nonempty affine subspace with at most $L^{1/3}$ points, at every sufficiently large L . Then $\text{NP} \subseteq \text{P/poly}$.

Theorem 1.4 (One and two witnesses). *Suppose a randomized nonuniform polynomial-size pruning transformation isolates every circuit with exactly one or exactly two satisfying assignments with probability at least $3/5 + \eta(L)$, where $\eta(L) \geq 1/\text{poly}(L)$, at every sufficiently large L . Then $\text{NP} \subseteq \text{P/poly}$.*

Theorem 1.4 improves the DKMW constant $2/3$ on the same input class. Table 1 places the three statements next to the known bounds.

Success probability	Inputs on which it is required	Status
$\Omega(1/n)$	all satisfiable circuits	achievable [1]
$10/\log L$	affine, at most $L^{1/3}$ witnesses	implies collapse (Thm. 1.3)
any fixed $\epsilon > 0$	affine, at most $2^{\lceil 2/\epsilon \rceil}$ witnesses	implies collapse (Thm. 1.1)
$3/5 + 1/\text{poly}$	one or two witnesses	implies collapse (Thm. 1.4)
$2/3 + 1/\text{poly}$	one or two witnesses	implies collapse [2]

Table 1: Isolation guarantees and their consequences. Pruning is always required on every input and every random seed; the table lists where the success guarantee is required.

1.2 Proof overview

The whole proof rests on one observation, which the rest of this subsection explains:

If a candidate circuit were unsatisfiable, the isolator would be forced to isolate a known region frequently. A low isolation frequency on that region is therefore evidence that the candidate is satisfiable.

We work with the promise problem USAT: a *YES* circuit has exactly one satisfying assignment and a *NO* circuit has none. It is standard that polynomial-size circuits for USAT give $\text{NP} \subseteq \text{P/poly}$ (Lemma 2.3).

The central tool: tag space. Our construction combines several circuits into one and adds d new variables to it, which we call the *tag* $t \in \mathbb{F}_2^d$. We call $\mathbb{F}_2^d$ *tag space*. In every circuit we build, each tag extends to at most one satisfying assignment, so a satisfying assignment is identified by its tag. Tag space does three jobs.

- (i) *It gives solutions fixed coordinates.* The witnesses of the circuits we combine are unknown, but tag space is not. Whatever satisfying assignments an output circuit of the isolator keeps, they are described by the set $T \subseteq \mathbb{F}_2^d$ of their tags. Questions about an arbitrary output circuit become questions about subsets of $\mathbb{F}_2^d$.
- (ii) *It keeps solutions apart.* Two solutions with different tags are different assignments, even if the witnesses inside them coincide or are all zero.
- (iii) *It lets each combined circuit own a piece of the space.* Each circuit is assigned an affine *region* of tag space, and the construction ensures that if that circuit is unsatisfiable, all surviving solutions lie over its region (Figure 1).

The rest of the proof uses this picture throughout: the test in Step 2 looks at one region, the averaging in Step 3 is a statement about subsets of tag space, and Step 5 is about how to choose the regions.



Figure 1: Tag space $\mathbb{F}_2^2$ for the two-circuit example below, with $K_1 = \{t_1 = 0\}$ (dashed) and $K_2 = \{t_2 = 0\}$ (dotted). A filled dot is a tag that carries a satisfying assignment; an empty dot is a tag that carries none. When E_1 is unsatisfiable, exactly the tags in its region survive.

Step 1: a tagged gadget. Take m distinct circuits $Q = \{E_1 < \dots < E_m\}$ of description length ℓ , in canonical order, and assign to each E_i an affine region $K_i \subseteq \mathbb{F}_2^d$ of tag space. We compile Q into one circuit $H_{\mathcal{F}}(Q)$. Its variables are a tag $t \in \mathbb{F}_2^d$ and one witness block for each linear equation defining each region. The block for an equation of K_i must be all zeros if t satisfies the equation, and must be a satisfying assignment of E_i if t violates it. Building this circuit requires only the descriptions of the E_i and of the regions, never a witness.

The example in Table 2 and Figure 1 shows what happens. Take $d = 2$, $K_1 = \{t_1 = 0\}$, $K_2 = \{t_2 = 0\}$, and suppose E_1, E_2 have unique witnesses a, b .

(t_1, t_2)	x^1	x^2	survives if E_1 is NO?
$(0, 0)$	0	0	yes
$(0, 1)$	0	b	yes
$(1, 0)$	a	0	no
$(1, 1)$	a	b	no

Table 2: The satisfying assignments of a two-region gadget.

If both circuits are YES, each tag has exactly one completion and the four rows are the whole satisfying set. If E_1 is NO, the rows with $t_1 = 1$ disappear, because they would need a witness for E_1 . The two remaining rows lie over K_1 and can be written down from b alone. Because the tag is part of the assignment, the four rows are distinct even if $a = b = 0$.

In general (Lemma 3.3), an all-YES pool gives a satisfying set with exactly one assignment over each tag, forming an affine space of dimension d . If E_j is the only NO member, the satisfying set is exactly the part over K_j , which we call $R_j(Q)$. This set can be listed from the other members' witnesses.

Step 2: a one-sided test. Suppose an advice string supplies $m - 1$ YES circuits X together with their witnesses, and we receive a further circuit y . Put $Q = X \cup \{y\}$ and let j be the position of y . Run the isolator on $H_{\mathcal{F}}(Q)$ for a fixed list of seeds, and for each output circuit check whether

it accepts exactly one point of $R_j(Q)$. Let $\text{acc}_X(y)$ be the fraction of seeds for which it does. We *accept* y when $\text{acc}_X(y)$ is *low*.

The direction of this test is the key point. If y is NO, then $R_j(Q)$ is the entire satisfying set of the gadget, so every seed that isolates the gadget leaves exactly one point in $R_j(Q)$. An isolation guarantee p therefore forces $\text{acc}_X(y) \geq p$, approximately, and a NO candidate is never accepted. A YES candidate may also be rejected by this particular test. That is fine: we only need some test to accept it.

Step 3: every all-YES pool has an accepted member. Take a pool Q of m YES circuits. For each position j , the *designation* j is the run of the Step 2 test in which member E_j plays the candidate and the other $m - 1$ members serve as the advice set: $y = E_j$ and $X = Q \setminus \{E_j\}$. A pool of m members therefore has m designations, one per member. Every designation forms the same pool $X \cup \{y\} = Q$. Hence every designation compiles to the *same* circuit $H_{\mathcal{F}}(Q)$ and runs the isolator on the same seeds, and all designations look at the same output circuits; only the region they inspect changes. For one output circuit, let $T \subseteq \mathbb{F}_2^d$ be the set of tags of the satisfying assignments that this output circuit keeps. Designation j sees a regional singleton exactly when $|T \cap K_j| = 1$. So the average score over the m designations equals the average, over seeds, of

$$f_{\mathcal{F}}(T) = \frac{1}{m} |\{j : |T \cap K_j| = 1\}|, \quad \text{which is at most} \quad c(\mathcal{F}) = \max_{T \subseteq \mathbb{F}_2^d} f_{\mathcal{F}}(T).$$

If $c(\mathcal{F})$ is below the acceptance threshold, at least one member scores below average and is accepted. This uses no isolation guarantee on the all-YES gadget, since the bound holds for every possible T . It does use that the compiled description is literally the same for all designations: the isolator may depend on syntax, so different descriptions of the same function could produce unrelated output distributions.

Step 4: from local tests to circuits. Step 3 says that in every m -subset of YES circuits some member is accepted by the tester built from the others. A counting argument in the style of Ko [7] and Ogihara [9] then shows that, among any large set of YES circuits, one fixed advice set is accepted by a $1/(2m)$ fraction. Removing those and repeating, $O(m\ell)$ advice sets cover all YES circuits of length ℓ . The OR of their tests decides USAT with polynomial advice.

Step 5: choosing the regions. Everything now depends on making $c(\mathcal{F})$ small. Disjoint regions fail: a set with one point in each region gives $f_{\mathcal{F}}(T) = 1$. Identical regions fail for the same reason. We use regions of every dimension $0, \dots, d - 1$, uniform within each dimension and with equal total weight per dimension. To bound c , couple random regions into a descending chain of affine subspaces, each a random hyperplane of the previous one. Along the chain the intersection with T only shrinks. Once it becomes a single point, each further cut keeps that point with probability $1/2$, and once it is lost it never returns. So the expected number of dimensions with a singleton intersection is at most $1 + \frac{1}{2} + \frac{1}{4} + \dots \leq 2$, and averaging over the d dimensions gives $c(\mathcal{F}) \leq 2/d$ (Lemma 5.1). Taking $d > 2/\epsilon$ proves Theorem 1.1. Taking $d \approx \log \ell$ and replacing the full family by a polynomial-size random sample proves Theorem 1.3. Theorem 1.4 uses the ten points and lines of $\mathbb{F}_2^2$, for which $c = 3/5$.

1.3 Scope and open questions

Our guarantees are stated in terms of the description length L , since efficiency and advice are measured in L . To compare with hashing, write $p_{\text{wit}}(n)$ for a success guarantee on every affine circuit with n variables, keeping the same efficiency convention. The gadgets we build have $n \leq L$ and $n \rightarrow \infty$, so $p_{\text{wit}}(n) \geq 10/\log n$ already gives the hypothesis of Theorem 1.3 on them. Padding can change L without changing n , so the two conventions should not be swapped in other contexts.

Two questions remain. As far as we know, the best general guarantee is still $\Theta(1/n)$ from affine hashing. Our results do not decide whether $p_{\text{wit}}(n) = \omega(1/n)$ with $p_{\text{wit}}(n) = o(1/\log n)$ is possible under $\text{NP} \not\subseteq \text{P/poly}$ alone. They also do not rule out every positive constant in front of $1/\log L$. Section 7 shows that the regional counting used here cannot go below $\Omega(1/\log L)$, so new ideas would be needed for either question.

Under stronger assumptions, the range is settled. Assuming indistinguishability obfuscation (iO) secure against nonuniform adversaries, a companion paper [14] shows that any affine-input guarantee $p_{\text{wit}}(n) = \omega(1/n)$ yields $2^{o(v)}$ -size circuits for 3SAT on v variables. So under iO and the nonuniform Exponential Time Hypothesis, the hashing scale $\Theta(1/n)$ is optimal on affine inputs. That result complements the present one rather than improving it: it reaches down to $1/n$, but it needs iO and its conclusion is weaker than $\text{NP} \subseteq \text{P/poly}$. Nothing in this paper uses a cryptographic assumption.

The companion report [17] studies cryptographic security-isolation reductions, bounds for checkable filters, and an oracle-output barrier. None of it is used here.

Organization. Section 2 fixes the model and collects standard tools. Section 3 defines the tagged gadget and computes its satisfying sets. Section 4 turns the gadget into polynomial-size circuits for USAT, given a region family with small $c(\mathcal{F})$. Section 5 constructs such families and proves Theorems 1.1 and 1.3. Section 6 treats small families, proves Theorem 1.4, and shows how an extra guarantee on all-YES gadgets lowers some thresholds. Section 7 gives matching tag filters for the small families and a lower bound for the regional counting method. Section 8 discusses related work.

2 Preliminaries

All logarithms are base two. “Efficient” means computed by a possibly nonuniform family of polynomial-size circuits, whose advice depends only on the input length. Table 3 lists the parameters used throughout.

Symbol	Meaning
ℓ	description length of the source circuits in a pool
m	number of circuits in a pool, equal to the number of regions
d	number of tag bits; regions are affine subspaces of $\mathbb{F}_2^d$
$s_{\max}$	largest dimension of a region
n	number of variables of the compiled gadget
$L = L(\ell)$	padded description length of the compiled gadget
N_ρ	number of seeds ρ in the fixed seed list
$c(\mathcal{F})$	largest fraction of regions that a tag set meets in exactly one point

Table 3: Parameters.

2.1 Circuits and pruning

We use a standard encoding of bounded fan-in Boolean circuits. A description gives the number of variables and then, for each gate in topological order, its type and the indices of its at most two inputs. Naming one wire of a circuit with $S \geq 2$ variables and gates takes $O(\log S)$ bits, so such a circuit has a description of length $O(S \log S)$. Conversely, a description of length L has at most L variables and at most L gates. Padding appends bits that the decoder ignores, so it changes neither the ordered variables nor the function. For a circuit H on n variables, $\text{Sol}(H) = \{x \in \{0, 1\}^n : H(x) = 1\}$. Unless stated otherwise, the parameter in an isolation guarantee is the description length, not the gate count or the number of variables.

Definition 2.1 (Pruning transformation). A *randomized polynomial-size pruning transformation* is a possibly nonuniform circuit family $\mathcal{I} = (\mathcal{I}_L)_L$ that maps a description $\langle H \rangle$ and a random string ρ to the description of a circuit $\mathcal{I}(H; \rho)$ on the same ordered variables. Its size, seed length, and output length are polynomial in L , and for *every* input and *every* seed,

$$\text{Sol}(\mathcal{I}(H; \rho)) \subseteq \text{Sol}(H).$$

It has *affine-input isolation guarantee* $p(L)$ if $\Pr_\rho[|\text{Sol}(\mathcal{I}(H; \rho))| = 1] \geq p(L)$ for every length- L circuit H whose satisfying set is a nonempty affine subspace of $\mathbb{F}_2^n$.

Three features of this definition matter later. Pruning is required on every input, while the success guarantee concerns only some inputs. Unsuccessful outputs, including empty ones and ones with several satisfying assignments, remain part of the experiment; we never condition on success. And the success event need not be recognizable: our tests evaluate a sampled output circuit at a few points, but never decide whether it has exactly one satisfying assignment overall.

2.2 Fixing the randomness

The tests in Section 4 must run the isolator on *one* list of seeds that is shared by all inputs of a given length. The standard averaging argument gives such a list.

Lemma 2.2 (One seed list for all inputs). *Let $\mathcal{I}$ be a randomized polynomial-size pruning transformation and write $q_H = \Pr_\rho[|\text{Sol}(\mathcal{I}(H; \rho))| = 1]$. For every $\gamma \geq 1/\text{poly}(L)$ there is a list of $N_\rho = O((L+1)/\gamma^2)$ seeds $\rho_1, \dots, \rho_{N_\rho}$ such that, for every description H of length L , at least a $q_H - \gamma$ fraction of the outputs $\mathcal{I}(H; \rho_r)$ isolate H .*

Proof. This is the fixed-list argument of DKMW [2, Lemma 2.2]. Choose N_ρ independent seeds. For a fixed H , Hoeffding's inequality bounds the probability that the empirical success frequency is below $q_H - \gamma$ by $\exp(-2N_\rho\gamma^2)$. There are at most 2^{L+1} descriptions of length at most L , so for a suitable constant in N_ρ the union bound is below one. $\square$

The list works simultaneously for all inputs, whatever their individual success probabilities. Every entry prunes because $\mathcal{I}$ prunes on every seed. Nobody needs to know which seeds succeed or which inputs satisfy a promise. We hardwire the list as advice.

2.3 From unique satisfiability to NP

Write $\text{USAT} = (\text{Yes}, \text{No})$ for the promise problem whose YES instances have exactly one satisfying assignment and whose NO instances have none.

Lemma 2.3 (Unique-satisfiability transfer [1, 2, 13]). *If USAT has nonuniform polynomial-size circuits, then $\text{NP} \subseteq \text{P/poly}$. A bounded-error randomized circuit family for the promise problem also suffices.*

Proof. This is the argument of DKMW [2, Lemma 3.2]; we include the size accounting. Fix a description length L for SAT instances. One application of the Valiant–Vazirani reduction turns a satisfiable input into a uniquely satisfiable pruning with inverse-polynomial probability. Repeating polynomially many times makes the probability that no trial succeeds less than 2^{-L-1} , and a union bound over the at most 2^L instances gives one fixed polynomial-size list of random strings that works for every satisfiable instance of length L . Every pruning of an unsatisfiable input is unsatisfiable.

Apply the USAT circuit to every pruning in the list and take the OR. A satisfiable input has at least one pruning in Yes, so the OR is one, whatever the circuit answers on prunings outside the promise. An unsatisfiable input has only prunings in No, so the OR is zero. All prunings have length $L^{O(1)}$. For a randomized USAT circuit, first amplify and fix its coins by the same union bound. Polynomial-time reductions to SAT give the result for every language in NP. $\square$

2.4 Two benchmarks

The first benchmark is the success probability that hashing achieves when the dimension of the target is unknown.

Lemma 2.4 (Affine hashing without the target dimension [1, 13]). *Let $0 \leq d \leq n$. Choose i uniformly from $\{0, \dots, d\}$ and a uniform linear map $h : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^i$, and keep the points with $h(x) = 0$. For every nonempty affine $W \subseteq \mathbb{F}_2^n$ of dimension at most d ,*

$$\Pr_{i,h}[|W \cap h^{-1}(0)| = 1] \geq \frac{1}{4(d+1)}.$$

The filter is chosen without looking at W .

Proof. Write $W = v + U$ with $\dim U = s \leq d$. With probability $1/(d+1)$ we have $i = s$. Then $h|_U$ is given by a uniform $s \times s$ matrix, and if it is invertible the equation $h(v + u) = 0$ has exactly one solution $u \in U$. The probability of invertibility is $\prod_{k=1}^s (1 - 2^{-k})$, which is 1 for $s = 0$ and at least $\frac{1}{2}(1 - \sum_{k \geq 2} 2^{-k}) = \frac{1}{4}$ for $s \geq 1$. $\square$

The loss of $1/(d+1)$ comes from guessing the scale. With $d = n$ this gives the $\Omega(1/n)$ benchmark.

The second benchmark is the DKMW threshold. Its proof needs only one- and two-point inputs, and these are affine, so the affine restriction alone does not lower it.

Proposition 2.5 (DKMW on affine inputs). *If a randomized polynomial-size pruning transformation isolates every nonempty affine input with probability at least $2/3 + \varepsilon(L)$, where $\varepsilon(L) \geq 1/\text{poly}(L)$ at all large lengths, then $\text{NP} \subseteq \text{P/poly}$. The guarantee is needed only on inputs with one or two satisfying assignments.*

Proof. DKMW [2, Theorem 3.4] need two properties: a unique witness survives with probability at least p_1 , and on an input with exactly two witnesses the output fails to keep both with probability at least p_2 . If $p_1 + p_2/2 \geq 1 + 1/\text{poly}(L)$, then $\text{NP} \subseteq \text{P/poly}$. Singletons and two-point sets are affine of dimension zero and one, so pruning and the guarantee give $p_1, p_2 \geq 2/3 + \varepsilon$, hence $p_1 + p_2/2 \geq 1 + 3\varepsilon/2$. Finitely many lengths are handled by advice. $\square$

Finite randomness. Several statements choose uniformly from a set whose size is not a power of two. These describe ideal distributions. Each can be approximated with fair coins to any inverse-polynomial accuracy at polynomial cost, and pruning holds on every seed in any case. For instance, Lemma 2.4 yields an $\Omega(1/n)$ guarantee with fair coins. Section 7 is the only place where the distinction between ideal and finite-coin probabilities matters, and it is handled there.

3 The tagged gadget

This section builds the circuit from Step 1 of the overview and proves the two properties it needs. When all pool members are YES, the satisfying set is affine with one point per tag. When exactly one member is NO, the satisfying set is a region that can be listed without that member's witness.

3.1 Normalization

Fix a source length ℓ and write Yes_ℓ , No_ℓ for the length- ℓ descriptions in Yes and No. A circuit E of length ℓ has at most ℓ variables. Extend it to exactly ℓ variables by forcing the new ones to zero:

$$\widehat{E}(z) = E(z_1, \dots, z_{n(E)}) \wedge \bigwedge_{k=n(E)+1}^{\ell} (z_k = 0).$$

For $E \in \text{Yes}_\ell$, write $w(E) \in \{0, 1\}^\ell$ for its unique extended witness. Let $Z_\ell(z) = \bigwedge_{k=1}^{\ell} (z_k = 0)$ be the predicate that accepts only 0^ℓ . Pools are always sorted by the original descriptions.

3.2 Region families

Definition 3.1 (Region family). A *region family* of tag dimension $d \geq 1$ is a finite sequence $\mathcal{F} = (K_1, \dots, K_m)$ of nonempty affine subspaces of $\mathbb{F}_2^d$, with repetitions allowed. Each region comes with a fixed presentation

$$K_i = \{t \in \mathbb{F}_2^d : \lambda_{i,a}(t) = \beta_{i,a} \text{ for } 1 \leq a \leq r_i\}, \quad r_i = d - \dim K_i,$$

by linearly independent linear forms $\lambda_{i,a}$ and bits $\beta_{i,a}$. Put $s_{\max} = \max_i \dim K_i$. For a set of tags $T \subseteq \mathbb{F}_2^d$, define

$$f_{\mathcal{F}}(T) = \frac{1}{m} |\{i : |T \cap K_i| = 1\}|, \quad c(\mathcal{F}) = \max_{T \subseteq \mathbb{F}_2^d} f_{\mathcal{F}}(T).$$

Repeated regions are counted with multiplicity.

So $f_{\mathcal{F}}(T)$ is the fraction of regions that T meets in exactly one point, and $c(\mathcal{F})$ is the worst case over all tag sets, affine or not. This single number controls the main theorems.

3.3 The gadget

Definition 3.2 (Tagged gadget). Let $Q = \{E_1 < \dots < E_m\}$ be distinct length- ℓ descriptions in canonical order, and let $\mathcal{F}$ be a region family with m regions. The variables are a tag $t \in \{0, 1\}^d$, listed first, and one block $x^{i,a} \in \{0, 1\}^\ell$ for each $1 \leq i \leq m$ and $1 \leq a \leq r_i$. Define

$$G_{i,a}(t, x^{i,a}) = ([\lambda_{i,a}(t) = \beta_{i,a}] \wedge Z_\ell(x^{i,a})) \vee ([\lambda_{i,a}(t) \neq \beta_{i,a}] \wedge \widehat{E}_i(x^{i,a})), \quad (3.1)$$

$$H_{\mathcal{F}}(Q) = \bigwedge_{i,a} G_{i,a}. \quad (3.2)$$

In words: if t satisfies an equation of K_i , the corresponding block must be zero; if t violates it, the block must be a witness for E_i . The compiler receives only $\mathcal{F}$ and Q . It receives no witnesses, and it receives no indication of which member will later be tested.

Size. Let $R = \sum_i r_i \leq md$ be the number of blocks. The gadget has $n = d + R\ell$ variables. Each block (i, a) needs $O(d)$ gates to compute the bit $\lambda_{i,a}(t) \oplus \beta_{i,a}$, which is an XOR of at most d tag bits; $O(\ell)$ gates for $Z_\ell(x^{i,a})$; a copy of $\hat{E}_i$ on $x^{i,a}$, consisting of the at most ℓ gates of E_i and $O(\ell)$ gates forcing the padding variables to zero; and $O(1)$ gates for the two branches of (3.1). The final conjunction over the R blocks needs $O(R)$ gates. So the gadget has

$$S = O(d + R(\ell + d))$$

variables and gates, and by the encoding of Section 2 its unpadded description has length $O(S \log S)$. The logarithmic factor is the cost of naming wires: inside E_i a wire is named with about $\log \ell$ bits, inside the gadget with about $\log S$ bits, so the copies of the E_i are re-encoded rather than pasted in verbatim. For a fixed family, d and R are constants, so $S = O(\ell)$ and the description has length $O(\ell \log \ell)$; we pad to $L(\ell) = \ell^2$ for large ℓ . The same padding is used for every pool at source length ℓ and adds no variables.

3.4 The satisfying sets

For a position j in the pool (later, the position of the member under test), define the affine map χ_j on tags, block by block, by

$$\chi_j^{i,a}(t) = \begin{cases} (\lambda_{i,a}(t) \oplus \beta_{i,a}) w(E_i), & i \neq j, \\ 0^\ell, & i = j, \end{cases} \quad R_j(Q) = \{(t, \chi_j(t)) : t \in K_j\}.$$

Here $(\lambda(t) \oplus \beta)w$ means w if the equation fails at t and 0^ℓ if it holds. The set $R_j(Q)$ can be computed from $\mathcal{F}$, j , and the witnesses $w(E_i)$ for $i \neq j$. It does not involve E_j at all.

Lemma 3.3 (Satisfying sets of the gadget). *Let Q be a pool in which every member other than E_j is YES.*

(a) *If E_j is also YES, then*

$$S_Q := \text{Sol}(H_{\mathcal{F}}(Q)) = \{(t, \chi(t)) : t \in \mathbb{F}_2^d\}, \quad \chi^{i,a}(t) = (\lambda_{i,a}(t) \oplus \beta_{i,a}) w(E_i).$$

This is an affine space of dimension d with exactly one point over each tag, and $R_j(Q)$ is the part of S_Q lying over K_j .

(b) *If E_j is NO, then $\text{Sol}(H_{\mathcal{F}}(Q)) = R_j(Q)$, an affine space of dimension $\dim K_j$.*

Proof. Fix a tag t and look at a single block (i, a) . The two disjuncts of (3.1) exclude each other, since one requires $\lambda_{i,a}(t) = \beta_{i,a}$ and the other the opposite. If the equation holds, the only completion of the block is 0^ℓ . If it fails, the completions are the satisfying assignments of $\hat{E}_i$: exactly $w(E_i)$ if E_i is YES, and none if E_i is NO. The blocks use disjoint variables, so t extends to a satisfying assignment of $H_{\mathcal{F}}(Q)$ if and only if every block has a completion, and then the extension is unique.

(a) Every block has the completion $\chi^{i,a}(t)$. Each coordinate of χ is an affine function of t , so S_Q is the graph of an affine map. Since the tag is part of the assignment, distinct tags give distinct

points, and S_Q has dimension d ; this holds even if several $w(E_i)$ coincide or equal 0^ℓ . For $t \in K_j$ every equation of K_j holds, so the blocks of member j are zero and $\chi(t) = \chi_j(t)$.

(b) A block of member j has a completion only if its equation holds. So all blocks of member j have completions exactly when $t \in K_j$, and then they are zero. All other blocks always have completions. Hence the satisfying set is $\{(t, \chi_j(t)) : t \in K_j\} = R_j(Q)$. $\square$

We call $H_{\mathcal{F}}(Q)$ a *one-NO gadget* if exactly one member of Q is NO and the rest are YES, and an *all-YES gadget* if every member is YES. By the lemma, a one-NO gadget whose NO member owns a region of dimension s has exactly 2^s satisfying assignments forming an affine space, and an all-YES gadget has 2^d .

4 From regional tests to small circuits

The argument has two parts. This section proves the complexity part: Theorem 4.1 below shows that any region family with small $c(\mathcal{F})$, together with an isolator whose success probability exceeds $c(\mathcal{F})$, yields polynomial-size circuits for USAT. The geometric part, constructing families with small $c(\mathcal{F})$, is in Section 5.

Theorem 4.1 (Regional collapse). *Let $(\mathcal{F}_\ell)_\ell$ be region families whose tag dimension, number of regions, and $2^{s_{\max}}$ are polynomial in ℓ ; the families need not be efficiently computable, only fixed for each ℓ . Let $L(\ell) \geq \ell$ be a polynomial padding length. Suppose a randomized polynomial-size pruning transformation isolates every one-NO gadget of $\mathcal{F}_\ell$, padded to length $L(\ell)$, with probability at least $p(\ell)$, and that for all large ℓ*

$$p(\ell) - c(\mathcal{F}_\ell) \geq \eta(\ell) \geq 1/\text{poly}(\ell). \quad (4.1)$$

Then $\text{NP} \subseteq \text{P/poly}$.

No success guarantee is assumed on all-YES gadgets, but the transformation must still prune them, as it must prune every input.

Where the region family comes from. The family $\mathcal{F}_\ell$ is part of the advice of the circuit we build for USAT at length ℓ . A presentation lists, for each of the m regions, at most d linear forms and right-hand sides, so it has $O(md^2)$ bits, which is polynomial in ℓ . Since P/poly allows any polynomial-length advice that depends only on the input length, it is enough that a suitable family exists at each length; nobody has to construct it. Three points make this choice legitimate.

- (i) The property the proof uses, a bound on $c(\mathcal{F}_\ell)$, belongs to the family alone. It bounds $f_{\mathcal{F}_\ell}(T)$ for every tag set T , hence for every output of every pruning transformation. So the family can be fixed before $\mathcal{I}$ is considered, and $\mathcal{I}$ cannot be adapted to defeat it.
- (ii) In the applications (Theorems 5.2 and 5.4) the hypothesis on $\mathcal{I}$ covers all affine inputs of the stated size, so the one-NO gadgets of whichever family is fixed are covered automatically.
- (iii) One family is fixed per length and used by every tester at that length. In particular, all designations of a pool compile with the same family, as Lemma 4.3 requires.

The order of choices is therefore: the transformation $\mathcal{I}$ and its advice; for each ℓ , a family $\mathcal{F}_\ell$ with small c ; then the seed list of Lemma 2.2 for length $L(\ell)$, which works for every description of that length and hence for every gadget of $\mathcal{F}_\ell$; finally the advice sets of the covering argument.

We prove the theorem in four steps: define a tester (Section 4.1), show it never accepts a NO input (Lemma 4.2), show that every all-YES pool has an accepted member (Lemma 4.3), and cover the YES inputs with polynomially many testers (Lemma 4.5).

Throughout, fix ℓ , put $\gamma = \eta/3$ and $p' = p - \gamma$, and fix the seed list of Lemma 2.2 at length $L(\ell)$. Write $C'_1, \dots, C'_{N_\rho}$ for the outputs of the isolator on a description H under this list.

4.1 The tester

An *advice set* X consists of $m - 1$ distinct circuits in Yes_ℓ together with their witnesses. On input a length- ℓ circuit y , the tester T_X does the following.

- (1) If $y \in X$, accept.
- (2) Otherwise, let $Q = X \cup \{y\}$ in canonical order and let j be the position of y . Compile $H_{\mathcal{F}}(Q)$, pad it, and compute the outputs $C'_1, \dots, C'_{N_\rho}$.
- (3) Using the witnesses in X , list the points of $R_j(Q)$. For each r , evaluate C'_r at each of these points, and set

$$\text{acc}_X(y) = \frac{1}{N_\rho} |\{r : |\text{Sol}(C'_r) \cap R_j(Q)| = 1\}|.$$

- (4) Accept if and only if $\text{acc}_X(y) < p'$.

A few remarks on this procedure. Each output circuit is sampled once and then evaluated at every point of the region; the isolator is not rerun between evaluations. The test is regional: on an all-YES gadget, C'_r may accept further points outside $R_j(Q)$, and the tester does not look at them. The points of $R_j(Q)$ are listed by solving the equations of K_j by Gaussian elimination; this basis is used only by the tester and is never given to the isolator. The comparison in step (4) uses the integer $\lceil N_\rho p' \rceil$, which can be hardwired. The tester has size $\text{poly}(\ell, m, d, L, N_\rho) \cdot 2^{s_{\max}}$. Off the promise its answers are arbitrary.

4.2 Soundness

Lemma 4.2 (Soundness). *Under the hypotheses of Theorem 4.1, every tester T_X rejects every $y \in \text{No}_\ell$.*

Proof. Step (1) does not apply because $X \subseteq \text{Yes}_\ell$. In step (2), y is the only NO member of Q , so by Lemma 3.3(b) the whole satisfying set of $H_{\mathcal{F}}(Q)$ is $R_j(Q)$. Every output prunes, so every output that isolates the gadget has exactly one point in $R_j(Q)$. By Lemma 2.2, at least a $p - \gamma = p'$ fraction of the outputs isolate. Hence $\text{acc}_X(y) \geq p'$ and T_X rejects. $\square$

4.3 Every all-YES pool has an accepted member

Fix an all-YES pool $Q = \{E_1 < \dots < E_m\}$. For $1 \leq j \leq m$, the *designation* j of Q is the run of the tester with candidate E_j and advice set $Q \setminus \{E_j\}$, that is, the computation of $T_{Q \setminus \{E_j\}}(E_j)$. Its score is $\text{acc}_{Q \setminus \{E_j\}}(E_j)$. The next lemma averages these m scores.

Lemma 4.3 (Regional averaging). *Let $Q = \{E_1 < \dots < E_m\}$ be an all-YES pool, let $D_r = \text{Sol}(C'_r)$ for the outputs on $H_{\mathcal{F}}(Q)$, and let $\tau(t, x) = t$ be the projection to the tag. Then*

$$\frac{1}{m} \sum_{j=1}^m \text{acc}_{Q \setminus \{E_j\}}(E_j) = \frac{1}{N_\rho} \sum_{r=1}^{N_\rho} f_{\mathcal{F}}(\tau(D_r)) \leq c(\mathcal{F}). \quad (4.2)$$

In particular, if $c(\mathcal{F}) < p'$, then some E_j is accepted by $T_{Q \setminus \{E_j\}}$.

Proof. Since the members of Q are distinct, no designation stops at step (1), and every designation compiles the same description $H_{\mathcal{F}}(Q)$ and sees the same outputs $C'_1, \dots, C'_{N_\rho}$. By pruning, $D_r \subseteq S_Q$. By Lemma 3.3(a), τ is a bijection from S_Q to $\mathbb{F}_2^d$ that maps $R_j(Q)$ onto K_j . Therefore $|D_r \cap R_j(Q)| = |\tau(D_r) \cap K_j|$. Summing the indicator of $|D_r \cap R_j(Q)| = 1$ over j and r in the two possible orders gives the equality, and the inequality is the definition of $c(\mathcal{F})$. If the average of the m scores is below p' , at least one score is below p' . $\square$

Remark 4.4 (Why the candidate must not be marked). The identity (4.2) swaps a sum over designations with a sum over outputs. This works only because the outputs do not depend on the designation. If the compiled circuit recorded which member is the candidate, the m designations would be m different inputs to the isolator, which may depend arbitrarily on syntax. Their output distributions would be unrelated, and the average on the left would no longer be an average of $f_{\mathcal{F}}$ over common outputs. Different pools, including one-NO pools, may of course produce entirely different output distributions; their guarantees are used separately.

Note that (4.2) holds for every output, isolating or not. That is why no success guarantee on all-YES gadgets is needed. The bound $c(\mathcal{F})$ also covers tag sets that are not affine, which matters because an isolator's outputs need not be affine.

4.4 Covering the YES inputs

Lemma 4.5 (Domination). *Let $S \subseteq \text{Yes}_\ell$ with $N = |S| \geq 2m$. Suppose every m -element subset $Q \subseteq S$ contains a member E such that $T_{Q \setminus \{E\}}$ accepts E . Then some $(m-1)$ -element $X \subseteq S$ has a tester accepting at least $N/(2m)$ members of $S \setminus X$.*

Proof. Choose Q uniformly among the m -subsets of S and then y uniformly in Q . By hypothesis, $T_{Q \setminus \{y\}}$ accepts y with probability at least $1/m$. The map $(Q, y) \mapsto (Q \setminus \{y\}, y)$ is a bijection onto pairs (X, y) with $|X| = m-1$ and $y \in S \setminus X$, and

$$\binom{N}{m} m = \binom{N}{m-1} (N - m + 1),$$

so the same experiment chooses a uniform $(m-1)$ -set X and then a uniform $y \in S \setminus X$. Some fixed X therefore accepts at least a $1/m$ fraction of $S \setminus X$, that is, at least $(N - m + 1)/m \geq N/(2m)$ members. $\square$

4.5 Proof of Theorem 4.1

The seed list. By Lemma 2.2, one list of $N_\rho = O((L+1)/\gamma^2)$ seeds gives every one-NO gadget an isolating fraction at least $p' = p - \gamma$. By (4.1), $c(\mathcal{F}) \leq p - \eta < p'$. So Lemma 4.2 applies, and by Lemma 4.3 every all-YES pool has an accepted member.

Greedy covering. Start with $S_0 = \text{Yes}_\ell$. While $|S_k| \geq 2m$, Lemma 4.5 gives an advice set $X_k \subseteq S_k$ whose tester accepts at least a $1/(2m)$ fraction of S_k ; let S_{k+1} be S_k minus the inputs that T_{X_k} accepts. Since $|S_0| \leq 2^\ell$ and $(1 - 1/(2m))^k 2^\ell \leq e^{-k/(2m)} 2^\ell < 1$ for $k \geq 2m\ell$, at most $2m\ell + 1$ rounds occur before fewer than $2m$ inputs remain.

The circuit. Hardwire the advice sets X_k with their witnesses, the seed list, the advice of $\mathcal{I}$, the family $\mathcal{F}_\ell$, and the final residual set of fewer than $2m$ YES descriptions. Accept y if it is in the residual set or some T_{X_k} accepts it. A NO input is rejected by every tester (Lemma 4.2) and is not in the residual set. A YES input was either removed in some round, so some tester accepts it, or is residual.

Size. There are $O(m\ell)$ testers, each with $O(m\ell)$ bits of advice, and the presentation of $\mathcal{F}_\ell$ adds $O(md^2)$ bits shared by all of them. Each compiles a gadget of length L , runs the isolator N_ρ times, and evaluates each output at most $2^{s_{\max}}$ times. All of these quantities are polynomial in ℓ , so USAT has polynomial-size circuits at length ℓ . Finitely many small lengths are handled by advice, and Lemma 2.3 gives $\text{NP} \subseteq \text{P/poly}$. $\square$

This is an existence argument. It shows that good advice exists at each length; it does not claim that the advice sets, their witnesses, or good seeds can be found efficiently. Storing YES instances and their witnesses in the advice is allowed for the same reason.

5 Multiscale regions and the main theorems

By Theorem 4.1, it remains to find region families with small $c(\mathcal{F})$ and small regions. This section gives the construction from Step 5 of the overview.

5.1 Why many scales

Two natural choices fail. If the regions are pairwise disjoint, a tag set with one point in each region meets every region once, so $c(\mathcal{F}) = 1$. If all regions are equal, a single point of that region gives $c(\mathcal{F}) = 1$ again. Overlap by itself does not help. What we need is a family in which no tag set, however it is chosen, meets many regions in exactly one point. Using regions at every scale achieves this, because a tag set that meets a region of dimension s in a single point tends to meet the regions below it in that same point or not at all.

5.2 The affine-flag bound

Lemma 5.1 (Multiscale affine-flag sum). *For each $0 \leq s < d$, let K_s be a uniformly random affine subspace of $\mathbb{F}_2^d$ of dimension s . For every fixed $T \subseteq \mathbb{F}_2^d$,*

$$\sum_{s=0}^{d-1} \Pr[|T \cap K_s| = 1] \leq 2 - 2^{1-d} \leq 2. \quad (5.1)$$

Proof. This is the affine-flag argument of [13]. Each term on the left depends only on the distribution of a single K_s , so we may couple the K_s in any way that keeps these distributions. Let K_{d-1} be a uniform affine hyperplane of $\mathbb{F}_2^d$, and for $s \geq 1$ let K_{s-1} be a uniform affine hyperplane of K_s . The distribution of this chain is invariant under the affine group of $\mathbb{F}_2^d$, which acts transitively on affine

subspaces of each dimension, so each K_s is uniform among the s -dimensional ones. We use two facts.

- (i) The intersections shrink: $T \cap K_{s-1} \subseteq T \cap K_s$. So a singleton intersection is followed by the same singleton or by the empty set, and an empty intersection stays empty.
- (ii) If $x \in K_s$ and $s \geq 1$, then $x \in K_{s-1}$ with probability exactly $1/2$, whatever happened earlier in the chain. Indeed, a hyperplane of K_s is chosen by picking a codimension-one direction subspace and then one of its two cosets in K_s , and exactly one of those cosets contains x .

Let Z be the number of dimensions s with $|T \cap K_s| = 1$, so that the left side of (5.1) is $\mathbb{E}[Z]$. If no intersection is a singleton, $Z = 0$. Otherwise, let j be the first dimension, going down the chain, at which $T \cap K_j = \{x\}$. By (i), the later intersections are $\{x\}$ while x survives and empty afterwards. By (ii), x survives k further cuts with probability 2^{-k} . Hence

$$\mathbb{E}[Z \mid j, x] = \sum_{k=0}^j 2^{-k} = 2 - 2^{-j} \leq 2 - 2^{1-d},$$

since $j \leq d - 1$. Averaging over j and x proves the lemma. $\square$

5.3 Every positive constant

To turn the lemma into a finite family, we must give the dimensions equal weight. Listing every affine subspace once would not do this, since there are many more subspaces of middle dimension than of dimension 0.

Let A_s be the number of affine s -dimensional subspaces of $\mathbb{F}_2^d$ and let $M = \text{lcm}(A_0, \dots, A_{d-1})$. Let $\mathcal{F}_d$ list every s -dimensional affine subspace M/A_s times, for $s = 0, \dots, d - 1$. Then each dimension contributes exactly M regions, and by Lemma 5.1, for every T ,

$$f_{\mathcal{F}_d}(T) = \frac{1}{d} \sum_{s=0}^{d-1} \Pr[|T \cap K_s| = 1] \leq \frac{2}{d}, \quad \text{so} \quad c(\mathcal{F}_d) \leq \frac{2}{d}. \quad (5.2)$$

For fixed d the family is huge but constant, independent of ℓ . Its regions have at most 2^{d-1} points.

Theorem 5.2 (Every positive constant). *Let $0 < \epsilon \leq 1$ and $d = \lfloor 2/\epsilon \rfloor + 1$. If a randomized polynomial-size pruning transformation isolates every affine input with at most 2^{d-1} satisfying assignments with probability at least ϵ , at every sufficiently large description length, then $\text{NP} \subseteq \text{P/poly}$.*

Proof. Use the constant family $\mathcal{F}_d$ at every source length, with $L(\ell) = \ell^2$. It depends only on ϵ , so it is a fixed finite object that can be hardwired, or computed by listing all affine subspaces of $\mathbb{F}_2^d$; no length-dependent advice is needed to describe it. Its tag dimension, size, and $2^{s_{\max}}$ are constants. Its one-NO gadgets are affine with 2^s satisfying assignments for some $s \leq d - 1$, so the hypothesis applies to them with $p = \epsilon$. Since $c(\mathcal{F}_d) \leq 2/d$ and $d > 2/\epsilon$, the margin satisfies $\epsilon - c(\mathcal{F}_d) \geq \eta := \epsilon - 2/d > 0$, a positive constant, so Theorem 4.1 applies. The isolation guarantee is used only on the one-NO gadgets of $\mathcal{F}_d$. $\square$

Corollary 5.3 (Constant-success equivalence). *For every fixed $0 < \epsilon \leq 1$, an efficient nonuniform pruning transformation with affine-input isolation guarantee ϵ exists if and only if $\text{NP} \subseteq \text{P/poly}$.*

Proof. One direction is Theorem 5.2. Conversely, polynomial-size circuits for SAT allow prefix self-reduction: test satisfiability, then fix each variable to 0 if a satisfying extension remains and to 1 otherwise. Output the circuit that accepts only the resulting lexicographically first witness, or the empty circuit if the input is unsatisfiable. This prunes on every input and has success one. $\square$

5.4 An inverse-logarithmic threshold

To let d grow with ℓ , we cannot use the full family $\mathcal{F}_d$, whose size grows much faster than any polynomial. A polynomial-size random sample from the same distribution keeps c within a constant factor.

Theorem 5.4 (Logarithmic threshold). *If a randomized polynomial-size pruning transformation isolates every affine input with at most $L^{1/3}$ satisfying assignments with probability at least $10/\log L$, at every sufficiently large description length L , then $\text{NP} \subseteq \text{P/poly}$.*

Proof. A small family with $c \leq 3/d$. Let $d = \lfloor \log \ell \rfloor$, $N = 2^d \leq \ell$, and $m = Nd^2$. Sample m regions independently, each by choosing s uniformly in $\{0, \dots, d-1\}$ and then a uniform affine s -dimensional subspace. Fix $T \subseteq \mathbb{F}_2^d$. The indicators of $|T \cap K_i| = 1$ are independent, and by Lemma 5.1 each has mean at most $2/d$. Hoeffding's inequality gives

$$\Pr[f_{\mathcal{F}}(T) > 3/d] \leq \exp(-2m/d^2) = e^{-2N}.$$

A union bound over the 2^N tag sets gives total failure probability at most $2^N e^{-2N} < 1$. So some family $\mathcal{F}_\ell$ of m regions satisfies $c(\mathcal{F}_\ell) \leq 3/d$ for every T at once. For each ℓ we fix one such family and supply it as advice, as Theorem 4.1 allows. Checking $c(\mathcal{F}) \leq 3/d$ directly means examining all 2^N tag sets, so we do not claim the family can be found efficiently; its existence is enough. The choice does not depend on the pruning transformation, because the bound holds for every tag set.

Sizes. The family has at most $md = Nd^3$ equations, so the gadget has $R \leq Nd^3 \leq \ell \log^3 \ell$ blocks. By the count after Definition 3.2, it has $S = O(R\ell) = O(\ell^2 \log^3 \ell)$ variables and gates and unpadded length $O(S \log S) = O(\ell^2 \log^4 \ell)$. Pad to $L = \ell^3$. The family's presentation has $O(md^2) = O(\ell \log^4 \ell)$ bits. Every region has at most $2^{d-1} \leq \ell/2 \leq L^{1/3}$ points, so the one-NO gadgets lie in the class where success is assumed, and the tester enumerates each region in polynomial time.

The margin. Put $x = \log L = 3 \log \ell$. Then $d \geq x/3 - 1$, so

$$c(\mathcal{F}_\ell) \leq \frac{3}{d} \leq \frac{9}{x-3}, \quad \frac{10}{x} - \frac{9}{x-3} = \frac{x-30}{x(x-3)} \geq \frac{1}{2x} \quad \text{for } x \geq 57.$$

Theorem 4.1 applies with $p = 10/\log L$ and $\eta = 1/(2 \log L)$. $\square$

The constant 10 is not claimed to be optimal. Since the gadgets have $n \leq L$ variables and $n \rightarrow \infty$, the hypothesis also follows from a witness-length guarantee $p_{\text{wit}}(n) \geq 10/\log n$.

6 Small families and a guarantee on all-YES gadgets

For small families, the threshold can be lowered if the isolator is also assumed to succeed on all-YES gadgets. An output that isolates an all-YES gadget retains a single tag, and a single tag can meet only a few regions. To use this, split $c(\mathcal{F})$ into two parts:

$$g(\mathcal{F}) = \max_{t \in \mathbb{F}_2^d} f_{\mathcal{F}}(\{t\}), \quad c'(\mathcal{F}) = \max_{|T| \neq 1} f_{\mathcal{F}}(T), \quad c(\mathcal{F}) = \max\{g, c'\}.$$

Here g bounds the contribution of outputs that isolate the all-YES gadget and c' bounds all other outputs, including empty ones.

6.1 The refined collapse theorem

Theorem 6.1 (Regional collapse with an all-YES guarantee). *In the setting of Theorem 4.1, suppose the transformation isolates every one-NO gadget with probability at least $p_{\text{No}}(\ell)$ and every all-YES gadget with probability at least $p_{\text{Yes}}(\ell)$, and that for all large ℓ*

$$p_{\text{No}} - \max\{g, p_{\text{Yes}}g + (1 - p_{\text{Yes}})c'\} \geq \eta(\ell) \geq 1/\text{poly}(\ell). \quad (6.1)$$

Then $\text{NP} \subseteq \text{P/poly}$. With $p_{\text{Yes}} = 0$ this is Theorem 4.1.

Proof. Only the averaging step changes. Put $\gamma = \eta/3$, $p' = p_{\text{No}} - \gamma$, and $\sigma = \max\{p_{\text{Yes}} - \gamma, 0\}$. By Lemma 2.2, applied once to all inputs, the common list has isolating fraction at least p' on every one-NO gadget and at least σ on every all-YES gadget. Soundness is unchanged.

Fix an all-YES pool Q and let $\sigma_Q \geq \sigma$ be the fraction of outputs that isolate $H_{\mathcal{F}}(Q)$. Such an output retains a single tag and contributes at most g to the right side of (4.2); any other output contributes at most c' . So the average score is at most $\sigma_Q g + (1 - \sigma_Q)c'$. This is affine in σ_Q , so over $\sigma_Q \in [\sigma, 1]$ it is at most $\max\{g, \sigma g + (1 - \sigma)c'\}$. Replacing p_{Yes} by $\sigma \geq p_{\text{Yes}} - \gamma$ raises $p_{\text{Yes}}g + (1 - p_{\text{Yes}})c'$ by at most γ (and not at all when $c' \leq g$). By (6.1),

$$\max\{g, \sigma g + (1 - \sigma)c'\} \leq p_{\text{No}} - \eta + \gamma < p_{\text{No}} - \gamma = p'.$$

So every all-YES pool has an accepted member, and the rest of the proof of Theorem 4.1 goes through unchanged. $\square$

When the same guarantee p is assumed on both classes, the condition is $p > \rho(\mathcal{F})$ with a $1/\text{poly}$ margin, where

$$\rho(\mathcal{F}) = \max\left\{g, \frac{c'}{1 + c' - g}\right\}$$

whenever $1 + c' - g > 0$. (In the degenerate case $g = 1$, $c' = 0$ there is no threshold in $[0, 1]$.)

Remark 6.2 (Disjoint regions). For $m \geq 2$ pairwise disjoint nonempty regions, $g = 1/m$ (a tag lies in at most one region) and $c' = 1$ (take one point from each region). The symmetric threshold is $\rho = m/(2m - 1)$, which equals $(k + 1)/(2k + 1)$ for a pool of $m = k + 1$ circuits and tends to $1/2$. With $m = 2$ it is DKMW's threshold $2/3$. This limit is a property of disjoint regions, not an obstacle for overlapping multiscale families.

6.2 Points and lines of the plane

Proposition 6.3 (Points and lines of $\mathbb{F}_2^2$). *Let $\mathcal{F}$ consist of the four points and six lines of $\mathbb{F}_2^2$, once each. Then*

$$f_{\mathcal{F}}(T) = \frac{|T|(5 - |T|)}{10}, \quad c = c' = \frac{3}{5}, \quad g = \frac{2}{5}.$$

Consequently, success $3/5 + \eta$ on all affine inputs with one or two satisfying assignments implies $\text{NP} \subseteq \text{P/poly}$, and so does a guarantee $1/2 + \eta$ on all affine inputs with one, two, or four satisfying assignments. In both cases $\eta \geq 1/\text{poly}$ at all large lengths.

Proof. Over $\mathbb{F}_2$ a line is just a pair of points. A set of u tags contains u point regions and meets exactly $u(4 - u)$ lines once, so $f_{\mathcal{F}}(T) = (u + u(4 - u))/10 = u(5 - u)/10$. The values for $u = 0, \dots, 4$ are $0, \frac{2}{5}, \frac{3}{5}, \frac{3}{5}, \frac{2}{5}$. The one-NO gadgets have one or two satisfying assignments, so the first claim is Theorem 4.1 with $c = 3/5$. For the second, the all-YES gadgets have four, and $c'/(1 + c' - g) = (3/5)/(6/5) = 1/2 > g$, so Theorem 6.1 applies. $\square$

This proves Theorem 1.4, since every input with one or two satisfying assignments is affine. For the first claim, no success is assumed on the four-assignment all-YES gadgets, but the transformation must still prune them.

6.3 All lines

Proposition 6.4 (All affine lines). *Let $d \geq 2$, $N = 2^d$, and let $\mathcal{F}$ contain each affine line of $\mathbb{F}_2^d$ once. Then*

$$f_{\mathcal{F}}(T) = \frac{|T|(N - |T|)}{\binom{N}{2}}, \quad g = \frac{2}{N}, \quad c' = \frac{N}{2(N - 1)}, \quad \rho_N = \frac{N^2}{3N^2 - 6N + 4} \leq \frac{1}{3} + \frac{4}{3N}.$$

If p_2 and p_N are success guarantees on affine inputs with 2 and N satisfying assignments and

$$p_2 + \frac{(N - 2)^2}{2N(N - 1)} p_N \geq \frac{N}{2(N - 1)} + \eta, \quad \eta \geq 1/\text{poly},$$

at all large lengths, then $\text{NP} \subseteq \text{P}/\text{poly}$. In particular a symmetric guarantee $\rho_N + \eta$ suffices.

Proof. Over $\mathbb{F}_2$ every two distinct points form a line, and a line meets T once exactly when one endpoint is in T and the other is not. This gives the formula for $f_{\mathcal{F}}$. For a singleton it equals $2/N$; over $|T| \neq 1$ it is largest at $|T| = N/2$, giving c' . Since $c' - g = (N - 2)^2/(2N(N - 1)) \geq 0$, the maximum in (6.1) is $p_N g + (1 - p_N)c' = c' - p_N(c' - g)$, and the displayed condition is (6.1). Setting $p_2 = p_N = p$ gives $p > c'/(1 + c' - g) = \rho_N$. For the upper bound, $\rho_N - \frac{1}{3} = \frac{6N - 4}{3(3N^2 - 6N + 4)}$ and $3N^2 - 6N + 4 \geq 3N^2/2$ for $N \geq 4$. For fixed d , all resource bounds of Theorem 4.1 hold. $\square$

For $N = 4, 8, 16$ the symmetric thresholds are $4/7, 16/37$, and $64/169$. So every fixed symmetric guarantee strictly above $1/3$ already implies the collapse on a suitable constant-witness affine class.

6.4 Exact values in small dimensions

More generally, one can give the affine s -dimensional subspaces total weight $a_s \geq 0$, with $\sum_s a_s = 1$, and each individual one weight a_s/A_s ; rational weights are realized exactly by repetition. Table 4 lists the resulting thresholds for $d = 2, 3, 4$. The c columns need success only on one-NO gadgets; the ρ columns also need it on all-YES gadgets, whose number of satisfying assignments is shown in parentheses.

The values were checked by enumerating all 2^{2^d} tag sets; the numbers of affine subspaces in the included dimensions are $(4, 6)$, $(8, 28, 14)$, and $(16, 120, 140, 30)$. The optimized values come with rational primal and dual certificates, included with the verification code. None of the asymptotic results depends on these computations.

d	Witness counts	c , uniform	c , optimized	ρ , uniform	ρ , optimized
2	1, 2 (4)	5/8	3/5	1/2	1/2
3	1, 2, 4 (8)	19/42	96/229	76/195	72/191
4	1, 2, 4, 8 (16)	589/1680	219/674	2356/7501	280/929

Table 4: Exact regional thresholds for affine subspaces of dimensions $0, \dots, d-1$. “Uniform” means $a_s = 1/d$; each optimized column optimizes its own weights.

7 Sharpness and the limit of regional counting

The proof uses only how retained tag sets meet the designated regions. This section asks how far that statistic can go. First, simple tag filters attain the small-family thresholds on the gadget classes, so those thresholds cannot be improved for those families by a better analysis. Second, any family whose regions can be enumerated efficiently has $c(\mathcal{F}) = \Omega(1/\log L)$. The filters below are restricted to the gadget classes; they are not isolators for all affine circuits with the same number of satisfying assignments.

Proposition 7.1 (Tag filters). *Fix a family $\mathcal{F}$ with $d = O(\log L)$ and a polynomial-size sampler for a distribution μ on subsets of $\mathbb{F}_2^d$, given as truth tables. On input H , sample $T \leftarrow \mu$ and return $H \wedge [t \in T]$, where t is the first d variables (return H if it has fewer than d variables). This is a polynomial-size pruning transformation. On a one-NO gadget whose NO member is in position j , it succeeds with probability $\Pr_\mu[|T \cap K_j| = 1]$; on an all-YES gadget, with probability $\Pr_\mu[|T| = 1]$.*

Proof. Conjoining a predicate always prunes, and a truth table on d bits gives a circuit of size $O(d2^d) = \text{poly}(L)$. On the gadgets the first d variables are the tag, and by Lemma 3.3 the projection to the tag is a bijection from the satisfying set onto K_j or onto $\mathbb{F}_2^d$. The transformation does not need to recognize gadgets; only its success guarantee is restricted to them. $\square$

Corollary 7.2 (Matching tag filters on the gadget classes). *The following ideal distributions on tag sets attain the indicated thresholds.*

- (i) *For the all-lines family of Proposition 6.4: with probability ρ_N a uniform point, otherwise a uniform affine hyperplane. Success is ρ_N on every one-NO and every all-YES gadget.*
- (ii) *For the points-and-lines family of Proposition 6.3: with probability $3/5$ a uniform 2-subset, otherwise a uniform 3-subset. Success is $3/5$ on every one-NO gadget.*
- (iii) *For the same family: a uniform singleton or a uniform 3-subset, each with probability $1/2$. Success is exactly $1/2$ on every one-NO and every all-YES gadget.*

The distributions in (i) and (ii) can be implemented with fair coins to within any inverse-polynomial $\delta(L)$ of the stated value, at polynomial cost. The distribution in (iii) is implemented exactly. No guarantee on other affine circuits is claimed.

Proof. (i) A uniform affine hyperplane $\{\lambda(t) = \beta\}$ meets a line $\{a, b\}$ once exactly when $\lambda(a \oplus b) = 1$, which has probability $(N/2)/(N-1) = c'$. A hyperplane has $N/2 \geq 2$ points, so it never isolates an all-YES gadget. A uniform point lies on a fixed line with probability $2/N = g$ and always isolates an all-YES gadget. So the success on lines is $\rho_N g + (1 - \rho_N)c' = \rho_N$, by the definition of ρ_N , and the success on all-YES gadgets is ρ_N .

(ii) A fixed point lies in a uniform 2-subset with probability $1/2$ and in a uniform 3-subset with probability $3/4$. A fixed line meets them once with probabilities $2/3$ and $1/2$. Both mixtures equal $\frac{3}{5} \cdot \frac{1}{2} + \frac{2}{5} \cdot \frac{3}{4} = \frac{3}{5} \cdot \frac{2}{3} + \frac{2}{5} \cdot \frac{1}{2} = \frac{3}{5}$.

(iii) A point succeeds with probability $\frac{1}{2}(\frac{1}{4} + \frac{3}{4}) = \frac{1}{2}$ and a line with $\frac{1}{2}(\frac{1}{2} + \frac{1}{2}) = \frac{1}{2}$; the full tag set is isolated exactly on the singleton branch. One bit chooses the branch and two bits choose a tag, used either as the singleton or as the point omitted from the 3-subset.

For (i) and (ii), the probabilities are not dyadic. Approximate each finite uniform choice by rejection sampling: drawing a uniform integer from a finite interval succeeds with probability above $1/2$ per trial, so $O(\log(1/\delta))$ trials leave fallback mass at most δ . Total variation distance bounds the change in each success probability, and the transformation still prunes on every seed, including fallback seeds. Exact finite-coin probabilities are dyadic, so a value such as $3/5$ is only approached, not attained. $\square$

The disjoint-region threshold of Remark 6.2 is also attained. Choose a point z_i in each of m disjoint regions. With probability s keep a uniform one of these points, and otherwise keep all m . The all-YES success is s , and each regional singleton probability is $1 - s + s/m$. At $s = m/(2m - 1)$ the two agree, so this counting inequality cannot go below that value for disjoint families. Multiscale families avoid this example because their regions overlap and c is uniformly small.

Lemma 7.3 (A lower bound for the regional counting method). *For every nonempty family with $s_{\max} < d$,*

$$c(\mathcal{F}) \geq \frac{1}{4(s_{\max} + 1)}, \quad \rho(\mathcal{F}) \geq \frac{1}{4(s_{\max} + 2)}.$$

Proof. For the first bound, take $T = h^{-1}(0)$ for the random hash of Lemma 2.4 with i uniform in $\{0, \dots, s_{\max}\}$. For each region K_j , conditioning on $i = \dim K_j$ shows that $|T \cap K_j| = 1$ with probability at least $1/4$. So $\mathbb{E}f_{\mathcal{F}}(T) \geq 1/(4(s_{\max} + 1))$, and some T attains at least the expectation.

For the second, let i be uniform in $\{0, \dots, s_{\max}\} \cup \{d\}$, a set of $s_{\max} + 2$ values, and put $q = 1/(4(s_{\max} + 2))$. As before $\mathbb{E}f_{\mathcal{F}}(T) \geq q$, and the choice $i = d$ gives $a := \Pr[|T| = 1] \geq q$. The denominator $1 + c' - g$ is positive: every region is proper, so a pair of points, one inside and one outside a region, shows $c' > 0$, while $g \leq 1$. Suppose $\rho(\mathcal{F}) < q$. Then $g < q$ and $qg + (1 - q)c' < q$. Splitting the sampled sets into singletons and non-singletons,

$$\mathbb{E}f_{\mathcal{F}}(T) \leq ag + (1 - a)c' \leq \max\{g, qg + (1 - q)c'\} < q,$$

a contradiction. (The middle step uses $a \geq q$ and that $ag + (1 - a)c'$ is affine in a .) $\square$

The tester enumerates a region, so polynomial cost forces $s_{\max} = O(\log L)$, and the lemma then gives an $\Omega(1/\log L)$ lower bound on the thresholds this method can certify, either with no all-YES guarantee or with the same guarantee on both classes. Theorem 5.4 matches this order. This is a limit of the proof technique, not a lower bound on witness isolation itself. It also does not cover asymmetric guarantees with a much larger all-YES success probability.

Lemma 7.3 charges the cost of enumerating large regions. A second lower bound depends only on the number of regions, so it persists even if regions could be tested for free.

Remark 7.4 (A lower bound from the number of regions). Let $K_1, \dots, K_m$ be any nonempty subsets of $\mathbb{F}_2^d$, affine or not, define $c(\mathcal{F})$ as in Definition 3.1, and put $r = \lceil \log m \rceil$. Then

$$c(\mathcal{F}) \geq \frac{1}{8(r + 1)}.$$

To see this, choose one point $u_j \in K_j$ for each j , let U be the set of these points, so $|U| \leq m$, and sample $T \subseteq U$ by choosing i uniformly from $\{0, \dots, r\}$ and keeping each point of U independently with probability 2^{-i} . Fix j and let $A = K_j \cap U$, which is nonempty and satisfies $T \cap K_j = T \cap A$. If $|A| = 1$, the choice $i = 0$ keeps exactly one point of A . If $a = |A| \geq 2$, the choice $i = \lceil \log a \rceil \leq r$ gives a keep probability $\theta = 2^{-i} \in (1/(2a), 1/a]$, and

$$\Pr[|T \cap A| = 1 \mid i] = a\theta(1 - \theta)^{a-1} \geq \frac{1}{2}(1 - 1/a)^{a-1} > \frac{1}{2e} > \frac{1}{8}.$$

So $\Pr[|T \cap K_j| = 1] \geq 1/(8(r+1))$ for every j , and averaging over j gives the bound for some T . For polynomially many regions this is again $\Omega(1/\log L)$, now independent of $s_{\max}$. The distribution is also an efficient tag filter in the sense of Proposition 7.1, up to finite-coin rounding, so on the gadget classes it is an actual pruning transformation meeting every regional test at this level. The objective $|\{j : |T \cap K_j| = 1\}|$ is the unique coverage problem, where logarithmic behavior of this kind is known [15, 16]; the elementary argument is included for completeness.

Remark 7.5 (What this method does not settle). Affine hashing achieves $\Theta(1/n)$, while $p_{\text{wit}}(n) \geq 10/\log n$ implies the collapse. The intermediate scales, and the best constant in front of $1/\log L$, remain open. Improving this regional counting argument below $1/\log L$ would require giving up at least one of its ingredients. Remark 7.4 shows that giving up explicit enumeration of the tested region is not enough while the regions form a polynomial-size public family and the analysis rests on the scalar bound $c(\mathcal{F})$. Neither Lemma 7.3 nor Remark 7.4 is a lower bound against every reduction; they do not rule out other gadgets, statistics, or advice constructions.

8 Related work

DKMW. DKMW’s two-circuit comparison [2, Theorems 3.3–3.4] gives the collapse from success above $2/3$ on one- and two-point inputs. In the language of this paper it uses two disjoint point regions, and pools of $k+1$ circuits with pairwise disjoint regions extend the same count to $(k+1)/(2k+1)$, which tends to $1/2$ (Remark 6.2). The tagged construction does more than sharpen that count: it changes the geometry of the regions while keeping one canonical description for every designation. Theorem 5.2 covers every positive constant, and Theorem 5.4 gives an inverse-logarithmic threshold, both without cryptographic assumptions. Since the guarantees are only required on affine inputs, the implications also apply to procedures that isolate every satisfiable circuit. The separate question of isolation modulo two is not addressed.

Domination. The conversion from local comparisons to polynomial advice is established. Ko [7] introduced the domination method. Naik, Rogers, Royer, and Selman [8, Section 2, Lemma 4] use polynomial families of small advice sets with their membership witnesses, attribute the multiway domination lemma to Ogihara [9], and describe the greedy averaging-and-deletion step. We rederive this conversion for the tagged tester, including the sampling bijection of Lemma 4.5 and the resource bounds when the number of regions grows. The contribution here is the tagged affine gadget and the multiscale consequence, not a new domination theorem.

Membership comparability. Ogihara [10] studies algorithms that exclude one possible vector of membership bits, and Sivakumar [11] shows that logarithmic membership comparability of SAT gives a polynomial-time algorithm for the unique-versus-zero promise. Our argument constructs no comparator and uses neither result. In particular, it does not rely on the general polynomial-arity membership-comparability circuit claim whose retraction is reported by Beigel, Fortnow, and Pavan [12, Section 1.1]. All transfer and domination steps are proved here.

Cryptographic barriers. Under iO, the companion paper [14] shows that any affine-input guarantee $\omega(1/n)$ gives $2^{o(v)}$ -size circuits for 3SAT, so iO together with the nonuniform Exponential Time Hypothesis rules out every improvement over the hashing scale on affine inputs. The present results need no such assumption and stop at order $1/\log n$.

The affine-flag estimate. Lemma 5.1 is the argument recorded in [13]. Its new use is inside a syntax-dependent comparison: the argument never relates the output distributions of different descriptions of the same function, which is what distinguishes it from presentation-invariant isolation.

Acknowledgements

We are grateful to Dieter van Melkebeek for his generosity with his time, and for many helpful suggestions and discussions, which improved the presentation and pointed to further directions for strengthening the results. We also thank him for the paper with Dell, Kabanets, and Watanabe [2], which raised the question studied here and started this work.

Artificial intelligence

We acknowledge ChatGPT and Claude for their extensive support in editing, surveying prior literature, developing code to verify the calculations, and troubleshooting flaws in the gadget’s formulation

References

- [1] L. G. Valiant and V. V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47:85–93, 1986. doi:10.1016/0304-3975(86)90135-0.
- [2] H. Dell, V. Kabanets, D. van Melkebeek, and O. Watanabe. Is Valiant–Vazirani’s isolation probability improvable? *Computational Complexity*, 22(2):345–383, 2013. doi:10.1007/s00037-013-0059-7. Author-hosted full version.
- [3] S. Toda. PP is as hard as the polynomial-time hierarchy. *SIAM Journal on Computing*, 20(5):865–877, 1991.
- [4] K. Mulmuley, U. V. Vazirani, and V. V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [5] K. Reinhardt and E. Allender. Making nondeterminism unambiguous. *SIAM Journal on Computing*, 29(4):1118–1131, 2000.
- [6] C. Calabro, R. Impagliazzo, V. Kabanets, and R. Paturi. The complexity of Unique k -SAT: An isolation lemma for k -CNFs. *Journal of Computer and System Sciences*, 74(3):386–393, 2008.
- [7] K.-I. Ko. On self-reducibility and weak P-selectivity. *Journal of Computer and System Sciences*, 26(2):209–221, 1983. doi:10.1016/0022-0000(83)90013-2.
- [8] A. V. Naik, J. D. Rogers, J. S. Royer, and A. L. Selman. A hierarchy based on output multiplicity. *Theoretical Computer Science*, 207(1):131–157, 1998. doi:10.1016/S0304-3975(98)00060-7. Author-hosted full version, Section 2, Lemma 4 and the following witness-advice construction.
- [9] M. Ogihara. Functions computable with limited access to NP. *Information Processing Letters*, 58(1):35–38, 1996. doi:10.1016/0020-0190(96)00030-0. The multiway domination attribution used here is explicit in [8, Lemma 4].
- [10] M. Ogihara. Polynomial-time membership comparable sets. *SIAM Journal on Computing*, 24(5):1068–1081, 1995. doi:10.1137/S0097539793258131. The general polynomial-arity P/poly claim is not used; see the retraction report in [12, Section 1.1].
- [11] D. Sivakumar. On membership comparable sets. *Journal of Computer and System Sciences*, 59(2):270–280, 1999. doi:10.1006/jcss.1999.1653.
- [12] R. Beigel, L. Fortnow, and A. Pavan. *Membership Comparable and p -selective Sets*. Author-hosted manuscript, accessed 12 September 2026. Section 1.1 reports the retraction of the general polynomial-arity membership-comparability circuit claim, citing Ogihara’s personal communication as reference [Ogi05].
- [13] S. Ben Daniel. *Obfuscation and the Limits of Witness Isolation*. Electronic Colloquium on Computational Complexity (ECCC), Report TR26-201, 2026.
- [14] S. Ben Daniel. *Obfuscation and the Inverse-Linear Barrier for Witness Isolation: From Improved Isolation to Subexponential Circuits for 3SAT*. Manuscript, October 2026. To be posted to the Electronic Colloquium on Computational Complexity (ECCC).
- [15] E. D. Demaine, U. Feige, M. T. Hajiaghayi, and M. R. Salavatipour. Combination can be hard: Approximability of the unique coverage problem. *SIAM Journal on Computing*, 38(4):1464–1483, 2008. doi:10.1137/060656048.
- [16] V. Guruswami and E. Lee. Nearly optimal NP-hardness of unique coverage. *SIAM Journal on Computing*, 2017. doi:10.1137/16M1070682.
- [17] S. Ben Daniel. *Cryptographic Security and the Limits of Witness Isolation: General Pruning, Checkable Filters, and Oracle Barriers*. Companion report, September 2026.