

Poly-logarithmic Frege depth lower bounds via an expander switching lemma

Toniann Pitassi*
University of Toronto
toni@cs.toronto.edu

Benjamin Rossman†
National Institute of Informatics
rossman@nii.ac.jp

Rocco A. Servedio‡
Columbia University
rocco@cs.columbia.edu

Li-Yang Tan§
Toyota Technological Institute
liyang@cs.columbia.edu

April 28, 2016

Abstract

We show that any polynomial-size Frege refutation of a certain linear-size unsatisfiable 3-CNF formula over n variables must have depth $\Omega(\sqrt{\log n})$. This is an exponential improvement over the previous best results [PBI93, KPW95, BS02] which give $\Omega(\log \log n)$ lower bounds.

The 3-CNF formulas which we use to establish this result are Tseitin contradictions on 3-regular expander graphs. In more detail, our main result is a proof that for every d , any depth- d Frege refutation of the Tseitin contradiction over these n -node graphs must have size $n^{\Omega((\log n)/d^2)}$. A key ingredient of our approach is a new switching lemma for a carefully designed random restriction process over these expanders. These random restrictions reduce a Tseitin instance on a 3-regular n -node expander to a Tseitin instance on a random subgraph which is a topological embedding of a 3-regular n' -node expander, for some n' which is not too much less than n . Our result involves $\Omega(\sqrt{\log n})$ iterative applications of this type of random restriction.

*Supported by NSERC. Part of this research was done while visiting the Simons Institute.

†Supported by JST-ERATO Kawarabayashi Large Graph Project. This research was carried out at the Simons Institute and during a visit to the University of Toronto.

‡Supported by NSF grants CCF-1319788 and CCF-1420349. Part of this research was done while visiting the Simons Institute.

§Part of this research was done while at the Simons Institute and during a visit to Columbia University.

Contents

1	Introduction	1
2	Definitions	5
2.1	Frege Systems	5
2.2	Restrictions and Decision Trees	6
2.3	DNFs and Formulas	7
2.4	Tseitin contradictions	7
2.5	Expanders and some of their useful properties	8
2.6	Topological embeddings	9
3	Our Tseitin instances and the Kleinberg–Rubinfeld random restrictions	9
3.1	Our Tseitin instances	9
3.2	The “atomic” Kleinberg–Rubinfeld random restrictions	10
4	Composing the Kleinberg–Rubinfeld random restrictions	12
4.1	The inductive definition of $\mathcal{A}^{(i)}$	13
5	Independent sets, closures, and pruning	14
5.1	Independent sets	15
5.2	“Pushing the contradiction” and closure of restrictions	15
5.3	Pruning	17
6	k-Evaluations	21
7	Proof of Theorem 1: Obtaining a k-evaluation from the switching lemma	24
7.1	Proof of Lemma 7.1 assuming Theorem 2	24
7.2	Proof of Lemma 7.2	27
8	The “atomic” Tseitin Switching Lemma	29
8.1	A simple switching lemma for r -clipped decision trees	30
8.2	A quantity $r_{\mathcal{H}_{n'',n'}}(m, s)$ that we need to bound	33
8.3	The Tseitin switching lemma	38
8.4	Auxiliary lemma for the proof of Lemma 8.17	40
8.5	Proof of Lemma 8.17	42
9	Safe trees become total under KR random restrictions	45
10	Final Tseitin switching lemma	47
10.1	Useful tools for Theorem 7	48
10.2	Proof of Theorem 7	49
A	Proof of Lemma 2.7	52

1 Introduction

Background. One of the main motivations for the field of propositional proof complexity is the following observation of Cook and Reckhow [CR79]: If there is a propositional proof system in which any unsatisfiable formula F has a short proof of unsatisfiability (of size polynomial in the size of F), then $\text{NP} = \text{coNP}$. Despite great progress over the last 25 years, our understanding of the complexity of propositional proof systems is still very limited. In particular, lower bounds for propositional proof systems lag significantly behind lower bounds for the analogous circuit classes.

For some very simple proof systems, exponential lower bounds have long been known. In particular, the complexity of Resolution (where lines in a proof are depth-1 circuits) is now very well understood. An early groundbreaking paper by Tseitin [Tse68] introduced the Tseitin formulas and proved that any regular Resolution proof of the Tseitin formulas requires exponential size. Almost 20 years later, Haken [Hak85] proved that any (unrestricted) Resolution proof of the pigeonhole principle (PHP) requires exponential size. This led to a flurry of results in the late 80's and early 90's, establishing similar exponential size lower bounds for other families of tautologies [Urq87, CS88], a simpler proof of Haken's result [BKPS98], and finally Ben-Sasson and Wigderson's [BSW99] celebrated work on Resolution width, establishing that clause width is the fundamental parameter characterizing Resolution proof complexity.

A natural next class of proof systems to contend with is *constant-depth Frege systems*. These are “textbook-style” proof systems, where the lines in the proof are restricted to formulas computable in AC^0 . A major breakthrough was made by Ajtai in 1994 [Ajt94], who used an ingenious blend of nonstandard model theory and combinatorics to prove that any constant-depth Frege proof of PHP must have super-polynomial size. While Ajtai did not work out the exact parameters in his lower bound, [BPU92] gave a purely combinatorial reformulation of his result, and showed that his proof implies that any polynomial-size Frege proof of PHP must have depth $\Omega(\log^* n)$.

Subsequently, [PBI93, KPW95] significantly strengthened Ajtai's bound, showing that any polynomial-size Frege proof of PHP must have depth $\Omega(\log \log n)$. Urquhart and Fu [UF96] simplified and adapted the previous proofs in order to prove similar lower bounds for Tseitin formulas over the complete graph, and Ben-Sasson [BS02], via a very clever reduction, proved exponential lower bounds for the Tseitin formulas over constant-degree expander graphs. Unfortunately, all of these lower bounds become trivial once the depth of formulas in the proof are allowed to exceed $\log \log n$. On the other hand, Buss [Bus87] proved that PHP has polynomial-size, $O(\log n)$ -depth Frege proofs, and similar upper bounds are known for the Tseitin formulas over any graph. Because of this exponential gap between $\log \log n$ and $\log n$, it has been a fairly longstanding open problem to break the “ $\log \log n$ depth barrier” for Frege proofs; see e.g. the first open problem in [Raz02].

Challenges in proving lower bounds for larger depth Frege. The combinatorics underlying small-depth Frege lower bounds is similar to the methodology that has been used to prove the celebrated $2^{\Omega(n^{1/d})}$ lower bounds for depth- d circuits computing the Parity function [Hås86]. Those arguments apply a restriction (obtained by a switching lemma and a union bound) in order to shrink the depth of the circuit by one; performing this iteratively d times shrinks the circuit to an r -DNF formula for Parity over more than r variables, giving a direct contradiction. If each variable is independently set to 0 or 1 each with probability $(1-p)/2$, and is left unset with probability p , then Håstad's switching lemma [Hås86] shows that an r -DNF fails to convert to an s -CNF under such a restriction with probability at most $O(pr)^s$. Thus, by setting $r = s = O(\log n)$ and $p = \Theta(1/\log n)$, by a union bound there is a restriction which, applied to a polynomial-size depth- d

circuit, shrinks it to a polynomial-size depth- $(d-1)$ circuit, still computing Parity on the remaining $\Omega(n/\log n)$ unset variables. With this choice of p , we can apply such a restriction $\Theta(\log n/\log \log n)$ times, to obtain an $\Omega(\log n/\log \log n)$ depth lower bound for polynomial-size circuits that compute the n -variable Parity function.

However, proving lower bounds against small-depth Frege proofs is considerably more difficult, for two reasons. First, the proof complexity arguments that are required are very subtle — unlike in the case of the Parity lower bound, in the proof complexity context when we shrink the depth of formulas in the proof, we can no longer preserve equivalence in the usual sense. Instead, equivalence is preserved only in a “local” sense. Second, the random restrictions used in switching lemmas in the proof complexity context must “preserve the structure” of the statement being proved, which makes the switching lemmas more involved. For example, the pigeonhole principle from $n+1$ pigeons to n holes has variables $p_{i,j}$, and expresses the (unsatisfiable) fact that there is a one-to-one mapping from $n+1$ pigeons to n holes. In the case of the Parity lower bound, the random restrictions are unbiased and coordinate-wise independent, hence very simple. However in the case of the pigeonhole principle, we must be very careful to set the variables so as to not obviously falsify the pigeonhole principle: for example, we must never choose a restriction that maps two pigeons to the same hole (since the pigeonhole principle, under the restriction, should still be an instance of the pigeonhole principle, or at least an instance of something that is nontrivial to prove!). Thus we are forced to use restrictions where there are significant correlations between the variables: setting $p_{i,j}$ to 1 means that we must set $p_{i',j}$ to 0 for all other pigeons $i' \neq i$. This issue is not unique to the pigeonhole principle; regardless of what unsatisfiable CNF propositional formula we start with, we are constrained to choose restrictions that do not obviously falsify any of the clauses of the formula, for the same reason sketched above.

These constraints makes it quite difficult to prove a switching lemma with strong parameters. Prior to the current work, every switching lemma that has been proved for small-depth Frege is for a distribution over random restrictions in which p (the probability that a given variable is unset) is $O(n^{-\varepsilon})$ for some constant $\varepsilon > 0$; having such a small p leads to lower bounds that are super-polynomial only up to depth $O(\log \log n)$. (To see this, note for example that if $\varepsilon = 1/2$, then after d rounds we are left with only $n^{1/2^d}$ unset variables, and thus we cannot take $d > \log \log n$.)

Our contributions. In this paper, we prove the first super-polynomial lower bounds for Frege proofs of depth up to $\Theta(\sqrt{\log n})$, for Tseitin formulas over 3-regular expander graphs. Our main result is the following:

Theorem 1. *There is a linear-size 3-CNF contradiction $\text{Tseitin}(\mathcal{G}_n[\alpha])$ (an instance of Tseitin on a 3-regular n -node expander graph, to be defined formally later) with the following property: for every d , any depth- d Frege refutation of $\text{Tseitin}(\mathcal{G}_n[\alpha])$ must have size $n^{\Omega((\log n)/d^2)}$. Consequently, any polynomial-size Frege refutation of $\text{Tseitin}(\mathcal{G}_n[\alpha])$ must have depth $d = \Omega(\sqrt{\log n})$.*

The main novelty in our approach is in developing a way to start with a sparse graph \mathcal{G}_n over $O(n)$ variables/edges (a 3-regular n -node expander with $3n/2$ variables/edges), and apply a restriction that both

- leaves a large fraction (roughly an $n^{-1/(2d)}$ fraction) of the variables unset, yet still
- maintains that the Tseitin constraints, under the restriction, are “just as hard” as the original Tseitin constraints (though now with respect to a smaller universe).

Roughly speaking, the chief difficulty in executing this approach is in designing a distribution over random restrictions that are “highly structured” (so that the reduced formula is still very difficult to refute), yet at the same time “have enough randomness” so that we can prove a switching lemma with good parameters.

An expander switching lemma. An important enabling ingredient in designing a distribution as described above is a beautiful structural result about expander graphs due to Kleinberg and Rubinfeld [KR96]. This result states that any graph G' with $n/\text{polylog}(n)$ nodes and edges can be embedded in a bounded-degree n -node expander; in other words, such an expander contains G' as a minor. This easily implies that an n -node 3-regular expander contains any $(n/\text{polylog}(n))$ -node 3-regular expander as a *topological* minor. Inspired by and building on this structural result, at a high level, we start with the Tseitin principle over an n -node 3-regular expander graph \mathcal{G}_n . Applying a carefully designed random restriction (which corresponds closely to the randomized algorithm given in [KR96] to establish that the embedding exists), the 3-regular n -node expander \mathcal{G}_n “collapses” to a random subgraph \mathbf{H} that is a topological embedding of \mathcal{G}_n , a 3-regular n' -node expander, in \mathcal{G}_n , where $n' = n^{1-1/(2d)}$. Note that the corresponding reduced Tseitin formula under the restriction is not quite as clean as the Tseitin formula over a smaller 3-regular expander $\mathcal{G}_{n'}$; instead it consists of the Tseitin constraints over the graph \mathbf{H} which may be viewed as $\mathcal{G}_{n'}$ with each of its edges sub-divided into a path. However, intuitively, the reduced Tseitin formula is just as hard as the original one, since it is over a graph that contains a 3-regular expander as a topological minor. Since $n' = n^{1-1/(2d)}$, we can repeat this d times and still have $n^{1/2}$ variables survive all d repetitions. The parameters of our switching lemma are such that we can take d to be as large as $\Theta(\sqrt{\log n})$, and hence we get a depth lower bound of $\Omega(\sqrt{\log n})$.

As alluded to above, our random restrictions are over variables that correspond to edges of a 3-regular expander, and the switching lemma that we prove for them may be viewed as a “projection switching lemma over expander graphs”. The term “projection” here is in the sense of [RST15], and alludes to the aforementioned issue that \mathbf{H} is not a 3-regular expander but contains a 3-regular expander $\mathcal{G}_{n'}$ as a topological minor — more precisely, \mathbf{H} is obtained from $\mathcal{G}_{n'}$ by sub-dividing each of its edges into a path (so all the resulting paths are node-disjoint except at their endpoints). The “projection” essentially amounts to viewing each such path, joining two degree-3 nodes in \mathbf{H} , as a “super-edge” corresponding to a single new variable.

As suggested by the above discussion, the random restrictions that our switching lemma deals with are quite complex (see Sections 3 and 4 where these restrictions are formally defined). To handle this complexity the proof of our switching lemma follows a different line of argument than we are aware of in previous switching lemma proofs. Section 8.1 gives an overview of these new ideas in a significantly simplified setting (and also lays essential groundwork for our real switching lemma proof).

Overview of the proof complexity arguments. Above we have sketched some aspects of our switching lemma; here we provide some intuition for the proof complexity arguments that use this switching lemma. Assume that for n sufficiently large, we have a small AC^0 -Frege refutation P of $\text{Tseitin}(\mathcal{G}_n[\alpha])$, the Tseitin contradiction over n -node expander \mathcal{G}_n (where $\alpha \in \mathbb{Z}_2^{V(\mathcal{G}_n)}$ is an “odd charge” assignment to the vertices $V(\mathcal{G}_n)$ of \mathcal{G}_n). We apply a random restriction ρ (ρ is essentially the concatenation of d “atomic” random restrictions, each of which is used to reduce depth by one via our switching lemma) to all of the sub-formulas occurring in P . The restriction ρ leaves us with an instance $\text{Tseitin}(\mathbf{H}[\alpha \upharpoonright \rho])$ where \mathbf{H} is a randomly generated topological embedding of $\mathcal{G}_{n_{\text{final}}}$ in \mathcal{G}_n , where $n_{\text{final}} = n/(n^{1/(2d)})^d = \sqrt{n}$ and $\alpha \upharpoonright \rho$ is an odd charge over $V(\mathbf{H})$ induced by ρ .

We apply the random restriction ρ to each formula in the Frege proof P to obtain a new proof $P \upharpoonright \rho$. (The restriction ρ applied to P simply substitutes in every formula of P the constants 0 or 1 for the variables that have been fixed by ρ .) Next, we argue that there is an outcome ρ of ρ such that $P' := P \upharpoonright \rho$ has a “ k -evaluation” for $k \ll n_{\text{final}}$. For every sub-formula A that occurs in P' , the k -evaluation assigns to A an associated decision tree of height at most k . These trees are constructed via an iterative process that involves d rounds, where each round both applies a random restriction and performs a simplification. (The analysis establishing the existence of a k -evaluation uses our switching lemma along with a union bound over the size of P to argue that all the formulas do indeed simplify.) These trees have some important properties. First, every clause of $\text{Tseitin}(H[\alpha \upharpoonright \rho])$ will be associated with a 1-tree (a decision tree with all leaves labelled by 1). Second, the final formula “0” (which is in P' since it is the final formula in P) will be associated with a 0-tree (all leaves labelled by 0). Finally, the trees associated with all formulas and their sub-formulas satisfy a certain type of “local consistency.” The existence of a k -evaluation for all sub-formulas in P' leads to a contradiction since it implies on the one hand that all formulas in P' are associated with 1-trees, but on the other hand, that the final “0” formula in P' is associated with a 0-tree.

Related Results and Discussion. As mentioned above, Ajtai [Ajt94] proved the first super-polynomial lower bound for bounded-depth Frege systems, establishing that depth $O(\log^* n)$ Frege proofs of PHP require super-polynomial size. This was improved by [PBI93, KPW95] who proved $\exp(\Omega(n^{5-d}))$ size lower bounds against depth- d Frege, thus establishing super-polynomial lower bounds up to depth $\Theta(\log \log n)$. Ben-Sasson [BS02] proved comparable lower bounds for certain Tseitin contradictions with 3-CNF encodings. All of these prior super-polynomial lower bounds hold for depths up to $\Theta(\log \log n)$, and thus our lower bound gives an exponential improvement in terms of depth. However, for constant-depth Frege the previous results give exponential size bounds, whereas our proof only gives quasi-polynomial size bounds. Thus in terms of size, previous results are better than ours for depths at most $\Theta(\log \log n)$. We conjecture that the optimal size lower bound for depth- d Frege is $\exp(\Omega(n^{1/d}))$ (which would match the state-of-the-art in Boolean circuit lower bounds [Hås86]) and we conjecture that this may be achievable by improving the analysis of our switching lemma (for the same random restrictions).

We view our result as a significant step towards establishing small-depth Frege lower bounds for randomly generated CNF formulas. Random CNF formulas are one of the most well-studied families of SAT formulas, both because of their strong connection to threshold phenomena in statistical physics, and because they are a canonical family of hard examples for SAT solvers. Proof complexity lower bounds for random 3-CNF formulas are important, as they demonstrate the limitations of the proof system (and corresponding limitations on SAT algorithms based on the proof system) for a large family of formulas, and not just for specially tailored hard instances. Super-polynomial lower bounds for random CNF formulas have been known for over ten years for Resolution [CS88] and Polynomial Calculus [BI99]. However, for small-depth Frege systems, progress has been quite slow. The best results to date are super-polynomial lower bounds for random CNFs in $\text{Res}(q)$, a weak generalization of Resolution where proof lines are disjunctions of terms of width $q \leq \sqrt{\log n / \log \log n}$ [SBI04, Ale11]. Thus it is even open to prove lower bounds for random CNFs for depth-2 Frege. To see the connection with Tseitin formulas, we first observe that Tseitin formulas over random graphs are a variant of random k -XOR formulas (subject to the additional constraint that each variable occurs in exactly two equations), and second, that Frege lower bounds for random k -XOR formulas imply Frege lower bounds for random k -CNF formulas

[BI99]. We are optimistic that our techniques can be extended to the case of random k -XOR and thus to random k -CNF.

Organization. Section 2 defines Frege proof systems, Tseitin formulas, restrictions, expanders, and other basic notions. In Section 3 we define our “atomic” random restrictions, and in Section 4 we explain how these atomic random restrictions are composed to obtain the actual random restrictions that we use. Section 5 establishes notions of independence (of sets of edges), closure (of restrictions), and pruning (of decision trees) which play an important role in our technical results. Section 6 introduces the key notion of a “ k -evaluation” from proof complexity, and Section 7 uses this notion to prove our main theorem assuming our Tseitin Switching Lemma. Finally, Sections 8 through 10 contains the proof of our switching lemma.

2 Definitions

2.1 Frege Systems

The underlying Frege system that we will use here is Shoenfield’s system, as presented in [UF96]. Because any two bounded-depth Frege systems over \wedge , \vee and \neg can polynomially simulate one another [CR79], our results hold more generally for any bounded-depth Frege system over this basis.

Our proof system uses binary disjunction \vee and \neg ; a conjunction $A \wedge B$ is treated as an abbreviation for the formula $\neg(\neg A \vee \neg B)$. In addition, we include the propositional constants 0 and 1, representing “false” and “true” respectively. If A, B_1, \dots, B_m are formulas over a sequence of variables p_1, \dots, p_m , then $A[B_1/p_1, \dots, B_m/p_m]$ is the formula resulting from A by substituting B_1, \dots, B_m for p_1, \dots, p_m .

A *rule* is defined to be a sequence of formulas, written $A_1, \dots, A_k \rightarrow A_0$. In the case where A_1, \dots, A_k is empty, the rule is referred to as an *axiom scheme*. The rule is *sound* if every truth assignment satisfying all of A_1, \dots, A_k also satisfies A_0 . If $A_1, \dots, A_k \rightarrow A_0$ is a Frege rule, then C_0 is *inferred from* C_1, \dots, C_k by this rule if there is a sequence of formulas B_1, \dots, B_m and variables p_1, \dots, p_m so that for all $i, 0 \leq i \leq k$, $C_i = A_i[B_1/p_1, \dots, B_m/p_m]$. (In other words $C_1, \dots, C_k \rightarrow C_0$ is a substitution instance of $A_1, \dots, A_k \rightarrow A_0$.)

Shoenfield’s system \mathcal{F} consists of the following rules:

- (Excluded Middle): $(p \vee \neg p)$;
- (Expansion rule): $p \rightarrow q \vee p$;
- (Contraction rule): $p \vee p \rightarrow p$;
- (Associative rule): $p \vee (q \vee r) \rightarrow (p \vee q) \vee r$;
- (Cut rule): $p \vee q, \neg p \vee r \rightarrow q \vee r$.

Let $A = C_1 \wedge \dots \wedge C_m$ be an unsatisfiable CNF formula. A *refutation* of A in \mathcal{F} is a finite sequence of formulas such that every formula in the sequence is one of C_1, \dots, C_m or inferred from earlier formulas in the sequence by a rule in \mathcal{F} , and the last formula is 0. \mathcal{F} is a sound and complete proof system: all of the rules are sound and thus if A has a refutation then A is unsatisfiable, and furthermore \mathcal{F} is complete since every unsatisfiable CNF formula has a refutation in \mathcal{F} .

We will be working with the above proof system which operates with *binary* connectives; however we want to measure formula depth using *unbounded* fan-in connectives. To do this, we will measure the depth of a formula by the number of alternations of \vee and \neg . More precisely, we can represent a formula by a tree in which each leaf is labeled with a propositional variable or a constant, and each interior node is labeled with \vee if it is the parent of two nodes or \neg if it is the parent of only one. A branch in the tree, when traversed from the root to the leaf, is labeled with a block of operators of one kind (say \neg) followed by a block of the other kind (say \vee), and so on, ending with a variable or constant. The *logical depth* of a branch is defined to be the number of blocks of operators labelling the branch. The *depth* of a formula is the maximum logical depth of the branches in its formation tree. This notion of depth is the same as the depth of the corresponding unbounded fan-in \wedge, \vee, \neg formula, up to small constant factors. A refutation in \mathcal{F} has depth d if the maximum depth of any formula in the proof is at most d . The *size* of a formula is the number of (binary) connectives in the formula. The size of a refutation Γ , denoted $|\Gamma|$, is the sum of the sizes of all formulas in the refutation.

2.2 Restrictions and Decision Trees

Given a set of variables E , a *restriction* β is an element of $(\mathbb{Z}_2 \cup \{*\})^E$. We may equivalently view β as an element of \mathbb{Z}_2^S for some $S \subseteq E$, and we say that S is the *support* of β , denoted $\text{supp}(\beta)$. Given restrictions $\beta, \beta' \in (\mathbb{Z}_2 \cup \{*\})^E$ we say that β' is a *sub-restriction* of β if $\text{supp}(\beta') \subseteq \text{supp}(\beta)$ and $\beta(e) = b \in \mathbb{Z}_2$ whenever $\beta'(e) = b \in \mathbb{Z}_2$. We say that β is an *extension* of β' if β' is a sub-restriction of β . We say that a set of restrictions $\{\beta_1, \dots, \beta_t\}$ is *mutually compatible* if for every $i, j \in [t]$ we have that $\beta_i(e) = \beta_j(e)$ whenever $e \in \text{supp}(\beta_i) \cap \text{supp}(\beta_j)$.

A decision tree is said to be *proper* if no variable is queried two or more times on any branch. Throughout the entire paper all decision trees are assumed to be proper unless explicitly stated otherwise.

Let T be a decision tree over a space of variables E and let β be a restriction. We write “ $T \upharpoonright \beta$ ” to denote the decision tree obtained by simplifying T according to β “in the obvious way.” More precisely, if $T = (e; T_0, T_1)$, meaning that T is the decision tree with variable e at the root and left (respectively, right) child T_0 (respectively, T_1), then we have the following:

- If $\beta(e) = *$ then $T \upharpoonright \beta = (e; T_0 \upharpoonright \beta; T_1 \upharpoonright \beta)$;
- If $\beta(e) = 0$ then $T \upharpoonright \beta = T_0 \upharpoonright \beta$;
- If $\beta(e) = 1$ then $T \upharpoonright \beta = T_1 \upharpoonright \beta$.

For $b \in \mathbb{Z}_2$, we write “ $T \upharpoonright \beta = b$ ” to indicate that $T \upharpoonright \beta$ is the 1-node tree consisting only of a leaf b .

A *branch* of a decision tree T over variable set E is a root-to-leaf path (which, more formally, is a sequence $\pi = \pi_1 \circ \dots \circ \pi_k$ where each π_i is an element of $E \times \mathbb{Z}_2$ and no element occurs twice in π). Clearly every branch in a decision tree T corresponds to a unique restriction and we will often take this perspective of viewing branches as restrictions.

Unless specified otherwise a decision tree is assumed to have every leaf labeled with an element of \mathbb{Z}_2 ; we sometimes refer to such trees as *total* to emphasize that every leaf is labeled with an element of \mathbb{Z}_2 . A *partial decision tree* is a decision tree in which each leaf is labeled with some element of $\mathbb{Z}_2 \cup \{\perp\}$. A partial decision tree is said to be *m-safe* if all leaves labeled \perp are at depth

$\geq m$. Given a total decision tree T , we write T^c to denote the tree obtained from T by replacing each leaf bit with its complement.

For $b \in (\mathbb{Z}_2 \cup \{\perp\})$ and a decision tree T let $\text{Branches}_b(T)$ denote the set of branches of T whose leaf is labeled b . Let $\text{Branches}(T) = \text{Branches}_0(T) \cup \text{Branches}_1(T)$, which is the set of all branches of T if T is total. Note that if $T \upharpoonright \beta = b$ for some $b \in \mathbb{Z}_2$ (i.e. $T \upharpoonright \beta$ is the constant- b tree of depth 0), then there exists a path $\pi \in \text{Branches}_b(T)$ such that β extends π .

A 1-tree is a decision tree in which every leaf is labeled by 1, and likewise 0-tree.

We close this subsection with a notational convention: many of our definitions, lemmas, etc. will involve restrictions. As a helpful notational convention, in such definitions, lemmas, etc. we use $\rho, \rho', \tilde{\rho}$, etc. to denote restrictions which should be thought of as “coming from” a random restriction process, and we use $\pi, \pi', \tilde{\pi}$, etc. to denote restrictions which should be thought of as “coming from” a branch in some decision tree. If the definition/lemma/etc. may arise in either context we use $\beta, \beta', \tilde{\beta}$, etc.

2.3 DNFs and Formulas

Observe that any element of $\text{Branches}_1(T)$, for any decision tree T , corresponds to a conjunction in an obvious way, so we may view $\text{Branches}_1(T)$ as a collection of conjunctions (i.e. of terms). Given a decision tree T we write “ $\text{Disj}(T)$ ” to denote the DNF whose terms correspond precisely to the 1-branches of T . Observe that if T_1, \dots are decision trees then “ $\bigvee_j \text{Disj}(T_j)$ ” is a DNF (whose terms are precisely the terms that occur in some $\text{Disj}(T_j)$).

For a term t , and a restriction β , $t \upharpoonright \beta$ is 0 if any literal in t is set to 0 by β , otherwise, $t \upharpoonright \beta$ is the conjunction of literals that are unset by β . (By definition, a term of size zero is equal to 1.) For an DNF formula $F = t_1 \vee \dots \vee t_m$, $F \upharpoonright \beta$ is the DNF formula obtained as follows: If any term is set to 1 by β , then $F \upharpoonright \beta = 1$, and otherwise $F \upharpoonright \beta = \bigvee_i (t_i \upharpoonright \beta)$. (In particular, $F \upharpoonright \beta = 0$ means that every term in F is set to 0 by β .)

For a Boolean formula A and a restriction β , $A \upharpoonright \beta$ denotes the formula obtained by performing simple variable substitution as dictated by β . If Γ is a collection of formulas $\Gamma = \{A_i\}$ then $\Gamma \upharpoonright \beta$ denotes $\{A_i \upharpoonright \beta\}$.

2.4 Tseitin contradictions

Given a labeling $\alpha \in \mathbb{Z}_2^V$ of the vertices of a graph $G = (V, E)$, we define a Boolean formula associated with (G, α) . We view each element of E as a formal Boolean variable. The Tseitin formula $\text{Tseitin}(G[\alpha])$ over the formal variables in E is defined to be

$$\text{Tseitin}(G[\alpha]) := \bigwedge_{v \in V} \underbrace{\left(\sum_{e \sim v} e = \alpha(v) \right)}_{\text{constraint}(v, \alpha)},$$

where “ $e \sim v$ ” indicates that edge e has vertex v as an endpoint. Note that if G has maximum degree 3, then for every $v \in V$, $\text{constraint}(v, \alpha)$ can be expressed as a 3-CNF with $2^{3-1} = 4$ clauses. Therefore for such an n -node graph, the formula $\text{Tseitin}(G[\alpha])$ is a 3-CNF with $4n$ clauses.

We sometimes refer to α as the “charge” of G , and to the value $\alpha(v)$ as being “odd” or “even.” The following fact is well known (see e.g. Lemma 18.16 of [Juk12]):

Fact 2.1. *If G is connected, then $\text{Tseitin}(G[\alpha])$ is satisfiable iff $\sum_{v \in V} \alpha(v) = 0$.*

For G a graph and α a charge we say that α is an *odd charge* if $\sum_{v \in V} \alpha(v) = 1$ and that it is an *even charge* otherwise. By Fact 2.1, if $\alpha \in \mathbb{Z}_2^V$ is any odd charge then the associated 3-CNF is unsatisfiable. We sometimes also consider charges $\alpha \in \mathbb{Z}_2^{V'}$ for some $V' \subseteq V$, as in the following definition and facts. (Here and throughout the paper, for a graph $G = (V, E)$ and $S \subseteq E$, we write $G - S$ to denote the graph $(V, E \setminus S)$.)

Definition 2.2 (α -consistency). *Let $G = (V, E)$ be a graph and $\alpha \in \mathbb{Z}_2^{V'}$ where $V' \subseteq V$. Let $S \subseteq E$ be a subset of the edges and $\rho \in \mathbb{Z}_2^S$. We say that ρ is α -consistent if*

$$\sum_{e \sim v} \rho(e) = \alpha(v) \quad \text{for all } v \in V' \text{ that are isolated in } G - S.$$

The following fact is an easy consequence of Fact 2.1:

Fact 2.3. *Let $G = (V, E)$ be a connected graph and $\alpha \in \mathbb{Z}_2^{V'}$ where $V' \subsetneq V$. Then there is an α -consistent assignment $\rho \in \mathbb{Z}_2^{E(G)}$.*

Fact 2.4. *Let $G = (V, E)$ be a connected graph and $\alpha \in \mathbb{Z}_2^{V'}$ where $V' \subsetneq V$. Then the set of all α -consistent assignments $\rho \in \mathbb{Z}_2^{E(G)}$ forms an affine subspace of $\mathbb{Z}_2^{E(G)}$.*

Proof. The set of α -consistent assignments is easily seen to be the set of solutions to a system of linear equations; this system is satisfiable (by Fact 2.3), and therefore the set of solutions forms an affine subspace. \square

The following notion of restricting (“toggling”) a charge by a partial assignment to the edges will be useful for us:

Definition 2.5 (Restricting a charge). *Let $G = (V, E)$ be a graph and $\alpha \in \mathbb{Z}_2^V$. Let $S \subseteq E$ be a subset of the edges and $\rho \in \mathbb{Z}_2^S$. We define $(\alpha \upharpoonright \rho) \in \mathbb{Z}_2^V$ by*

$$(\alpha \upharpoonright \rho)(v) = \alpha(v) - \sum_{e \sim v} \mathbf{1}[e \in S \ \& \ \rho(e) = 1] \quad \text{for all vertices } v \in V.$$

We say that a connected component $C = (V(C), E(C))$ of an n -node graph is *giant* if $V(C) > n/2$. Note that there can only be one giant component in any given graph. For $\alpha \in \mathbb{Z}_2^V$ a charge of G and $C = (V(C), E(C))$ a connected component of G , we say the α -charge of C is $\sum_{v \in C} \alpha(v)$.

We mention that sometimes we will write things like $H[\alpha]$ when $\alpha \in \mathbb{Z}_2^{V(G)}$ and H is a subgraph of G where $V(H)$ is a proper subset of $V(G)$; in such a usage “ α ” should be interpreted as α restricted to $V(H)$.

2.5 Expanders and some of their useful properties

Definition 2.6 (Expansion). *Fix $\gamma > 0$. A simple undirected graph $G = (V, E)$ is said to be a γ -expander if for every set $X \subset V$ of at most half the vertices, we have $|\delta(X)| \geq \gamma \cdot |X|$, where $\delta(X)$ denotes the set of edges that have one end in X and the other end in $V \setminus X$.*

It is well known that there exists an infinite family of graphs $\{\mathcal{G}_n = (V(\mathcal{G}_n), E(\mathcal{G}_n))\}_{\text{even } n \geq N_0}$ such that each \mathcal{G}_n is a 3-regular γ -expander for some absolute constant $\gamma > 0$. We fix such a family; these are the graphs we shall work with for the remainder of the paper.

Let $S \subseteq E(\mathcal{G}_n)$ be a subset of the edges of \mathcal{G}_n . As will be clear in Section 5.2, we require the fact that if S is not too large, then $\mathcal{G}_n - S$ does not contain too many bridges – in particular, the number of bridges should not be too much more than $|S|$. While it should be possible to show that $\mathcal{G}_n - S$ has at most $O(|S|)$ bridges, a weaker bound suffices for our purposes. The following lemma, based on natural intuitions about expander graphs, is proved in Appendix A.

Lemma 2.7. *For every $S \subseteq E(\mathcal{G}_n)$, the graph $\mathcal{G}_n - S$ contains at most $C_1 \cdot |S| \cdot \log^2 n$ bridges, where C_1 is a constant depending only on the expansion parameter γ .*

An easy corollary is the following:

Corollary 2.8. *Fix $S \subseteq E(\mathcal{G}_n)$ and let B denote the set of all bridges in $\mathcal{G}_n - S$. Then $\mathcal{G}_n - (B \cup S)$ has a connected component of size at least $n - \lambda|S|$ where $\lambda = O(\log^2 n)$.*

Proof. By Lemma 2.7 we have $|B \cup S| \leq (1 + C_1 \log^2 n)|S|$. The corollary now follows from the well known fact that removing t edges from an n -node γ -expander results in a graph with a connected component of size at least $n - C_\gamma t$ where C_γ is a constant depending only on γ (see e.g. Exercise 12 of [Tao11] or Lemma 1.2 of [Tre11]). \square

2.6 Topological embeddings

An important notion for our approach is that of the topological embedding of one graph in another. We briefly recall the definition of topological embedding:

Definition 2.9 (Topological embedding). *Let $G = (V(G), E(G))$, $G' = (V(G'), E(G'))$ and $H = (V(H), E(H))$ be three graphs. We say that H is a topological embedding of G' in G if H is a subgraph of G and*

1. *There is a one-to-one map $\phi : V(G') \rightarrow V(H)$,*
2. *There is a one-to-one map $\varphi : E(G') \rightarrow \{\text{simple paths in } H\}$ such that for all $(u, v) \in E(G')$, $\varphi((u, v))$ is a simple path from $\phi(u)$ to $\phi(v)$ in H and no two paths $\varphi(e_1)$ and $\varphi(e_2)$ share any non-endpoint vertices,*
3. *Every $e \in E(H)$ is contained in $\varphi(e')$ for a unique $e' \in E(G')$.*

We refer to $\phi(V(G')) \subseteq V(H)$ as the *real vertices* of H , and the remaining vertices $V(H) \setminus \phi(V(G'))$ as the *path vertices*. For every $e \in E(H)$, we write $\text{super}_H(e)$ to denote the unique path $\varphi(e')$ containing e , and refer to it as the *super-edge of H containing e* .

A useful intuitive view of (1)–(3) above is that by subdividing edges of G' it is possible to obtain an isomorphic copy of H .

3 Our Tseitin instances and the Kleinberg–Rubinfeld random restrictions

3.1 Our Tseitin instances

Fix $\mathcal{G}_n = \{(V(\mathcal{G}_n), E(\mathcal{G}_n))\}_{\text{even } n \geq N_0}$ to be any particular fixed 3-regular γ -expander family. For each even n we view the vertex set $V(\mathcal{G}_n)$ of \mathcal{G}_n as $[n] = \{1, \dots, n\}$; since \mathcal{G}_n is 3-regular we have that $|E(\mathcal{G}_n)| = 3n/2$.

Throughout this paper, we reserve $\alpha \in \mathbb{Z}_2^{V(\mathcal{G}_n)}$ to denote an odd charge over $V(\mathcal{G}_n)$. Our hard instances will be $\text{Tseitin}(\mathcal{G}_n[\alpha])$.

Notational convention. We reserve d^* to denote the value $c\sqrt{\log n}$, the largest depth bound for which we will establish a result (here c is a small absolute constant). Let $\tau(\cdot, \cdot)$ denote the function $\tau(i, n) = n^{1-i/(2d^*)}$. For brevity we will sometimes write n' where the intended meaning is $\tau(i, n)$ for some $i \in [0, d^* - 1]$, and we will sometimes write n'' where the intended meaning is $\tau(i + 1, n)$. Note that we always have $n^{1/2} \leq n'' < n' \leq n$ throughout the paper.

3.2 The “atomic” Kleinberg–Rubinfeld random restrictions

Let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. In this section we define a distribution $\mathcal{F}_{n'', n'}$ over pairs $(\boldsymbol{\rho}, \mathbf{H})$ where $\boldsymbol{\rho} \in (\mathbb{Z}_2 \cup \{*\})^{E(\mathcal{G}_{n'})}$ and $\mathbf{H} = (V(\mathbf{H}), E(\mathbf{H}))$ where \mathbf{H} is a topological embedding of $\mathcal{G}_{n''}$ in $\mathcal{G}_{n'}$. Furthermore $E(\mathbf{H}) = \boldsymbol{\rho}^{-1}(*)$ and $\boldsymbol{\rho}$ is α' -consistent.

The definition of $\mathcal{F}_{n'', n'}$ is as follows. A draw from $\mathcal{F}_{n'', n'}$ is obtained in two stages:

1. First draw $\mathbf{H} \sim \mathcal{H}_{n'', n'}$. Here $\mathcal{H}_{n'', n'}$ is a distribution over subgraphs of $\mathcal{G}_{n'}$ that have $\mathcal{G}_{n''}$ as a topological minor. We describe this distribution $\mathcal{H}_{n'', n'}$ in detail in Section 3.2.1 below.
2. Then draw an assignment $\boldsymbol{\rho} \in (\mathbb{Z}_2 \cup \{*\})^{E(\mathcal{G}_{n'})}$ as follows:
 - $\boldsymbol{\rho}(e) = *$ for all $e \in E(\mathbf{H})$, and $\boldsymbol{\rho}(e') \in \mathbb{Z}_2$ for all $e' \in E(\mathcal{G}_{n'}) \setminus E(\mathbf{H})$.
 - Viewing $\text{supp}(\boldsymbol{\rho})$ as an element of $\mathbb{Z}_2^{E(\mathcal{G}_{n'}) \setminus E(\mathbf{H})}$, $\boldsymbol{\rho}$ is drawn by selecting a uniform random element of the set of all α' -consistent assignments in $\mathbb{Z}_2^{E(\mathcal{G}_{n'}) \setminus E(\mathbf{H})}$.

The output of the draw is $(\boldsymbol{\rho}, \mathbf{H})$. Observe that \mathbf{H} is connected and $\boldsymbol{\rho}$ is α' -consistent, hence the $(\alpha' \upharpoonright \boldsymbol{\rho})$ -charge of \mathbf{H} is odd and $\text{Tseitin}(\mathbf{H}[\alpha' \upharpoonright \boldsymbol{\rho}])$ is unsatisfiable (just like $\text{Tseitin}(\mathcal{G}_{n'}[\alpha'])$).

3.2.1 The distribution $\mathcal{H}_{n'', n'}$ over random topological embeddings

Let $n'', n' \leq n$ be as described at the end of Section 3.1 and let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. We define a distribution $\mathcal{H}'_{n'', n'}$ over connected subgraphs of $\mathcal{G}_{n'}$. We remark that the distribution $\mathcal{H}'_{n'', n'}$ given below is closely derived from a probabilistic argument (given in [KR96]; see also [BFU94]) establishing the existence of a $\mathcal{G}_{n''}$ -minor in the n' -node expander $\mathcal{G}_{n'}$. (Looking ahead, as we discuss in detail below $\mathcal{H}'_{n'', n'}$ should be viewed as an auxiliary distribution for our ultimate purposes; given a draw of \mathbf{H}' from $\mathcal{H}'_{n'', n'}$, some additional massaging will yield a subgraph $\mathbf{H} \subset \mathbf{H}'$ which has $\mathcal{G}_{n''}$ as a topological minor. Later our description of the distribution $\mathcal{H}_{n'', n'}$ from which \mathbf{H} is drawn will build on $\mathcal{H}'_{n'', n'}$.) A draw of $\mathbf{H}' \sim \mathcal{H}'_{n'', n'}$ is obtained as follows:

1. (**Choose “special” vertices.**) For $i = 1, \dots, n''$ do the following:
 - Pick a uniform random vertex $\mathbf{v}_i \in [n']$ conditioned on \mathbf{v}_i having distance at least $3\kappa_1 \ln \ln n'$ in $\mathcal{G}_{n'}$ from all of $\mathbf{v}_1, \dots, \mathbf{v}_{i-1}$. (Here κ_1 should be viewed as a large constant.)

We say that the elements of the set $\{\mathbf{v}_1, \dots, \mathbf{v}_{n''}\}$ are the *special vertices* of \mathbf{H}' (\mathbf{H}' also contains other vertices, described below).

2. (**Choose “bundles” linking up special vertices.**) Independently for each edge $\{a_{i,1}, a_{i,2}\} \in E(\mathcal{G}_{n''}) \subset \binom{[n'']}{2}$, $i = 1, \dots, 3n''/2$, construct a *bundle* $\mathbf{B}(i)$ of $r = (\ln n')^2$ paths $\mathbf{P}_{i,1}, \dots, \mathbf{P}_{i,r}$ connecting $\mathbf{v}_{a_{i,1}}$ and $\mathbf{v}_{a_{i,2}}$ as follows. The j -th path $\mathbf{P}_{i,j}$ in the bundle is constructed as follows (independently for all j):

- (a) Let $\tau = \kappa_2 \ln n'$.
- (b) Choose a midpoint $\mathbf{x}_{i,j} \in [n']$ uniformly at random (note that the uniform distribution over $[n']$ is the stationary distribution over $V(\mathcal{G}_{n'})$ since $\mathcal{G}_{n'}$ is 3-regular). Choose a random walk $\mathbf{W}'_{i,j}$ from all length- τ walks that start at $\mathbf{v}_{a_{i,1}}$ and end at $\mathbf{x}_{i,j}$. And choose a random walk $\mathbf{W}''_{i,j}$ from all length- τ walks that start at $\mathbf{v}_{a_{i,2}}$ and end at $\mathbf{x}_{i,j}$. We define the walk $\mathbf{W}_{i,j}$ to be the concatenation of $\mathbf{W}'_{i,j}$ and the reversal of $\mathbf{W}''_{i,j}$, and we define the path $\mathbf{P}_{i,j}$ to be $\mathbf{W}_{i,j}$ with any cycles removed.

Note that each path in each bundle both starts and ends at a special vertex. The vertex set $V(\mathbf{H}')$ is the set of all vertices that occur anywhere on any path in any bundle (hence it includes all the special vertices). The edge set $E(\mathbf{H}')$ is the union of all the edges in all the paths in all the bundles.

Actually, it will be convenient for us to view a draw of $\mathbf{H}' \sim \mathcal{H}'_{n'',n'}$ as being not just the graph described above, but as the entire transcript of the random draw described above (so a draw of \mathbf{H}' includes the structure of the bundles, walks and paths; we will use this structure below).

As stated earlier, the distribution $\mathcal{H}'_{n'',n'}$ defined above is an auxiliary distribution; we are really interested in a different distribution $\mathcal{H}_{n'',n'}$ over connected subgraphs H of $\mathcal{G}_{n'}$ which we now describe. Informally, every draw of $\mathbf{H} \sim \mathcal{H}_{n'',n'}$ is a graph with n'' distinguished nodes (we call them *real* nodes) that have degree three; all other nodes in \mathbf{H} have degree two and lie on vertex-disjoint paths between the real nodes (we call these degree-2 nodes *path* nodes). The edges of \mathbf{H} are the edges constituting these paths. If each path is contracted to a single edge joining the two real nodes at its endpoints, we obtain a graph isomorphic to $\mathcal{G}_{n''}$. In other words, for every \mathbf{H} in the support of $\mathcal{H}_{n'',n'}$, the graph $\mathcal{G}_{n''}$ is a topological minor of \mathbf{H} .

A draw of $\mathbf{H} \sim \mathcal{H}_{n'',n'}$ is obtained as follows:

- 1'. Draw $\mathbf{H}' \sim \mathcal{H}'_{n'',n'}$. This defines special vertices and bundles, walks and paths joining special vertices as described above.
- 2'. As described in [KR96], a sequence of steps “pruning” each bundle down to a single path results in a graph $\tilde{\mathbf{H}}$ (a subgraph of \mathbf{H}' and hence of $\mathcal{G}_{n'}$) with the following property: for each $\{a_{i,1}, a_{i,2}\} \in E(\mathcal{G}_{n''})$, $i \in [3n''/2]$, $\tilde{\mathbf{H}}$ contains a simple path which we denote $\mathbf{Z}_{a_{i,1},a_{i,2}}$ joining $v_{a_{i,1}}$ and $v_{a_{i,2}}$, and every edge in $\tilde{\mathbf{H}}$ belongs to some such path (possibly to more than one path). Moreover, this collection of paths has the following property: If two of these paths intersect at any node then (i) they must share an endpoint, and (ii) they can only intersect within their first $\kappa_1 \ln \ln n$ steps from the common endpoint.

(We omit the details of this pruning process, as they are involved and are not important to us. We only mention that it is a deterministic process that discards paths from bundles but does not alter paths; each path $\mathbf{Z}_{a_{i,1},a_{i,2}}$ is some path $\mathbf{P}_{i,j}$ in the bundle $\mathbf{B}(i)$ in H' , and no bundle contributes more than one such path to $\tilde{\mathbf{H}}$. There is an $o(1)$ failure probability for the process (over the draw of $\mathbf{H}' \sim \mathcal{H}'_{n'',n'}$ in step 1 above). If this failure event occurs we go back and repeat step 1 until this step does not fail.)

3'. Given that the failure event does not occur, the analysis of [KR96] establishes that there is an embedding of $\mathcal{G}_{n''}$ in $\tilde{\mathbf{H}}$ in which the i -th vertex of $\mathcal{G}_{n''}$ corresponds to the union $\mathbf{Z}_{i,j_1}^{\text{init}} \cup \mathbf{Z}_{i,j_2}^{\text{init}} \cup \mathbf{Z}_{i,j_3}^{\text{init}}$, where $\mathbf{Z}_{i,j}^{\text{init}}$ is the set of the first $\kappa_1 \ln \ln n$ edges of the path $\mathbf{Z}_{i,j}$ starting from \mathbf{v}_i , and $\{i, j_1\}, \{i, j_2\}, \{i, j_3\}$ are the three edges incident to node i in $\mathcal{G}_{n''}$ (and the $\{i, j\}$ edge in $\mathcal{G}_{n''}$ corresponds to the sub-path $\mathbf{Z}_{i,j} \setminus (\mathbf{Z}_{i,j}^{\text{init}} \cup \mathbf{Z}_{j,i}^{\text{init}})$; the analysis of [KR96] establishes that all these $3n''/2$ sub-paths are pairwise vertex disjoint.) Hence $\mathcal{G}_{n''}$ is a minor of $\tilde{\mathbf{H}}$; now, since $\mathcal{G}_{n''}$ is 3-regular and is a minor of $\tilde{\mathbf{H}}$, it must in fact be a *topological* minor of $\tilde{\mathbf{H}}$ (this is a standard fact, see e.g. Proposition 1.7.4(ii) of [Die10]). This implies that there is a subgraph \mathbf{H} of $\tilde{\mathbf{H}}$ as described earlier. More precisely \mathbf{H} has n'' “real” nodes $\mathbf{u}_1, \dots, \mathbf{u}_{n''}$ of degree 3 (where each \mathbf{u}_i is at distance at most $\kappa_1 \ln \ln n$ from \mathbf{v}_i), and for each edge $\{a_{i,1}, a_{i,2}\} \in E(\mathcal{G}_{n''}) \subset \binom{[n'']}{2}$, $i = 1, \dots, 3n''/2$, \mathbf{H} has a simple path connecting $\mathbf{v}_{a_{i,1}}$ and $\mathbf{v}_{a_{i,2}}$; all these paths are node-disjoint except at the endpoints and hence any non-endpoint node occurring on any such path has degree 2 in \mathbf{H} . Such vertices of degree 2 are called “path” nodes.

This completes our description of the draw of $\mathbf{H} \sim \mathcal{H}_{n'',n'}$. We sometimes write “ $(\mathbf{H}, \mathbf{H}') \sim \mathcal{H}_{n'',n'}$ ” where \mathbf{H}' is the draw from $\mathcal{H}_{n'',n'}$ obtained in Step 1' that is subsequently massaged to yield \mathbf{H} .

Some terminology will be useful going forward: for a given edge $\{a_{i,1}, a_{i,2}\} \in E(\mathcal{G}_{n''}) \subset \binom{[n'']}{2}$, $i \in [3n''/2]$ and an $H \in \text{supp}(\mathcal{H}_{n'',n'})$, we refer to the set of edges comprising the simple path in H connecting $v_{a_{i,1}}$ and $v_{a_{i,2}}$ as the $\{a_{i,1}, a_{i,2}\}$ -*super-edge* of H .

4 Composing the Kleinberg–Rubinfeld random restrictions

In this section we use the “atomic” Kleinberg–Rubinfeld random restrictions (described in Section 3.2) to define a sequence of distributions $\mathcal{A}^{(0)}, \mathcal{A}^{(1)}, \dots, \mathcal{A}^{(d)}$ (for any $d \leq d^*$) of random restrictions over the variable set $E(\mathcal{G}_n)$, the $(i+1)$ -st of which may be viewed as an extension of the i -th. As sketched in the Introduction, our first random restriction applies the atomic Kleinberg–Rubinfeld random restriction to “collapse” the n -node expander \mathcal{G}_n to a random subgraph $\mathbf{H}^{(1)}$ that is a topological embedding of $\mathcal{G}_{n'}$ in \mathcal{G}_n where $n' = n^{1-1/(2d)}$: this graph $\mathbf{H}^{(1)}$ has n' many degree-3 nodes (each of which corresponds to a vertex in $\mathcal{G}_{n'}$) that are joined by simple paths (each of which corresponds to an edge in $\mathcal{G}_{n'}$). Conceptually, our second random restriction “contracts” these paths in $\mathbf{H}^{(1)}$ to convert it into $\mathcal{G}_{n''}$, and applies another atomic Kleinberg–Rubinfeld random restriction to $\mathcal{G}_{n''}$ to obtain a random subgraph $\mathbf{H}^{(2)}$ of $\mathbf{H}^{(1)}$, a topological embedding of $\mathcal{G}_{n''}$ in $\mathbf{H}^{(1)}$ (and hence also in our original expander \mathcal{G}_n) where $n'' = n'/n^{1/(2d)} = n^{1-2/(2d)}$. Our successive random restrictions continue in this fashion.

The i -th distribution of random restrictions over $E(\mathcal{G}_n)$ that we consider is denoted $\mathcal{A}^{(i)}$. (We emphasize that these restrictions are always applied over edge set $E(\mathcal{G}_n)$ and not over $E(\mathcal{G}_{n'})$ for smaller values $n' < n$.) A draw from $\mathcal{A}^{(i)}$ yields a pair $(\boldsymbol{\rho}^{(i)}, \mathbf{H}^{(i)})$ with the following properties:

- A. $\boldsymbol{\rho}^{(i)} \in (\mathbb{Z}_2 \cup \{*\})^{E(\mathcal{G}_n)}$ is α -consistent,
- B. $\mathbf{H}^{(i)} = (V(\mathbf{H}^{(i)}), E(\mathbf{H}^{(i)}))$ is a random subgraph of \mathcal{G}_n which is a topological embedding of $\mathcal{G}_{\tau(i,n)}$ in \mathcal{G}_n (recall the $\tau(\cdot, \cdot)$ function defined in Section 3.1);
- C. The edge set $E(\mathbf{H}^{(i)})$ is precisely $(\boldsymbol{\rho}^{(i)})^{-1}(\ast)$.

Observe that since $\mathbf{H}^{(i)}$ is connected and $\boldsymbol{\rho}^{(i)}$ is α -consistent, we have that the induced $(\alpha \upharpoonright \boldsymbol{\rho}^{(i)})$ -charge of $\mathbf{H}^{(i)}$ is odd and hence $\text{Tseitin}(\mathbf{H}^{(i)}[\alpha \upharpoonright \boldsymbol{\rho}^{(i)}])$ is unsatisfiable (just like $\text{Tseitin}(\mathcal{G}_n[\alpha])$). Notationally we write $(\boldsymbol{\rho}^{(i)}, H^{(i)})$ (without boldface) to denote an element of $\text{supp}(\mathcal{A}^{(i)})$.

4.1 The inductive definition of $\mathcal{A}^{(i)}$

We now give a precise definition of our sequence $\{\mathcal{A}^{(i)}\}_{i \in \{0, \dots, d\}}$ of random restrictions, from which it will be clear it has the properties (A)–(C) above. $\mathcal{A}_n^{(0)}$ is supported on the single element $(\boldsymbol{\rho}^{(0)} = \{*\}^{E(\mathcal{G}_n)}, H^{(0)} = \mathcal{G}_n)$. Given $\mathcal{A}^{(i)}$ we define $\mathcal{A}^{(i+1)}$ in a sequence of steps as follows. Let $n' = \tau(i, n)$ and $n'' = \tau(i+1, n)$ (recall the definition of $\tau(\cdot, \cdot)$ in Section 3.1).

Fix an element $(\boldsymbol{\rho}^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$. We first define an auxiliary distribution $\mathcal{A}^{(i+1)}(\boldsymbol{\rho}^{(i)}, H^{(i)})$, where a draw $(\boldsymbol{\rho}^{(i+1)}, \mathbf{H}^{(i+1)}) \sim \mathcal{A}^{(i+1)}(\boldsymbol{\rho}^{(i)}, H^{(i)})$ will be such that $\boldsymbol{\rho}^{(i+1)}$ extends $\boldsymbol{\rho}^{(i)}$, and $\mathbf{H}^{(i+1)}$ is a topological embedding of $\mathcal{G}_{n''}$ in $H^{(i)}$ (and hence in \mathcal{G}_n). Recall that $\text{Tseitin}(H^{(i)}[\alpha \upharpoonright \boldsymbol{\rho}^{(i)}])$ is unsatisfiable. Roughly speaking, we first “contract” $\text{Tseitin}(H^{(i)}[\alpha \upharpoonright \boldsymbol{\rho}^{(i)}])$ into a corresponding Tseitin instance $\text{Tseitin}(\mathcal{G}_{n'}[\alpha'])$. This charge $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ induced by $(\boldsymbol{\rho}^{(i)}, H^{(i)})$ will have the following crucial property: since the initial charge $\alpha \in \mathbb{Z}_2^{V(\mathcal{G}_n)}$ is odd and $\boldsymbol{\rho}^{(i)}$ is α -consistent, α' will also be odd (and hence $\text{Tseitin}(\mathcal{G}_{n'}[\alpha'])$ is unsatisfiable). We will then use a draw $(\boldsymbol{\rho}, \mathbf{H}) \sim \mathcal{F}_{n'', n'}$ — an “atomic” Kleinberg–Rubinfeld random restriction that “collapses” $\text{Tseitin}(\mathcal{G}_{n'}[\alpha'])$ into $\text{Tseitin}(\mathbf{H}[\alpha' \upharpoonright \boldsymbol{\rho}])$ — to define our next random restriction $(\boldsymbol{\rho}^{(i+1)}, \mathbf{H}^{(i+1)}) \sim \mathcal{A}^{(i+1)}(\boldsymbol{\rho}^{(i)}, H^{(i)})$ that “collapses” $\text{Tseitin}(H^{(i)}[\alpha \upharpoonright \boldsymbol{\rho}^{(i)}])$ into $\text{Tseitin}(\mathbf{H}^{(i+1)}[\alpha \upharpoonright \boldsymbol{\rho}^{(i+1)}])$.

Recall that $H^{(i)}$ is a topological embedding of $\mathcal{G}_{n'}$ in \mathcal{G}_n , and hence there are maps $\phi : V(\mathcal{G}_{n'}) \rightarrow V(H^{(i)})$ and $\varphi : E(\mathcal{G}_{n'}) \rightarrow \{\text{simple paths in } H^{(i)}\}$ satisfying properties (1)–(3) of Definition 2.9. For each super-edge $\varphi(e)$ in $H^{(i)}$, we fix a canonical edge $e^* \in \varphi(e)$, i.e. $\text{super}_{H^{(i)}}(e^*) = \varphi(e)$ (any choice is fine).

Definition of $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$. Fix $v \in V(\mathcal{G}_{n'})$, and suppose v is incident to edges $e_{v,1}, e_{v,2}, e_{v,3} \in E(\mathcal{G}_{n'})$. For each $j \in [3]$,

- Let $e_{\phi(v),j}$ denote the canonical edge of the super-edge $\varphi(e_{v,j})$. That is, $\{e_{\phi(v),j}\}_{j=1,2,3}$ are the three canonical edges of the three super-edges $\{\varphi(e_{v,j})\}_{j=1,2,3}$ that are incident to $\phi(v)$ in $H^{(i)}$.
- Let $e'_{\phi(v),j}$ denote the edge in $\varphi(e_{v,j})$ that is incident to $\phi(v)$. Note that $e_{\phi(v),j}$ and $e'_{\phi(v),j}$ lie on the same super-edge $\varphi(e_{v,j})$ (they may even be the same edge).
- For each assignment $b(e_{\phi(v),j}) \in \mathbb{Z}_2$ to $e_{\phi(v),j}$, there is a unique extension $\mathbb{Z}_2^{\varphi(e_{v,j})}$ that satisfies all $(\alpha \upharpoonright \boldsymbol{\rho}^{(i)})$ -constraints of path vertices in $\varphi(e_{v,j})$. Let $b(e'_{\phi(v),j}) \in \mathbb{Z}_2$ denote the assignment to $e'_{\phi(v),j}$ under this unique extension.

There is a unique “offset bit” $\omega_v \in \mathbb{Z}_2$ such that for every assignment $(b(e_{\phi(v),j}))_{j=1,2,3} \in \mathbb{Z}_2^3$, we have

$$\omega_v = \sum_{j=1}^3 (b(e_{\phi(v),j}) + b(e'_{\phi(v),j})).$$

We define $\alpha'(v) := \alpha(\phi(v)) + \omega_v$. We have the following key property of $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$: for every assignment $(b(e_{v,j}))_{j=1,2,3} \in \mathbb{Z}_2^3$, if $\{e_{\phi(v),j}\}_{j=1,2,3}$ are set the same way (i.e. $b(e_{\phi(v),j}) = b(e_{v,j})$),

then the induced assignment $(b(e'_{v,j}))_{j=1,2,3} \in \mathbb{Z}_2^3$ satisfies the following:

$$b(e_{v,1}) + b(e_{v,2}) + b(e_{v,3}) = \alpha(v) \iff b(e'_{\phi(v),1}) + b(e'_{\phi(v),2}) + b(e'_{\phi(v),3}) = \alpha'(\phi(v)).$$

This concludes our description of the charge $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$.

As above we continue to consider a fixed element $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$. We are now ready to define the auxiliary distribution $\mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)})$:

Definition 4.1. Fix $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$. We define the distribution over random restrictions $\mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)})$ as follows. A draw $(\boldsymbol{\rho}^{(i+1)}, \mathbf{H}^{(i+1)}) \sim \mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)})$ is obtained by

- Draw $(\boldsymbol{\rho}, \mathbf{H})$ from $\mathcal{F}_{n'', n'}$ where the underlying charge on $V(\mathcal{G}_{n'})$ is the just-described α' . Recall that $\boldsymbol{\rho} \in (\mathbb{Z}_2 \cup \{*\})^{E(\mathcal{G}_{n'})}$ and that \mathbf{H} is a topological embedding of $\mathcal{G}_{n''}$ in $\mathcal{G}_{n'}$.
- For every $e \in \mathbf{E}(\mathcal{G}_n)$, if $\rho^{(i)}(e) \in \mathbb{Z}_2$ then $\boldsymbol{\rho}^{(i+1)}(e) = \rho^{(i)}(e)$ (so $\boldsymbol{\rho}^{(i+1)}$ extends $\rho^{(i)}$).
- Now consider an edge $e \in \boldsymbol{\rho}^{-1}(\mathbb{Z}_2) \subseteq E(\mathcal{G}_{n'})$ and its super-edge $\varphi(e)$ in $H^{(i)}$. Let e^* be the canonical edge in $\varphi(e)$. Set $\boldsymbol{\rho}^{(i+1)}(e^*) = \boldsymbol{\rho}(e)$, and extend $\boldsymbol{\rho}^{(i+1)}$ to fix all edges in $\varphi(e)$ in the unique way that satisfies all $(\alpha \upharpoonright \boldsymbol{\rho}^{(i)})$ -constraints of path vertices in $\varphi(e)$. We do this for every $e \in \boldsymbol{\rho}^{-1}(\mathbb{Z}_2)$.
- Let $\mathbf{H}^{(i+1)} = (\boldsymbol{\rho}^{(i+1)})^{-1}(\ast)$. Note that $\mathbf{H}^{(i+1)}$ is indeed a topological embedding of \mathbf{H} in $H^{(i)}$ (and hence $\mathbf{H}^{(i+1)}$ is a topological embedding of $\mathcal{G}_{n''}$ in \mathcal{G}_n).

Given Definition 4.1, at last we can define $\mathcal{A}^{(i+1)}$ quite simply. A draw from $\mathcal{A}^{(i+1)}$ is obtained as follows: Draw $(\boldsymbol{\rho}^{(i)}, \mathbf{H}^{(i)}) \sim \mathcal{A}^{(i)}$, and output a draw from $\mathcal{A}^{(i+1)}(\boldsymbol{\rho}^{(i)}, \mathbf{H}^{(i)})$. It is clear from the preceding discussion that indeed Properties (A)–(C) from above are indeed satisfied.

5 Independent sets, closures, and pruning

For technical reasons, we will be especially interested in graphs G and charges $\alpha \in \mathbb{Z}_2^{V(G)}$ that satisfy the following:

Definition 5.1 (Nice graphs and charges). Let G be a graph and $\alpha \in \mathbb{Z}_2^{V(G)}$ be an odd charge. We say that the pair (G, α) is nice if

1. G has a giant component, and
2. The α -charge of a component C of G is odd iff C is giant.

Note that the pair (\mathcal{G}_n, α) from $\text{Tseitin}(\mathcal{G}_n[\alpha])$ is (trivially) nice since \mathcal{G}_n is connected and α is odd. Intuitively, if (G, α) is nice then “the contradiction of $\text{Tseitin}(G[\alpha])$ is trapped in the giant component of G ” in the following sense: any assignment $x \in \mathbb{Z}_2^{E(G)}$ has to violate an α -charge constraint somewhere in the giant component of G .

In the following subsections, we introduce notions of “ G -independence” and “ G -closure”. For readers who are well-versed in matroid theory, we note that these are the standard notions of closure and independence in what is commonly known as the bond matroid (or co-graphic bond matroid) of the graph G .

5.1 Independent sets

We will work extensively with special kinds of sets of edges that we refer to as *independent*. Very roughly and intuitively, an independent set of edges (i.e. variables) is one in which no variable's assignment is “forced” given an assignment to the other variables in the set (we will make this much more precise later).

Definition 5.2 (Independent sets and decision trees). *Let $G = (V, E)$ be a connected graph. A set $I \subseteq E$ is G -independent if $G - I$ is connected. A G -independent decision tree T is a decision tree querying edges in E such that every branch of T queries a G -independent set.*

If G is connected and $\alpha \in \mathbb{Z}_2^{V(G)}$ is odd (hence (G, α) is nice), it is easy to see that an independent set $I \subseteq E$ is one for which for *any* restriction $\rho \in \mathbb{Z}_2^I$, the pair $(G - I, \alpha \upharpoonright \rho)$ is nice.

At various points in the argument it will be important for the decision trees we work with to be of bounded depth, hence the following definition:

Definition 5.3 (Good trees). *Let $G = (V, E)$ be a connected graph and T be a decision tree over E . We say that T is (k, G) -good if T is total, has depth $< k$, and is G -independent.*

To motivate the following definition, we may think of the assignment to a bridge variable in $G - S$ as being “forced” by an assignment to the variables in S (this too will be made much more precise later):

Definition 5.4 (Closure of set). *Fix a graph $G = (V, E)$ and a subset $S \subseteq E$. The G -closure of S , denoted $\text{closure}_G(S)$, is the set $S \cup B$ where B is the set of all bridges in $G - S$. We say that S is G -closed if $S = \text{closure}_G(S)$ (i.e. $G - S$ is bridgeless).*

With this terminology in place we may restate Lemma 2.7 and Corollary 2.8, about our 3-regular expanders \mathcal{G}_n , as follows, recalling that $\lambda = O(\log^2 n)$. (These will be used to satisfy the conditions of Proposition 5.12 below.)

Fact 5.5. *For every $S \subseteq E(\mathcal{G}_n)$ we have $\text{closure}_{\mathcal{G}_n}(S) \leq \lambda|S|$.*

Fact 5.6. *For every $S \subseteq E(\mathcal{G}_n)$, the number of vertices disconnected from the largest component in $\mathcal{G}_n - \text{closure}_{\mathcal{G}_n}(S)$ is at most $\lambda|S|$.*

Define λ' to be $(3\kappa_2 \log n)^{d^*+1}$ for the constant κ_2 from Step 2(a) of Section 3.2.1 (the definition of the distribution $\mathcal{H}_{n', n'}$, so every super-edge in any $H^{(i)}$ graph is always a path in \mathcal{G}_n of length at most λ'). Let $\lambda^* = 3\lambda'/2$. The following analogue of Fact 5.6 holds for $H^{(i)}$ graphs and is an immediate consequence of Fact 5.6:

Fact 5.7. *For any $0 \leq i \leq d^*$ and any $H^{(i)}$ such that $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$ for some $\rho^{(i)}$, and any $S \subseteq E(H^{(i)})$, the number of vertices of $H^{(i)}$ that are disconnected from the largest component in $H^{(i)} - \text{closure}_{H^{(i)}}(S)$ is at most $\lambda^*|S|$.*

5.2 “Pushing the contradiction” and closure of restrictions

Recall that a connected component C in an n -node graph G is said to be *giant* if $|V(C)| > n/2$.

Given a nice pair (G, α) , we will be interested in restrictions π such that $(G - \text{supp}(\pi), \alpha \upharpoonright \pi)$ is nice:

Definition 5.8 (Push the contradiction). *Let (G, α) be nice. We say that a restriction $\pi \in (\mathbb{Z}_2 \cup \{*\})^{E(G)}$ pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$ if $(G - \text{supp}(\pi), \alpha \upharpoonright \pi)$ is nice.*

Note that any π which pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$ must be α -consistent, because every isolated vertex in $G - \text{supp}(\pi)$ is a non-giant component. As remarked after the definition of independent sets, if G is connected and α is odd (hence (G, α) is nice), then *every* restriction π to an independent set $I \subseteq E$ pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$. (By the two preceding observations, any restriction π to an independent set is α -consistent.) We will use the following fact extensively:

Fact 5.9. *Let (G, α) be nice.*

- *If π is a restriction that pushes the contradiction of α into the giant component of $G - \text{supp}(\alpha)$, then for all sub-restrictions π' of π , we have that π' pushes the contradiction of α into the giant component of $G - \text{supp}(\pi')$.*
- *If π is a restriction that does not push the contradiction of α into the giant component of $G - \text{supp}(\alpha)$, then for all extensions π'' of π , we have that π'' does not push the contradiction of α into the giant component of $G - \text{supp}(\pi'')$.*

Let (G, α) be nice. If e is not a bridge in G , it is straightforward to see that $(G - \{e\}, \alpha \upharpoonright e \rightarrow b)$ is nice for both values $b \in \mathbb{Z}_2$. If e is a bridge and $G - \{e\}$ has a giant component, then the next fact implies that there is a unique assignment $b \in \mathbb{Z}_2$ such that $(G - \{e\}, \alpha \upharpoonright e \rightarrow b)$ is nice. (If $G - \{e\}$ does not have a giant component then clearly $(G - \{e\}, \alpha \upharpoonright e \rightarrow b)$ is not nice for either $b \in \mathbb{Z}_2$.)

Fact 5.10 (Bridges are forced). *Let $G = (V, E)$ be a graph, $\alpha \in \mathbb{Z}_2^V$, and $e \in E$ be a bridge in G . Let C be the component in G that contains e , and C_1 and C_2 be the two connected components of $C - \{e\}$. Then for every $b' \in \mathbb{Z}_2$ there exist a unique $b \in \mathbb{Z}_2$ such that*

$$\sum_{v \in V(C_1)} (\alpha \upharpoonright (e \rightarrow b)) = b'.$$

In particular, if the α -charge of C is odd then there is a unique assignment b to e so that the $(\alpha \upharpoonright e \rightarrow b)$ -charge (the “induced charge”) of C_1 is odd (and hence the induced charge of C_2 is even). Likewise, if the α -charge of C is even then there is a unique assignment b to e so that the induced charges of both C_1 and C_2 are even.

The following notion is crucial for us:

Definition 5.11 (Closure of restriction). *Let $G = (V, E)$ and $\alpha \in \mathbb{Z}_2^V$ be an odd charge. Let $\pi \in (\mathbb{Z}_2 \cup \{*\})^E$ be a restriction such that $G - \text{closure}_G(\text{supp}(\pi))$ has a giant component. A (G, α) -closure of π , denoted $\text{closure}_{G, \alpha}(\pi) \in \mathbb{Z}_2^{\text{closure}_G(\text{supp}(\pi))}$, is the unique extension of π with the following properties: Fix any $e \in \text{closure}_G(\text{supp}(\pi)) \setminus \text{supp}(\pi)$, and recall that e is a bridge in $G - \text{supp}(\pi)$. Let C be the component of $G - \text{supp}(\pi)$ that contains e , and let C_1 and C_2 be the two disjoint components of $C - \{e\}$ where $|V(C_1)| \geq |V(C_2)|$. (In the following, recall from Fact 5.10 that there is indeed a unique assignment as claimed in (1) and (2) below.)*

- If the $(\alpha \upharpoonright \pi)$ -charge of C is even, then $(\text{closure}_{G,\alpha}(\pi))(e) = b$ where $b \in \mathbb{Z}_2$ is the unique assignment such that the $((\alpha \upharpoonright \pi) \upharpoonright (e \rightarrow b))$ -charges of both C_1 and C_2 are even.
- If the $(\alpha \upharpoonright \pi)$ -charges of C is odd, then $(\text{closure}_{G,\alpha}(\pi))(e) = b$ where $b \in \mathbb{Z}_2$ is the unique assignment such that the $((\alpha \upharpoonright \pi) \upharpoonright (e \rightarrow b))$ -charge in C_1 is odd and C_2 is even.

It is easy to see from the above definition that for any $G = (V, E)$, any charge $\alpha \in \mathbb{Z}_2^V$, and any restriction $\pi \in (\mathbb{Z}_2 \cup \{*\})^E$, there exists a unique (G, α) -closure of π . We have the following crucially useful property of closure:

Proposition 5.12 (Key property of closure). *Let (G, α) be nice. Let $\pi \in (\mathbb{Z}_2 \cup \{*\})^{E(G)}$ be a restriction that pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$, and suppose $G - \text{closure}_{G,\alpha}(\pi)$ has a giant component. Then $\text{closure}_{G,\alpha}(\pi) \in \mathbb{Z}_2^{\text{closure}_G(\text{supp}(\pi))}$ pushes the contradiction of α into the giant component of $G - \text{closure}_G(\text{supp}(\pi))$.*

Proof. The proof is by considering a forest T whose vertices are the connected components of $G - \text{closure}_G(\text{supp}(\pi))$ and whose edges are the elements of $\text{closure}_G(\text{supp}(\pi)) \setminus \text{supp}(\pi)$ (i.e. the bridges in $G - \text{supp}(\pi)$). We view the tree with the giant component as rooted at the giant component (other trees may be rooted at an arbitrary vertex). The definition of $\text{closure}_{G,\alpha}(\pi)$, applied repeatedly to all leaf edges, gives that each leaf vertex (which is a non-giant component) has even charge under $\alpha \upharpoonright \pi$. Repeatedly applying the definition of $\text{closure}_{G,\alpha}(\pi)$ to tree edges “working up from the leaves,” we infer that each non-giant component has even charge under $\alpha \upharpoonright \pi$. Finally, since each component other than the giant component has even charge and the total charge is odd, the charge of the giant component under $\alpha \upharpoonright \pi$ is odd, and we have satisfied the definition of “ $\text{closure}_{G,\alpha}(\pi) \in \mathbb{Z}_2^{\text{closure}_G(\text{supp}(\pi))}$ pushes the contradiction of α into the giant component of $G - \text{closure}_G(\text{supp}(\pi))$.” \square

5.3 Pruning

Recall that a decision tree is good if it is total, has height at most k , and is independent. In the previous section, we have argued that if a tree is good, then for any branch π in the tree the closure of π pushes the contradiction into the giant component (and thus leaves us with an induced instance of Tseitin that is still hard).

Given a graph \tilde{G} , let us consider what happens to a good tree \tilde{T} over $E(\tilde{G})$ when we apply a restriction $\rho \in (\mathbb{Z}_2 \cup \{*\})^{E(\tilde{G})}$ to it. After applying the restriction ρ , the relevant graph is now $G := \tilde{G} - \text{supp}(\rho)$; for this intuitive explanation, one should think of G as connected, as indeed will be the case in our setting. First, note that there may be branches π in $T := \tilde{T} \upharpoonright \rho$ that are not G -independent; to obtain a tree that is G -independent, these branches need to be “pruned”. Additionally, there may be branches π' in T that do not push the contradiction of $\alpha \upharpoonright \rho$ into the giant component of $G - \text{supp}(\pi')$; these branches need to be “marked as invalid”. With the above as motivation, below we describe our deterministic *pruning* process for decision trees.

5.3.1 Definition of the pruning process

The input to our process is a total decision tree T over the edge set $E(G)$ of an n -node graph $G = (V(G), E(G))$, where T is not necessarily G -independent. Let α be an odd charge $\alpha \in \mathbb{Z}_2^{V(G)}$. The pruned variant of T is the decision tree $\text{Prune}_{G,\alpha}(T)$ defined as follows.

$\text{Prune}_{G,\alpha}(T) :$

1. If $T = b$ for some $b \in \mathbb{Z}_2$ then output b if G has a giant component and \perp otherwise.
2. If $T = (e; T_0, T_1)$, then
 - If e is not a bridge in G , output the tree $(e; \text{Prune}_{G-\{e\},\alpha|(e \rightarrow 0)}(T_0), \text{Prune}_{G-\{e\},\alpha|(e \rightarrow 1)}(T_1))$.
 - If e is a bridge of G , let C be the connected component in G that e belongs to, and C_1 and C_2 be the two disjoint components of $C - \{e\}$ where $|V(C_1)| \geq |V(C_2)|$. (In the following, recall from Fact 5.10 that there is indeed a unique assignment as claimed in (1) and (2) below.)
 - (a) If the α -charge of C is even, then shortcut e according to the unique assignment b to e such that the $(\alpha \upharpoonright (e \rightarrow b))$ -charge of both C_1 and C_2 are even, i.e. output $\text{Prune}_{G-\{e\},\alpha|(e \rightarrow b)}(T_b)$.
 - (b) If the α -charge of C is odd and C_1 is the giant component of $G - \{e\}$, then shortcut e according to the unique assignment b to e such that the $(\alpha \upharpoonright (e \rightarrow b))$ -charge of C_1 is odd, i.e. output $\text{Prune}_{G-\{e\},\alpha|(e \rightarrow b)}(T_b)$.
 - (c) If the α -charge of C is odd and C_1 is not a giant component of $G - \{e\}$ (so both C_1 and C_2 are non-giant components), output \perp .

The following facts are self-evident:

Fact 5.13. *For every connected graph G and any T we have that $\text{Prune}_{G,\alpha}(T)$ is G -independent. Moreover, if T is G -independent then $\text{Prune}_{G,\alpha}(T) = T$.*

Fact 5.14. *Let (G, α) be nice and T be a tree over $E(G)$. Then for all $b \in \mathbb{Z}_2$ and $\pi \in \text{Branches}_b(\text{Prune}_{G,\alpha}(T))$, the graph $G - \text{supp}(\pi)$ has a giant component, and moreover, π pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$.*

Fact 5.15. *Let (G, α) be nice. Let T be a decision tree over $E(G)$, and suppose both T and $\text{Prune}_{G,\alpha}(T)$ are total. Then for all restrictions $\beta \in (\mathbb{Z}_2 \cup \{*\})^{E(G)}$ and $b \in \mathbb{Z}_2$, if $T \upharpoonright \beta = b$ then additionally $\text{Prune}_{G,\alpha}(T) \upharpoonright \beta = b$.*

The next two lemmas show that our pruning process and definition of the closure of a restriction (Definition 5.11) “sync up”:

Lemma 5.16. *Let (G, α) be nice. Let T be a decision tree over $E(G)$ and $\pi \in \text{Branches}_\perp(\text{Prune}_{G,\alpha}(T))$. Then $G - \text{closure}_G(\text{supp}(\pi))$ does not have a giant component.*

Proof. We proceed by induction on the depth of T . The base case is that T has depth 0, i.e. T is a leaf bit. Since α is odd G must have a giant component, hence $\text{Branches}_\perp(\text{Prune}_{G,\alpha}(T))$ is the empty set and the base case holds.

For the inductive step, we write T as $T = (e; T_0, T_1)$. We consider two mutually exhaustive cases. The first case is that e is not a bridge in G . In this case we have

$$\text{Prune}_{G,\alpha}(T) = (e; \text{Prune}_{G-\{e\},\alpha|e \rightarrow 0}(T_0); \text{Prune}_{G-\{e\},\alpha|e \rightarrow 1}(T_1)).$$

Therefore $e \in \text{supp}(\pi)$, and we write $\pi = (e, b) \circ \tilde{\pi}$ for some $b \in \mathbb{Z}_2$ and $\tilde{\pi} \in \text{Branches}_\perp(\text{Prune}_{G-\{e\}, \alpha|_{e \rightarrow b}}(T_b))$. We observe that

$$\text{closure}_G(\text{supp}(\pi)) = \{e\} \cup \text{closure}_{G-\{e\}}(\text{supp}(\tilde{\pi})).$$

By the inductive hypothesis, $G - \{e\} - \text{closure}_{G-\{e\}}(\text{supp}(\tilde{\pi}))$ does not have a giant component, so the above equality gives that $G - \text{closure}_G(\text{supp}(\pi))$ does not have a giant component, and we are done in this case.

It remains to consider the case when e is a bridge in G . In this case $\text{Prune}_{G, \alpha}(T)$ shortcuts it according to some $b \in \mathbb{Z}_2$, i.e.

$$\text{Prune}_{G, \alpha}(T) = \text{Prune}_{G-\{e\}, \alpha|_{e \rightarrow b}}(T_b),$$

and so $e \notin \text{supp}(\pi)$ and $\pi \in \text{Branches}_\perp(\text{Prune}_{G-\{e\}, \alpha|_{e \rightarrow b}}(T_b))$. Since e is a bridge in G , it is certainly a bridge in $G - \text{supp}(\pi)$, i.e. $e \in \text{closure}_G(\pi)$. Note that we again have that

$$\text{closure}_G(\text{supp}(\pi)) = \{e\} \cup \text{closure}_{G-\{e\}}(\text{supp}(\pi)),$$

because removing a bridge cannot create new bridges or cause edges that were previously bridges to become non-bridges. By the inductive hypothesis, $G - \{e\} - \text{closure}_{G-\{e\}}(\text{supp}(\pi))$ does not have a giant component, but note that this is exactly $G - \text{closure}_G(\text{supp}(\pi))$, so again we are done. \square

Lemma 5.17. *Let (G, α) be nice and T be a decision tree over $E(G)$. Suppose π is a restriction such that $\text{Prune}_{G, \alpha}(T) \upharpoonright \pi = b$ for some $b \in \mathbb{Z}_2$, π pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$, and $G - \text{closure}_G(\text{supp}(\pi))$ has a giant component. Then $T \upharpoonright \text{closure}_{G, \alpha}(\pi) = b$. Moreover, $\text{closure}_{G, \alpha}(\pi)$ pushes the contradiction of α into the giant component of $G - \text{closure}_G(\text{supp}(\pi))$.*

Proof. The last sentence of the lemma statement follows from the preceding portion of the lemma by Proposition 5.12. For the preceding portion, we proceed by induction on the depth of T , noting that the base case when T has depth 0 is trivial. For the inductive step, let $T = (e; T_0, T_1)$. We consider two cases, depending on whether or not e is a bridge in G .

Case 1: e is not a bridge in G . Then

$$\text{Prune}_{G, \alpha}(T) = (e; \text{Prune}_{G-\{e\}, \alpha|_{(e \rightarrow 0)}}(T_0), \text{Prune}_{G-\{e\}, \alpha|_{(e \rightarrow 1)}}(T_1)).$$

Therefore $e \in \text{supp}(\pi)$, and we write $\pi = (e, b') \circ \tilde{\pi}$ for some $b' \in \mathbb{Z}_2$ and restriction $\tilde{\pi}$ such that $\text{Prune}_{G-\{e\}, \alpha|_{e \rightarrow b'}}(T_{b'}) \upharpoonright \tilde{\pi} = b$. We claim that

$$\text{closure}_{G, \alpha}(\pi) = (e, b') \circ \text{closure}_{G-\{e\}, \alpha|_{e \rightarrow b'}}(\tilde{\pi}), \tag{1}$$

noting first that the two restrictions have the same support. Clearly the edges in $\text{supp}(\pi) = \{e\} \cup \text{supp}(\tilde{\pi})$ are fixed the same way by the two restrictions; as for the edges in $\text{closure}_{G, \alpha}(\pi) - \text{supp}(\pi)$, the fact that they are fixed the same way follows from the from the fact that

$$\begin{aligned} G - \text{supp}(\pi) &= G - \{e\} - \text{supp}(\tilde{\pi}) \\ \alpha \upharpoonright \pi &= (\alpha \upharpoonright e \rightarrow b') \upharpoonright \tilde{\pi}. \end{aligned}$$

Given (1), we have that

$$\begin{aligned} T \upharpoonright \text{closure}_{G,\alpha}(\pi) &= T \upharpoonright (e, b') \circ \text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\tilde{\pi}) \\ &= T_{b'} \upharpoonright \text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\tilde{\pi}) \end{aligned}$$

and the inductive step then follows by the induction hypothesis applied to $G - \{e\}$, $\alpha \upharpoonright e \rightarrow b'$, $T_{b'}$, and $\tilde{\pi}$.

Case 2: e is a bridge in G . In this case $\text{Prune}_{G,\alpha}(T)$ shortcuts it according to some $b' \in \mathbb{Z}_2$, i.e.

$$\text{Prune}_{G,\alpha}(T) = \text{Prune}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(T_{b'}).$$

We consider two cases, depending on whether e is in $\text{supp}(\pi)$. If $e \in \text{supp}(\pi)$ we first note that $\pi(e) = b'$ (i.e. π fixes e the same way it is shortcut in $\text{Prune}_{G,\alpha}(T)$), since otherwise by Fact 5.9 it cannot be the case that π pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$. We may therefore write $\pi = (e, b') \circ \tilde{\pi}$ where $\tilde{\pi}$ is a restriction such that $\text{Prune}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(T_{b'}) \upharpoonright \tilde{\pi} = b$, and the remainder of the argument proceeds exactly as in Case 1 above.

If $e \notin \text{supp}(\pi)$, we claim that

$$\text{closure}_{G,\alpha}(\pi) = (e, b') \circ \text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\pi), \quad (2)$$

To establish (2), we note first that the two restrictions have the same support (since removing a bridge cannot create new bridges or cause edges that were previously bridges to become non-bridges). Next, by our assumptions that π pushes the contradiction of α into the giant component of $G - \text{supp}(\pi)$ and that $G - \text{closure}_{G,\alpha}(\pi)$ has a giant component, we may apply Proposition 5.12 to get that:

(†) $\text{closure}_{G,\alpha}(\pi)$ pushes the contradiction of α into the giant component of $G - \text{supp}(\text{closure}_G(\pi))$.

Clearly the edges in $\text{supp}(\pi)$ are fixed the same way by both restrictions in (2). As for the edges in $\text{closure}_{G,\alpha}(\pi) \setminus \text{supp}(\pi)$, we first consider e . Note that $(\text{closure}_{G,\alpha}(\pi))(e) = b'$ (i.e. that $\text{closure}_{G,\alpha}(\pi)$ sets e the same way it is shortcut in $\text{Prune}_{G,\alpha}(T)$), since otherwise (†) cannot hold. It remains to consider the edges in $\text{closure}_{G,\alpha}(\pi) - (\text{supp}(\pi) \cup \{e\})$. Fix any such edge e' . We claim that

$$(\text{closure}_{G,\alpha}(\pi))(e') = (\text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\pi))(e'). \quad (3)$$

By (†) along with the fact that $(e, b') \circ \pi$ is a subrestriction of $\text{closure}_{G,\alpha}(\pi)$, we have that π pushes the contradiction of $\alpha \upharpoonright e \rightarrow b'$ into the giant component of $(G - \{e\}) - \text{supp}(\pi)$. Applying Proposition 5.12, it then follows that:

(††) $\text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\pi)$ pushes the contradiction of $\alpha \upharpoonright e \rightarrow b'$ into the giant component of $G - (\{e\} \cup \text{closure}_{G-\{e\}}(\pi))$.

Again we let C be the component in $G - \text{supp}(\pi)$ that contains e' , and let C_1 and C_2 be the two components of $C - \{e'\}$ where $|V(C_1)| \geq |V(C_2)|$.

- If the $(\alpha \upharpoonright \pi)$ -charge of C is even, then C is non-giant (and so are C_1 and C_2) and $(\text{closure}_{G,\alpha}(\pi))(e') = b^* \in \mathbb{Z}_2$ where b^* is the unique assignment to e such that the induced charges in C_1 and C_2 are both even. If $\text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\pi)$ sets e' to \bar{b}^* , the induced $(\alpha \upharpoonright \pi \circ (e \rightarrow b') \circ (e' \rightarrow \bar{b}^*))$ -charges of both C_1 and C_2 would be odd and this contradicts (††), the fact that $\text{closure}_{G-\{e\}, \alpha \upharpoonright e \rightarrow b'}(\pi)$ pushes the contradiction of $\alpha \upharpoonright e \rightarrow b'$ into the giant component of $G - (\{e\} \cup \text{closure}_{G-\{e\}}(\pi))$.

- If the $(\alpha \upharpoonright \pi)$ -charge of C is odd, then C is giant and $|V(C_1)| > n/2$ since G -closure $_G(\text{supp}(\pi))$ has a giant component. By the same argument as above, we have that (3) holds.

Thus we have established (2). Having established (2), we conclude that

$$\begin{aligned} T \upharpoonright \text{closure}_{G,\alpha}(\pi) &= T \upharpoonright (e, b') \circ \text{closure}_{G-\{e\},\alpha \upharpoonright e \rightarrow b'}(\pi) \\ &= T_{b'} \upharpoonright \text{closure}_{G-\{e\},\alpha \upharpoonright e \rightarrow b'}(\pi), \end{aligned}$$

and the inductive step follows by the induction hypothesis applied to $G - \{e\}$, $\alpha \upharpoonright e \rightarrow b'$, $T_{b'}$, and π (note that $\text{Prune}_{G-\{e\},\alpha \upharpoonright e \rightarrow b'}(T_{b'}) \upharpoonright \pi = b$). \square

6 k -Evaluations

In the previous section, we defined good decision trees, which are tailored to the Tseitin formulas. The variables that are queried and set along any path are required to be independent, but we view the associated restriction as not just the assignment to these variables, but the unique assignment for the closure. Moreover, we defined a pruning procedure where we truncate any path that could quickly lead to a contradiction. Specifically, if a path (and its associated restriction) leaves us with a graph containing a small component of odd charge, then this path will be truncated since the Tseitin contradiction under this partial assignment has become “too easy.”

In this section we define what it means for a good decision tree to represent a formula. We stress that a decision tree representing a formula is in no way truth functionally equivalent to the formula. In fact, the original Tseitin formula (which is unsatisfiable) will be represented by a 1-tree — a shallow tree where all leaves are labelled by 1 — and indeed this is essential to the proof complexity argument.

The sense in which a decision tree represents a formula is purely local: if a verifier checks the soundness of any given step in the proof (using the locally consistent decision trees in place of the active subformulas for that inference step), no inconsistency will be detected. This means that if we follow a branch π down the tree $\mathcal{T}(A)$ (this is the good tree that will be associated with the formula A) and it leads to a leaf labeled 1, then any branch in the tree $\mathcal{T}(\neg A)$ that is consistent with π will lead to a leaf labeled 0. Similarly, if we follow a branch down the tree $\mathcal{T}(A \vee B)$ for a formula $A \vee B$ and the branch reaches a 1-leaf, then there is either a consistent branch in $\mathcal{T}(A)$ leading to a 1-leaf or there is a consistent branch in $\mathcal{T}(B)$ leading to a 1-leaf.

The following definition makes this precise.

Definition 6.1 ($(H^{(i)}, \alpha')$ -represents.). *Fix $0 \leq i \leq d < d^*$, $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$, and an odd charge $\alpha \in \mathbb{Z}_2^{V(\mathcal{G}_n)}$, and let $\alpha' = \alpha \upharpoonright \rho^{(i)}$. Let T_1, \dots, T_m be $H^{(i)}$ -independent decision trees over variable set $E(H^{(i)})$. A $(\tau(i, n)/(3\lambda^*), H^{(i)})$ -good¹ decision tree T is said to $(H^{(i)}, \alpha')$ -represent $\bigvee_{j=1}^m T_j$ if for all $b \in \mathbb{Z}_2$,*

$$\pi \in \text{Branches}_b(T) \implies (\bigvee_j \text{Disj}(T_j)) \upharpoonright \text{closure}_{H^{(i)},\alpha'}(\pi) = b.$$

We want to show that (after applying a suitable restriction) we can associate a good decision tree with each subformula in the (restricted) proof. Such a collection of trees, called a k -evaluation,

¹This depth bound on T implies that using Fact 5.7, $\text{closure}_{H^{(i)},\alpha'}(\pi)$ will be well defined for any branch π in T .

must satisfy some important properties. Namely, the initial Tseitin clauses are mapped to 1-trees, the final (identically 0) formula in the proof is mapped to a 0-tree, and all subformulas in the proof map to good decision trees satisfying the local consistency property mentioned above. The remainder of this section gives the formal definition of a k -evaluation, and shows that if we can construct a k -evaluation for an alleged Frege proof of the Tseitin formula, then it implies our lower bound.

Definition 6.2 (k -evaluation). *As in Definition 6.1, fix $0 \leq i \leq d < d^*$, $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$, and an odd charge $\alpha \in \mathbb{Z}_2^{V(\mathcal{G}_n)}$, and let $\alpha' = \alpha \upharpoonright \rho^{(i)}$.*

Let Γ be a sequence of formulas over the variable set $E(H^{(i)})$ such that (i) Γ is closed under sub-formulas, and (ii) Γ includes all the clauses of $\text{Tseitin}(H^{(i)}[\alpha'])$. Let $k < \tau(i, n)/(3\lambda^)$.*

A k -evaluation for Γ is a mapping $\mathcal{T}(\cdot)$ which assigns to each formula $A \in \Gamma$ a total, $(k, H^{(i)})$ -good decision tree $\mathcal{T}(A)$ satisfying the following properties:

1. $\mathcal{T}(b) = b$ for $b \in \mathbb{Z}_2$;
2. $\mathcal{T}(\neg A) = \mathcal{T}(A)^c$;
3. If $A = \vee_j A_j$ then $\mathcal{T}(A)$ $(H^{(i)}, \alpha')$ -represents $\vee_j \mathcal{T}(A_j)$;
4. If A is a clause of $\text{Tseitin}(H^{(i)}[\alpha'])$ then $\mathcal{T}(A)$ is a 1-tree;
5. For every tree $T_0 = \mathcal{T}(A_0)$, every collection of at most 6 other trees $T_i = \mathcal{T}(A_i)$, $i = \{1, \dots, 6\}$ where $A_0, \dots, A_6 \in \Gamma$, and every branch π_0 in $\text{Branches}(T_0)$, there exists a restriction π^* , extending π_0 , such that $\text{closure}_{H^{(i)}, \alpha'}(\pi^*) = \pi^*$ and π^* pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi^*)$. Furthermore for each $i \in \{0, 1, \dots, 6\}$, $T_i \upharpoonright \pi^* = b_i$ for some $b_i \in \mathbb{Z}_2$.

The high level context of how the next lemma will eventually be applied (in Section 7) is as follows. Starting with a small refutation of $\text{Tseitin}(\mathcal{G}_n[\alpha])$ (i.e. a sequence of formulas P), let $P \upharpoonright \rho^{(d)}$ denote the sequence of formulas obtained by applying a restriction $\rho^{(d)}$ to every formula in the proof, where $(\rho^{(d)}, H^{(d)})$ is in the support of our final $\mathcal{A}^{(d)}$ distribution. The ‘ P ’ of Lemma 6.3 will be such a $P \upharpoonright \rho^{(d)}$. Since proofs are closed under restrictions, $P \upharpoonright \rho^{(d)}$ will be a small refutation of $\text{Tseitin}(H^{(d)}[\alpha^{(d)}])$, where $\alpha^{(d)}$ is $\alpha \upharpoonright \rho^{(d)}$ restricted to $V(H^{(d)})$. (Recall from Section 2.3 that by “a restriction to the proof,” what we mean is just a substitution of variables by their assigned values, without any further simplification.) The P^* of Lemma 6.3 will be the set of all subformulas of $P \upharpoonright \rho$.

Lemma 6.3. *Fix $(\rho^{(d)}, H^{(d)}) \in \text{supp}(\mathcal{A}^{(d)})$ where $d \leq d^*$ and an odd charge $\alpha \in \mathbb{Z}_2^{V(\mathcal{G}_n)}$, and let $\alpha^{(d)} = \alpha \upharpoonright \rho^{(d)}$. Let P be a sequence of formulas over variable set $E(H^{(d)})$ containing the clauses of $\text{Tseitin}(H^{(d)}[\alpha^{(d)}])$, and let P^* be the set of all sub-formulas of P . If P^* has a k -evaluation where $k < \tau(d, n)/(3\lambda^*)$, then P cannot be a Frege refutation of $\text{Tseitin}(\mathcal{G}_n[\alpha])$.*

Proof. We will show that P cannot be a Frege refutation of $\text{Tseitin}(\mathcal{G}_n[\alpha])$ by proving that for every formula A in P , $\mathcal{T}(A)$ must be a 1-tree (recall that this means that all of the leaves of $\mathcal{T}(A)$ are labelled by 1). This will yield a contradiction since the final formula in the derivation is 0, and by property (1) of Definition 6.2, $\mathcal{T}(0)$ is a 0-tree.

The proof is by induction on the number of inference steps used to derive A . For the base case, there are no inference steps and thus A must be an initial clause of $\text{Tsetin}(H^{(d)}[\alpha^{(d)}])$. (If A is an instance of the axiom scheme, then A is derived by one inference step; such formulas will be handled in the inductive step.) When A is an initial clause, property (4) of Definition 6.2 immediately implies that $\mathcal{T}(A)$ is a 1-tree.

For the inductive step, assume that $\mathcal{T}(F)$ is a 1-tree for each formula F in P that was derived by at most i inference steps. Now consider a formula A that is derived in at most $i + 1$ inference steps. Such an A is derived from zero, one or two previously derived formulas, which were themselves each derived in at most i inference steps, using some rule from Section 2.1. There are different cases depending on which inference rule was used to derive A , but the proof will be essentially the same in all these cases.

Suppose that the rule used to derive A derives C_0 from C_1 and C_2 . That is, consider an instance of a rule:

$$A_1[B_1/p_1, \dots, B_m/p_m], A_2[B_1/p_1, \dots, B_m/p_m] \rightarrow A_0[B_1/p_1, \dots, B_m/p_m].$$

Let $\Gamma = \{C_0, C_1, \dots, C_j\}$ be the set of all formulas $D[B_1/p_1, \dots, B_m/p_m]$, where $D[p_1, \dots, p_m]$ is a sub-formula of some A_i ; note that Γ only includes the sub-formulas in which each B_i is substituted for p_i and does not include “lower-level sub-formulas”. (These are called the *active* subformulas of the inference.) For our proof system \mathcal{F} we have that $|\Gamma|$ is always at most 7, since the associative rule and the cut rule both have 7 active sub-formulas and the other rules have fewer.

We will use the following running example to demonstrate the idea for the associative rule (similar reasoning goes through for any of the other rules in \mathcal{F}). Suppose that the formula $C_0 = (F \vee G) \vee H$ is derived from $C_1 = F \vee (G \vee H)$ by the associative rule: $p_1 \vee (p_2 \vee p_3) \rightarrow (p_1 \vee p_2) \vee p_3$. Thus in the above notation we have that $B_1 = F$, $B_2 = G$ and $B_3 = H$, and the set of active subformulas is $\Gamma = \{C_0 = (F \vee G) \vee H, C_1 = F \vee (G \vee H), C_2 = F, C_3 = G, C_4 = H, C_5 = F \vee G, C_6 = G \vee H\}$. The subformulas in Γ are the only active ones because while there are other subformulas (for example F , G , and H may have proper subformulas), the soundness of the rule does not depend on these subformulas.

Assume for sake of contradiction that $\mathcal{T}(C_0)$ is not a 1-tree, and therefore, there exists a path π_0 in this tree such that $\mathcal{T}(C_0) \upharpoonright \pi_0 = 0$. By Property (5) of Definition 6.2 there exists a restriction π^* extending π_0 such that $\text{closure}_{H^{(i)}, \alpha'}(\pi^*) = \pi^*$ and π^* pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi^*)$. Furthermore, for each $i \in \{0, 1, \dots, 6\}$, $\mathcal{T}(C_i) \upharpoonright \pi^* = b_i$ for some $b_i \in \mathbb{Z}_2$. By the definition of a k -evaluation, the restriction π^* must satisfy the following properties:

- If $\neg A \in \Gamma$, then $\mathcal{T}(\neg A) \upharpoonright \pi^* = 1$ implies that $\mathcal{T}(A) \upharpoonright \pi^* = 0$. This is by property (2) of the definition of a k -evaluation.
- If $(A \vee B) \in \Gamma$, then by property (3) in the definition of a k -evaluation (recalling $\text{closure}_{H^{(i)}, \alpha'}(\pi^*) = \pi^*$), if $\mathcal{T}(A \vee B) \upharpoonright \pi^* = 1$ then $\mathcal{T}(A) \upharpoonright \pi^* = 1$ or $\mathcal{T}(B) \upharpoonright \pi^* = 1$. Otherwise, if $\mathcal{T}(A \vee B) \upharpoonright \pi^* = 0$ then $\mathcal{T}(A) \upharpoonright \pi^* = 0$ or is a 0-tree, and likewise $\mathcal{T}(B) \upharpoonright \pi^*$; recalling property (5) in the definition of a k -evaluation, they are in fact both 0.

By induction we know that each of $\mathcal{T}(C_1), \dots, \mathcal{T}(C_6)$ are 1-trees. Note that by property (5) of Definition 6.2 we have that $\mathcal{T}(C_0) \upharpoonright \pi^*$ is either 0 or 1; we now show that $\mathcal{T}(C_0) \upharpoonright \pi^*$ cannot be 0. Since $C_0 = C_5 \vee C_4$, by the second bullet above, $\mathcal{T}(C_0) \upharpoonright \pi^* = 0$ implies that $\mathcal{T}(C_5) \upharpoonright \pi^* = 0$, and the same for $\mathcal{T}(C_4)$, but this contradicts the earlier assertion that each of $\mathcal{T}(C_1), \dots, \mathcal{T}(C_6)$

are 1-trees. Thus $\mathcal{T}(C_0) \upharpoonright \pi^* = 1$, contradicting the assumption that $\mathcal{T}(C_0) \upharpoonright \pi_0 = 0$ (recall that π^* extends π_0).

We can apply similar reasoning for other rules; for example, consider an application of a rule “ C_1, C_2 implies C_0 .” By induction, we know that both $T(C_1)$ and $T(C_2)$ are 1-trees. By the above properties, the truth value assignments for $\{T(C_i) \upharpoonright \pi^*\}_{i=0,\dots,6}$ respect the usual rules of logic, and thus because the rule is sound and only involves these sub-formulas, it follows that $T(C_0) \upharpoonright \pi^*$ must also be 1. But this contradicts our assumption that $T(C_0) \upharpoonright \pi^* = 0$. \square

7 Proof of Theorem 1: Obtaining a k -evaluation from the switching lemma

Let \mathcal{P} be a depth- d refutation of $\text{Tseitin}(\mathcal{G}_n[\alpha])$ where $d \leq d^* = c\sqrt{\log n}$, and let \mathcal{P}^* be the set of all subformulas of \mathcal{P} (note that $|\mathcal{P}^*|$ is polynomially related to $|\mathcal{P}|$). The main result of this section (Lemma 7.1) uses our final switching lemma (Theorem 2, stated below) to prove that if \mathcal{P}^* is “small”, then there exists a pair $(\rho^{(d)}, H^{(d)}) \in \text{supp}(\mathcal{A}^{(d)})$ such that $\mathcal{P}^* \upharpoonright \rho^{(d)}$ has a k -evaluation where $k = (\log n)/200d$. Recall that $\mathcal{P} \upharpoonright \rho^{(d)}$ is a depth- d refutation of $\text{Tseitin}(H^{(d)}[\alpha^{(d)}])$ where $\alpha^{(d)} = \alpha \upharpoonright \rho^{(d)}$, and observe that $\mathcal{P}^* \upharpoonright \rho^{(d)}$ is the set of all subformulas of $\mathcal{P} \upharpoonright \rho^{(d)}$. Applying Lemma 6.3 (with its P being $\mathcal{P} \upharpoonright \rho^{(d)}$ and its P^* being $\mathcal{P}^* \upharpoonright \rho^{(d)}$), it follows that $\mathcal{P} \upharpoonright \rho^{(d)}$ cannot be a refutation of $\text{Tseitin}(H^{(d)}[\alpha^{(d)}])$. This contradiction implies that \mathcal{P}^* cannot be “small”, and hence neither can \mathcal{P} .

Lemma 7.1. *Let Γ be a set of depth- d formulas over $E(\mathcal{G}_n)$, closed under subformulas, where $d \leq d^*$. For some absolute constant $c_1 > 0$, if $|\Gamma| < n^{c_1(\log n)/d^2}$, then there exists a pair $(\rho^{(d)}, H^{(d)}) \in \text{supp}(\mathcal{A}^{(d)})$ such that $\Gamma \upharpoonright \rho^{(d)}$ has a k -evaluation where $k = (\log n)/200d$.*

As discussed above, Theorem 1 follows directly from Lemmas 6.3 and Lemma 7.1, observing that $(\log n)/200d \ll \tau(d, n)/(3\lambda^*)$.

We now state our final switching lemma. The rest of this paper after Section 7 is devoted to proving Theorem 2. (In Theorem 2, $\varepsilon > 0$ is a small absolute constant).

Theorem 2 (Final switching lemma). *Let $0 \leq i \leq d < d^*$. Fix $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$. Let T_1, \dots, T_M be $(k, H^{(i)})$ -good decision trees over $E(H^{(i)})$ where $k = (\log n)/200d$ and let $\alpha' = \alpha \upharpoonright \rho^{(i)}$. Then except with failure probability at most $n^{-\varepsilon(\log n)/d^2}$ over a draw of $(\rho^{(i+1)}, \mathbf{H}^{(i+1)}) \sim \mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)})$, there is a $(k, \mathbf{H}^{(i+1)})$ -good decision tree over $E(\mathbf{H}^{(i+1)})$ that $(\mathbf{H}^{(i+1)}, \alpha' \upharpoonright \rho^{(i+1)})$ -represents $\bigvee_j \text{Prune}_{\mathbf{H}^{(i+1)}, \alpha' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)})$.*

7.1 Proof of Lemma 7.1 assuming Theorem 2

Given $i \in \{0, \dots, d\}$, let $\Gamma_i \subseteq \Gamma$ denote the set of all formulas in Γ of depth at most i . We will first prove, by induction on i , that there exists a pair $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$ such that the following holds: there is a mapping $\mathcal{T}^{(i)}(\cdot)$ which assigns to each formula $A \in \Gamma_i \upharpoonright \rho^{(i)}$ (note that these are formulas over the variable set $E(H^{(i)})$) a $(k := (\log n)/200d, H^{(i)})$ -good decision tree satisfying properties (1)–(3) of the definition of a k -evaluation (Definition 6.2).

For our base case when $i = 0$, recall that $\mathcal{A}^{(0)}$ only contains the trivial pair $(\{*\}^{E(\mathcal{G}_n)}, \mathcal{G}_n)$. The only formulas we need to consider are constants and literals of form e or $\neg e$. For $b \in \mathbb{Z}_2$, we

define $\mathcal{T}^{(0)}(b) = b$ (the one-node tree comprising a single leaf b). Next suppose that $e \in E(\mathcal{G}_n)$ is a bridge in \mathcal{G}_n connecting components C_1 and C_2 where $|V(C_1)| \geq |V(C_2)|$. (In fact, C_1 is a giant component since removing a single edge from \mathcal{G}_n leaves a giant component; see the proof of Corollary 2.8.) In this case $\mathcal{T}^{(0)}(e) = 1$ if the induced $(\alpha \upharpoonright e \rightarrow 1)$ -charge of C_1 is odd, and $\mathcal{T}^{(0)}(e) = 0$ otherwise. Finally, suppose e is not a bridge in \mathcal{G}_n . In this case $\mathcal{T}^{(0)}(e) = (e; 0, 1)$, and $\mathcal{T}^{(0)}(\neg e) = (e; 1, 0) = \mathcal{T}^{(0)}(e)^c$. It is easy to check that these trees are (trivially) (k, \mathcal{G}_n) -good and properties (1)–(3) hold (property (3) is vacuous in this case).

Now assume by the inductive hypothesis that there exists a pair $(\rho^{(i)}, H^{(i)}) \in \text{supp}(H^{(i)})$ and a mapping $\mathcal{T}^{(i)}$ as described above. Note that Γ_{i+1} is the disjoint union of Γ_i , Υ_{i+1} , and Ξ_{i+1} , where Υ_{i+1} are the formulas in Γ_{i+1} of depth exactly $i + 1$ of the form $\bigvee_j A_j$, and Ξ_{i+1} are the formulas in Γ_{i+1} of depth exactly $i + 1$ of the form $\neg A$.

- (a) First consider a formula $A \in \Upsilon_{i+1}$ of the form $A = \bigvee_j A_j$. By the induction hypothesis, for all j the tree $\mathcal{T}^{(i)}(A_j \upharpoonright \rho^{(i)})$ is a $(k, H^{(i)})$ -good decision tree satisfying properties (1)–(3) of Definition 6.2. Applying Theorem 2 and a union bound over $|\Upsilon_{i+1}| < n^{c_1(\log n)/d^2}$ many A 's in Υ_{i+1} (with a suitable choice of c_1 relative to ε), we get there exists some $(\rho^{(i+1)}, H^{(i+1)}) \in \text{supp}(\mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)}))$ satisfying the following: for all $A \in \Upsilon_{i+1}$, there is a $(k, H^{(i+1)})$ -good decision tree that $(H^{(i+1)}, \alpha')$ -represents $\bigvee_j \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A_j \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)})$, where $\alpha' \in \mathbb{Z}_2^{V(H^{(i+1)})}$ denotes $\alpha \upharpoonright \rho^{(i+1)}$ restricted to the vertices in $V(H^{(i+1)})$. We define $\mathcal{T}^{(i+1)}(A \upharpoonright \rho^{(i+1)})$ to be this tree (and observe that it is $(k, H^{(i+1)})$ -good as desired).
- (b) Next, consider a formula $\neg A \in \Xi_{i+1}$. For the extension $\rho^{(i+1)}$ of $\rho^{(i)}$ whose existence is asserted in the bullet above, we define

$$\mathcal{T}^{(i+1)}(\neg A \upharpoonright \rho^{(i+1)}) = \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)})^c,$$

which is $(k, H^{(i+1)})$ -good (since $\text{Prune}_{H^{(i+1)}, \alpha'}(\cdot)$ yields an $H^{(i+1)}$ -independent tree, and by our induction hypothesis $\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)})$ has depth at most k).

- (c) Finally, consider a formula $A \in \Gamma_i$. In this case we define

$$\mathcal{T}^{(i+1)}(A \upharpoonright \rho^{(i+1)}) = \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)}),$$

where again $\rho^{(i+1)}$ is the extension of $\rho^{(i)}$ whose existence is asserted in the first bullet above. By the same reasons as above, this tree is $(k, H^{(i+1)})$ -good.

It remains to argue that this map $\mathcal{T}^{(i+1)}(\cdot)$ as defined above satisfies properties (1)–(3) of Definition 6.2. Property (1) is immediate. For property (2), we would like to show that $\mathcal{T}^{(i+1)}(\neg A \upharpoonright \rho^{(i+1)}) = \mathcal{T}^{(i+1)}(A \upharpoonright \rho^{(i+1)})^c$ for all $\neg A \in \Gamma_{i+1}$. If $\neg A \in \Gamma_{i+1}$ we have:

$$\begin{aligned} \mathcal{T}^{(i+1)}(\neg A \upharpoonright \rho^{(i+1)}) &= \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(\neg A \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)}) \\ &= \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)})^c \upharpoonright \rho^{(i+1)}) \\ &= \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}((A \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)})^c) \\ &= \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)})^c \\ &= \mathcal{T}^{(i+1)}(A \upharpoonright \rho^{(i+1)})^c, \end{aligned}$$

where the first equality is by the definition of $\mathcal{T}^{(i+1)}$ (specifically, (b) if $\neg A \in \Xi_{i+1}$, and (c) if $\neg A \in \Gamma_i$); the second equality holds because property (2) holds for $\mathcal{T}^{(i)}$ by the induction hypothesis; the third equality holds because applying a restriction and toggling the leaf bits are commutative; the fourth equality holds because pruning and toggling the leaf bits are commutative; and the last equality holds by the definition of $\mathcal{T}^{(i+1)}$.

We will now verify property (3). If $A \in \Upsilon_{i+1}$, it follows from (a) and (c). Next we consider $A \in \Gamma_i$. Let $A = \vee_j A_j$, and $\alpha'' \in \mathbb{Z}_2^{V(H^{(i)})}$ denote $\alpha \upharpoonright \rho^{(i)}$ restricted to $V(H^{(i)})$. By the induction hypothesis, we have that $\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)})$ ($H^{(i)}, \alpha''$)-represents $\vee_j \mathcal{T}^{(i)}(A_j \upharpoonright \rho^{(i)})$. By two applications of (c), it suffices to prove that

$$\text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)}) \text{ } (H^{(i+1)}, \alpha')\text{-represents } \vee_j \text{Prune}_{H^{(i+1)}, \alpha'}(\mathcal{T}^{(i)}(A_j \upharpoonright \rho^{(i)}) \upharpoonright \rho^{(i+1)}),$$

which follows from the following lemma (recall that $\alpha' = \alpha'' \upharpoonright \rho^{(i+1)}$ restricted to $V(H^{(i+1)})$):

Lemma 7.2. *Let $k = (\log n)/200d$ and T_1, \dots, T_ℓ be $(k, H^{(i)})$ -good decision trees over $E(H^{(i)})$, and let T be a $(k, H^{(i)})$ -good decision tree that $(H^{(i)}, \alpha'')$ -represents $\vee_j T_j$. Then $\text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T \upharpoonright \rho^{(i+1)})$ ($H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}$)-represents $\vee_j \text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)})$.*

Given Lemma 7.2 (the proof of which we defer to the next subsection), we have shown that properties (1)–(3) hold for all $i \in \{0, 1, \dots, d\}$.

Properties (4) and (5). It remains to argue that for each formula $A \in \Gamma_d \upharpoonright \rho^{(d)}$, the tree $\mathcal{T}^{(d)}(A)$ satisfies properties (4) and (5) of Definition 6.2.

For property (4), consider a clause A of $\text{Tseitin}(H^{(d)}[\alpha \upharpoonright \rho^{(d)}])$. Such a clause is one of the four length-3 clauses whose AND gives a constraint from some vertex $v \in V(\mathcal{G}_n)$; let this vertex be v and let the clause A be $e_1 \vee e_2 \vee e_3$, and let the charge at v be $\alpha(v) = 1$. (Other possibilities can be handled similarly.) From the base case we have that for each $j \in [3]$, $\mathcal{T}^{(0)}(e_j \upharpoonright \rho^{(0)})$ is the depth-1 tree $(e_j; 0, 1)$, and from the $i = 0$ case of (a) we have that $\mathcal{T}^{(1)}(A \upharpoonright \rho^{(1)})$ is a $(k, H^{(1)})$ -good decision tree that $(H^{(1)}, \alpha \upharpoonright \rho^{(1)})$ -represents $\vee_{j=1}^3 \text{Prune}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}((e_j; 0, 1) \upharpoonright \rho^{(1)})$. We will claim that already $\mathcal{T}^{(1)}(A \upharpoonright \rho^{(1)})$ is a 1-tree, from which property (4) follows. Seeking a contradiction, suppose there exists $\pi \in \text{Branches}_0(\mathcal{T}^{(1)}(A \upharpoonright \rho^{(1)}))$. By Definition 6.1, for all three values $j \in [3]$ we have that

$$\text{Prune}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}((e_j; 0, 1) \upharpoonright \rho^{(1)}) \upharpoonright \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi) = 0 \quad \text{for } j = 1, 2, 3. \quad (4)$$

Since $\mathcal{T}^{(1)}(A \upharpoonright \rho^{(1)})$ is $(k, H^{(1)})$ -good, we have that $\text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ pushes the contradiction of $\alpha \upharpoonright \rho^{(1)}$ into the giant component of $H - \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ (Fact 5.14 and Proposition 5.12); this implies that $\rho \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ is α -consistent. However, by Lemma 5.17 applied to Equation (4) (noting that the closure of a closed set is itself) we have that

$$((e_j; 0, 1) \upharpoonright \rho^{(1)}) \upharpoonright \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi) = 0 \quad \text{for } j = 1, 2, 3.$$

This implies that $(\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi))(e_j) = 0$ for $j = 1, 2, 3$, which contradicts the fact that $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ is α -consistent (since $\alpha(v) = 1$).

It remains to prove property (5) which follows from the following proposition:

Proposition 7.3. *Fix $(\rho^{(i)}, H_n^{(i)}) \in \text{supp}(\mathcal{A}_n^{(i)})$ and let $\alpha' = \alpha \upharpoonright \rho^{(i)}$. Let T_0, T_1, \dots, T_6 be $(k, H_n^{(i)})$ -good decision trees where $k = (\log n)/200d$. Then for every $\pi_0 \in \text{Branches}(T_0)$, there exist $\pi_i \in \text{Branches}(T_i)$ for $i \in [6]$ such that*

1. $\{\pi_0, \dots, \pi_6\}$ is mutually compatible.

2. $\pi^* := \text{closure}_{H^{(i)}, \alpha'}(\pi_0 \circ \dots \circ \pi_6)$ pushes the contradiction of α' into the giant component of $H - \text{supp}(\pi^*)$.

Proof. Since π_0 is $H^{(i)}$ -independent, we have that π_0 pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi)$. Equivalently, the $(\alpha' \upharpoonright \pi_0)$ -charge of a component C of $H^{(i)} - \text{supp}(\pi_0)$ is odd iff C is giant (note that $H^{(i)} - \text{supp}(\pi_0)$ is connected).

Consider $T_1 \upharpoonright \pi_0$, a decision tree over $E(H^{(i)} - \text{supp}(\pi_0))$. Observe that

$$\text{Branches}_\perp(\text{Prune}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(T_1 \upharpoonright \pi_0)) = \emptyset$$

by Lemma 5.16 since $(H^{(i)} - \text{supp}(\pi_0)) - \text{closure}_{H^{(i)} - \text{supp}(\pi_0)}(\tilde{\pi}_1)$ has a giant component (by Fact 5.7 using $|H^{(i)}| \geq \sqrt{n}$ and $k = (\log n)/200d$). Therefore, we may fix

$$\tilde{\pi} \in \text{Branches}_b(\text{Prune}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(T_1 \upharpoonright \pi_0))$$

for some $b \in \mathbb{Z}_2$, and by Fact 5.14, we have that $\tilde{\pi}_1$ pushes the contradiction of $\alpha' \upharpoonright \pi_0$ into the giant component of $(H^{(i)} - \text{supp}(\pi_0)) - \text{supp}(\tilde{\pi}_1)$. By Proposition 5.12, we have that $\text{closure}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(\tilde{\pi}_1)$ pushes the contradiction of $\alpha' \upharpoonright \pi_0$ into the giant component of $(H^{(i)} - \text{supp}(\pi_0)) - \text{closure}_{H^{(i)} - \text{supp}(\pi_0)}(\tilde{\pi}_1)$. Equivalently,

(\dagger) $\pi_0 \circ \text{closure}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(\tilde{\pi}_1)$ pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi_0) - \text{closure}_{H^{(i)} - \text{supp}(\pi_0)}(\tilde{\pi}_1)$.

Recall that $\tilde{\pi}_1 \in \text{Branches}_b(\text{Prune}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(T_1 \upharpoonright \pi_0))$. By Lemma 5.17, we have that

$$(T_1 \upharpoonright \pi_0) \upharpoonright \text{closure}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(\tilde{\pi}_1) = b.$$

This implies the existence of a path $\pi_1 \in \text{Branches}_b(T_1)$ such that π_1 is a subrestriction of $\pi_0 \circ \text{closure}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(\tilde{\pi}_1)$. Note that π_1 is consistent with π_0 , and since $\pi_0 \circ \pi_1$ is a subrestriction of $\pi_0 \circ \text{closure}_{H^{(i)} - \text{supp}(\pi_0), \alpha' \upharpoonright \pi_0}(\tilde{\pi}_1)$ it follows from (\dagger) that $\pi_0 \circ \pi_1$ pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi_0 \circ \pi_1)$.

The existence of π_2, \dots, π_6 with the claimed properties follows by repeating the above argument, where to get the argument started we observe that (as we have just shown) $\pi_0 \circ \pi_1$ pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi_0 \circ \pi_1)$. Since $|\pi_0 \cup \dots \cup \pi_6| \leq 7k = 7(\log n)/200d$, we apply Fact 5.7 and Proposition 5.12 to conclude that $\pi^* := \text{closure}_{H^{(i)}, \alpha'}(\pi_0 \circ \dots \circ \pi_6)$ pushes the contradiction of α' into the giant component of $H^{(i)} - \text{supp}(\pi^*)$. \square

This concludes the proof of Lemma 7.1, modulo the proof of Lemma 7.2.

7.2 Proof of Lemma 7.2

We will need the following lemma:

Lemma 7.4. Fix $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$ and $(\rho^{(i+1)}, H^{(i+1)}) \in \text{supp}(\mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)}))$, and let $\alpha'' = \alpha \upharpoonright \rho^{(i)}$. Let π be a restriction to the edges in $E(H^{(i+1)})$ such that $\text{supp}(\pi)$ is an $H^{(i+1)}$ -independent set of size at most $\tau(i+1, n)/3\lambda^*$, where λ^* is the λ^* from Fact 5.7. Let β be a sub-restriction of $\rho \circ \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi)$ where $\text{supp}(\beta)$ is a $H^{(i)}$ -independent set of size at most $\tau(i+1, n)/3\lambda^*$. Then $\rho \circ \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi)$ extends $\text{closure}_{H^{(i)}, \alpha''}(\beta)$.

Proof. For notational clarity we prove the $i = 0$ case, noting that the proof of the general case proceeds along essentially identical lines. In this case $\rho^{(0)} = \{*\}^{E(\mathcal{G}_n)}$, $H^{(0)} = \mathcal{G}_n$, and $\alpha'' = \alpha$. We write n' for $\tau(1, n)$, the number of real vertices of $H^{(1)}$.

Note that $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ extends β , and so the lemma statement is equivalent to the claim that every edge $e \in (\text{closure}_{\mathcal{G}_n}(\text{supp}(\beta)) \setminus \text{supp}(\beta))$ is fixed to the same constant $b \in \mathbb{Z}_2$ by $\text{closure}_{\mathcal{G}_n, \alpha}(\beta)$ and by $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$. Fix any such edge e , and recall that e is a bridge in $\mathcal{G}_n - \text{supp}(\beta)$. Since $\mathcal{G}_n - \text{supp}(\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)) = H^{(1)} - \text{closure}_{H^{(1)}}(\text{supp}(\pi))$ is bridgeless and $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ extends β , it follows that e is also fixed by $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$; it remains to argue that $\text{closure}_{\mathcal{G}_n, \alpha}(\beta)$ and $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ fix e to the same constant.

Since $|\text{supp}(\pi)| \leq n'/3\lambda^*$, it follows from Fact 5.7 that $H^{(1)} - \text{closure}_{H^{(1)}}(\text{supp}(\pi))$ has a giant component. We may then apply Proposition 5.12 (recall that $\text{supp}(\pi)$ is $H^{(1)}$ -independent) to get that

(\dagger) $\text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ pushes the contradiction of $\alpha \upharpoonright \rho^{(1)}$ into the giant component of $H^{(1)} - \text{closure}_{H^{(1)}}(\text{supp}(\pi))$ (of size $> |V(H^{(1)})|/2$).

Since $\text{supp}(\beta)$ is \mathcal{G}_n -independent, we have that $\mathcal{G}_n - \text{supp}(\beta)$ is connected and has e as bridge. Let C_1 and C_2 be the two components of $\mathcal{G}_n - (\text{supp}(\beta) \cup \{e\})$ where $|V(C_1)| \geq |V(C_2)|$ and $|V(C_1)| + |V(C_2)| = |V(\mathcal{G}_n)| = n$. Since $|\text{supp}(\beta)| \leq n'/3\lambda^*$, it follows from Fact 5.7 that $|V(C_2)| \leq n'/3 < |V(H^{(1)})|/3$ and $|V(C_1)| \geq n - n'/3 > n/2$. Recall that the $(\alpha \upharpoonright \beta)$ -charge of $\mathcal{G}_n - \text{supp}(\beta)$ is odd since $\text{supp}(\beta)$ is \mathcal{G}_n -independent. Therefore $\text{closure}_{G, \alpha}(\beta)$ fixes e according to the unique assignment $b \in \mathbb{Z}_2$ so that the induced $(\alpha \upharpoonright \beta \circ e \rightarrow b)$ -charge of C_1 is odd and that of C_2 is even. We claim that $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ fixes e to the same value b . If not, $\rho^{(1)} \circ \text{closure}_{H^{(1)}, \alpha \upharpoonright \rho^{(1)}}(\pi)$ extends $(\beta \circ e \rightarrow \bar{b})$ where $(\beta \circ e \rightarrow \bar{b})$ does *not* push the contradiction of α into the giant component of $\mathcal{G}_n - (\text{supp}(\beta) \cup \{e\})$; by Fact 5.9 this contradicts (\dagger). This completes the proof of Lemma 7.4. \square

We are now ready to prove Lemma 7.2:

Proof of Lemma 7.2. Let $T^* := \text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T \upharpoonright \rho^{(i+1)})$. We first recall that “ $T^*(H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)})$ -represents $\vee_j \text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)})$ ” means:

$$\pi \in \text{Branches}_b(T^*) \implies (\vee_j \text{Disj}(\text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)}))) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = b.$$

We consider two cases depending on whether $b = 1$ or 0 . If $b = 1$, it suffices to argue that

$$\pi \in \text{Branches}_1(T^*) \implies \exists j: \text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 1.$$

Fix $\pi \in \text{Branches}_1(T^*)$. Since T^* has height at most $k = (\log n)/200d$ it follows from Fact 5.7 that $H^{(i+1)} - \text{closure}_{H^{(i+1)}}(\text{supp}(\pi))$ has a giant component. Hence we can apply Lemma 5.17 and get that

$$(T \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 1.$$

This implies that there exists $\pi' \in \text{Branches}_1(T)$ such that π' is a sub-restriction of $\rho^{(i+1)} \circ \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi)$.

Recall our assumption that $T(H^{(i)}, \alpha'')$ -represents $\vee_j T_j$. This means that

$$\pi' \in \text{Branches}_1(T) \implies \exists j: T_j \upharpoonright \text{closure}_{H^{(i)}, \alpha''}(\pi') = 1.$$

Fix such an index j . By Lemma 7.4, we have that $\text{closure}_{H^{(i)}, \alpha''}(\pi')$ is a sub-restriction of $\rho^{(i+1)} \circ \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi)$, and so we have that

$$(T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 1.$$

Observe that $T_j \upharpoonright \rho^{(i+1)}$ is total and has depth at most k . It follows from Lemma 5.16, Fact 5.6, and inspection of the Prune procedure that $\text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)})$ is total as well. Hence we may apply Fact 5.15 to obtain that

$$\text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 1,$$

which completes the proof of the $b = 1$ case.

For the $b = 0$ case, fix $\pi \in \text{Branches}_0(T^*)$. In this case it suffices to argue that for all j , either

$$\text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 0 \quad (5)$$

or

$$\text{Prune}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) \text{ is a 0-tree.} \quad (6)$$

Since T^* has height at most $k = (\log n)/200d$ it follows from Fact 5.7 that $H^{(i+1)}$ -closure $_{H^{(i+1)}}(\text{supp}(\pi))$ has a giant component. Hence we can apply Lemma 5.17 and get that

$$(T \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 0.$$

This implies that there exists $\pi' \in \text{Branches}_0(T)$ such that π' is a sub-restriction of $\rho^{(i+1)} \circ \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi)$.

Recall our assumption that T ($H^{(i)}, \alpha''$)-represents $\bigvee_j T_j$. This means that

$$\pi' \in \text{Branches}_0(T) \implies \forall j: \text{either } T_j \upharpoonright \text{closure}_{H^{(i)}, \alpha''}(\pi') = 0, \text{ or } T_j \upharpoonright \text{closure}_{H^{(i)}, \alpha''}(\pi') \text{ is a 0-tree.}$$

By Lemma 7.4, we have that $\text{closure}_{H^{(i)}, \alpha''}(\pi')$ is a sub-restriction of $\rho^{(i+1)} \circ \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi)$, and so we have that for all j , either

$$(T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) = 0 \text{ or } (T_j \upharpoonright \rho^{(i+1)}) \upharpoonright \text{closure}_{H^{(i+1)}, \alpha'' \upharpoonright \rho^{(i+1)}}(\pi) \text{ is a 0-tree.}$$

Again by Fact 5.15, it follows that either (5) or (6) holds, and the proof is complete. \square

8 The ‘‘atomic’’ Tseitin Switching Lemma

We begin in Section 8.1 by giving a simplified version of the switching lemma in which there is no underlying graph and no charge – it is for the most basic setting of having n Boolean variables x_1, \dots, x_n with no additional structure (e.g., the setting for proving depth lower bounds for polynomial-size circuits computing the n -variable Parity function). This simplified switching lemma is quantitatively somewhat weaker than Håstad’s switching lemma, yielding only an $\Omega(\sqrt{\log n})$ depth lower bound for Parity rather than the optimal $\Omega(\log n / \log \log n)$ bound that follows from Håstad’s switching lemma. However, the advantage of the new approach is that we are able to extend it (in Sections 8.2 through 8.5) to prove the actual Tseitin switching lemma over expander graphs which is our ultimate goal, whereas we were unable to extend previous known switching

lemma arguments to this complicated context. We also feel that some of the ideas underlying our expander switching lemma are much easier to follow in the simple setting of Section 8.1 than in the much more complex setting of Sections 8.2 through 8.5. (Note though that Section 8.1 serves more than just an expository role, since in it we establish some key definitions and technical results which are used later in our main expander switching lemma.)

8.1 A simple switching lemma for r -clipped decision trees

We introduce the following natural distribution over random walks down from the root of a decision tree T :

Definition 8.1 (Distribution $\mathcal{W}(T)$). *For a decision tree T , let $\mathcal{W}(T)$ be the probability distribution on $\text{Branches}(T)$ under which each $\pi \in \text{Branches}(T)$ has mass $2^{-|\pi|}$, where $|\pi|$ denotes the number of edges on the branch π . This corresponds to a uniform random walk down from the root of T . (If T has depth 0 (it is simply a constant), a draw from $\mathcal{W}(T)$ simply outputs the empty branch.)*

The following notion of an r -clipped decision tree is a key ingredient in our proof:

Definition 8.2 (r -clipped). *A decision tree T is r -clipped if every node in T has distance $\leq r$ from a leaf.*

The following lemma is simple but crucial for us:

Lemma 8.3 (r -DNF to r -clipped tree). *Every r -DNF is equivalent to an r -clipped decision tree.*

Proof. Build the decision tree term-by-term. Read the variables in the current term one-by-one, moving on to the next term once the current term is falsified. If the current term is ever satisfied, halt and output 1. If all terms are falsified, halt and output 1. It is clear that the resulting tree is r -clipped since at any internal node v , the branch that satisfies the current term terminates at a leaf at distance at most r from v . \square

Lemma 8.4 (Moment Bound). *If T is r -clipped, then $\mathbf{E}_{\pi \sim \mathcal{W}(T)} \binom{|\pi|}{s} \leq (20r2^r)^s$.*

Proof. Let $\mathbf{X} \in \{r, r+1, \dots\}$ be the first time that r consecutive heads come up in a sequence $\mathbf{C}_1, \mathbf{C}_2, \dots$ of i.i.d. unbiased coin flips. Note that \mathbf{X} stochastically dominates $|\pi|$.

We claim that $\mathbf{E}[\mathbf{X}^s] \leq (7rs2^r)^s$. Arguing by induction, assume that $\mathbf{E}[\mathbf{X}^i] \leq (7ri2^r)^i$ for all $i < s$. We have

$$\begin{aligned} \mathbf{E}[\mathbf{X}^s] &= \Pr[\mathbf{C}_1 = \dots = \mathbf{C}_r = \text{heads}]r^s + \sum_{k=1}^r \Pr[\mathbf{C}_1 = \dots = \mathbf{C}_{k-1} = \text{heads} \ \& \ \mathbf{C}_k = \text{tails}] \mathbf{E}[(k + \mathbf{X})^s] \\ &= \frac{r^s}{2^r} + \sum_{k=1}^r \frac{1}{2^k} \sum_{i=0}^s \binom{s}{i} k^i \mathbf{E}[\mathbf{X}^{s-i}] \\ &= \frac{r^s}{2^r} + \left(1 - \frac{1}{2^r}\right) \mathbf{E}[\mathbf{X}^s] + \sum_{i=1}^s \binom{s}{i} \mathbf{E}[\mathbf{X}^{s-i}] \sum_{k=1}^r \frac{k^i}{2^k}. \end{aligned}$$

Therefore,

$$\begin{aligned}
\mathbf{E}[\mathbf{X}^s] &= r^s + 2^r \sum_{i=1}^s \binom{s}{i} \mathbf{E}[\mathbf{X}^{s-i}] \sum_{k=1}^r \frac{k^i}{2^k} \\
&\leq r^s + 2^r \sum_{i=1}^s \left(\frac{e s}{i}\right)^i (7r(s-i)2^r)^{s-i} r^i \\
&\leq r^s + (rs2^r)^s 7^{s-1} \sum_{i=1}^s \left(\frac{e}{i}\right)^i \\
&\leq r^s + (rs2^r)^s 7^{s-1} 6 \\
&\leq (7rs2^r)^s.
\end{aligned}$$

Finally, we have

$$\mathbf{E}_{\pi \sim \mathcal{W}(T)} \binom{|\pi|}{s} \leq (e/s)^s \mathbf{E}_{\pi \sim \mathcal{W}(T)} [|\pi|^s] \leq (e/s)^s \mathbf{E}[\mathbf{X}^s] \leq (e/s)^s (7rs2^r)^s \leq (20r2^r)^s. \quad \square$$

We recall the standard notion of coordinate-wise independent random restrictions:

Definition 8.5 (Random restrictions \mathcal{R}_p). *For $p \in (0, 1)$, let \mathcal{R}_p be the distribution on restrictions ρ which independently sets each variable to $*$ with probability p and to $0, 1$ with probability $(1-p)/2$*

Now we are ready for our simple switching lemma, which states that any r -clipped decision tree — regardless of its depth — is unlikely to have large depth after it is hit by a random restriction.

Lemma 8.6 (Switching Lemma for r -Clipped Decision Trees). *Suppose T is an r -clipped decision tree. Then*

$$\Pr_{\rho \sim \mathcal{R}_p} [\text{depth}(T \upharpoonright \rho) \geq s] \leq (40pr2^r)^s.$$

Proof. For every restriction ρ , we have

$$\text{depth}(T \upharpoonright \rho) \geq s \implies \Pr_{\sigma \sim \mathcal{W}(T \upharpoonright \rho)} [|\sigma| \geq s] \geq 2^{-s}.$$

Using Markov's inequality and Lemma 8.4, we have

$$\begin{aligned}
\Pr_{\rho \sim \mathcal{R}_p} [\text{depth}(T \upharpoonright \rho) \geq s] &\leq \Pr_{\rho \sim \mathcal{R}_p} \left[\Pr_{\sigma \sim \mathcal{W}(T \upharpoonright \rho)} [|\sigma| \geq s] \geq 2^{-s} \right] \\
&\leq 2^s \mathbf{E}_{\rho \sim \mathcal{R}_p} \left[\Pr_{\sigma \sim \mathcal{W}(T \upharpoonright \rho)} [|\sigma| \geq s] \right] \\
&= 2^s \mathbf{E}_{\pi \sim \mathcal{W}(T)} \left[\Pr_{Y \sim \text{Bin}(|\pi|, p)} [Y \geq s] \right] \\
&\leq 2^s \mathbf{E}_{\pi \sim \mathcal{W}(T)} \left[p^s \binom{|\pi|}{s} \right] \\
&\leq (40pr2^r)^s. \quad \square
\end{aligned}$$

The crucial step in the above, marked by $=$, is justified by the following lemma:

Lemma 8.7. *Let T be a proper decision tree (no variable occurs twice on any path). The following two distributions are equivalent:*

1. $\mathcal{D}_1(T)$: Draw $\rho \sim \mathcal{R}_p$ and consider $T \upharpoonright \rho$. Output $\sigma \sim \mathcal{W}(T \upharpoonright \rho)$.
2. $\mathcal{D}_2(T)$: Draw $\pi = \langle \pi_1, \dots, \pi_k \rangle \sim \mathcal{W}(T)$. Output the sub-list of π obtained by going through π and independently including each element π_i with probability p .

Note that \mathcal{D}_1 corresponds to the LHS of the crucial $=$, and \mathcal{D}_2 to the RHS.

Proof. The proof is by induction on the depth of T . The base case is simple: if T has depth 0 then both $\mathcal{D}_1(T)$ and $\mathcal{D}_2(T)$ output the empty list with probability 1.

For the inductive step, let $T = (e; T_0, T_1)$. We first analyze the distribution $\mathcal{D}_1(T)$. By the definitions of $T \upharpoonright \rho$ and of \mathcal{W} , we have that for $\rho \sim \mathcal{R}_p$, the distribution of $\sigma \sim \mathcal{W}(T \upharpoonright \rho)$ is:

- (a) with probability $\frac{1-p}{2}$, σ is distributed as $\mathcal{W}(T_0 \upharpoonright \rho)$; by the inductive hypothesis applied to T_0 , this is the same distribution as $\langle (\text{draw from } \mathcal{D}_2(T_0)) \rangle$;
- (b) with probability $\frac{1-p}{2}$, σ is distributed as $\mathcal{W}(T_1 \upharpoonright \rho)$; by the inductive hypothesis applied to T_1 , this is the same distribution as $\langle (\text{draw from } \mathcal{D}_2(T_1)) \rangle$;
- (c) with probability $\frac{p}{2}$, σ is distributed as $\langle (e, 0), (\text{draw from } \mathcal{W}(T_0 \upharpoonright \rho)) \rangle$; by the inductive hypothesis applied to T_0 , this is the same distribution as $\langle (e, 0), (\text{draw from } \mathcal{D}_2(T_0)) \rangle$;
- (d) with probability $\frac{p}{2}$, σ is distributed as $\langle (e, 1), (\text{draw from } \mathcal{W}(T_1 \upharpoonright \rho)) \rangle$; by the inductive hypothesis applied to T_1 , this is the same distribution as $\langle (e, 1), (\text{draw from } \mathcal{D}_2(T_1)) \rangle$.

Now we analyze the distribution $\mathcal{D}_2(T)$. By inspection the distribution of a draw from $\mathcal{D}_2(T)$ is

- (e) with probability $\frac{1-p}{2}$, π is distributed as $\langle (\text{draw from } \mathcal{D}_2(T_0)) \rangle$;
- (f) with probability $\frac{1-p}{2}$, π is distributed as $\langle (\text{draw from } \mathcal{D}_2(T_1)) \rangle$;
- (g) with probability $\frac{p}{2}$, π is $\langle (e, 0), (\text{draw from } \mathcal{D}_2(T_0)) \rangle$;
- (h) with probability $\frac{p}{2}$, π is $\langle (e, 1), (\text{draw from } \mathcal{D}_2(T_1)) \rangle$;

Items (a)–(d) synch up exactly with (e)–(h) respectively, and the lemma is proved. \square

Using Lemma 8.3 we immediately obtain the following corollary:

Corollary 8.8 (Weak Switching Lemma for r -DNFs). *If F is an r -DNF, then*

$$\Pr_{\rho \sim \mathcal{R}_p} [F \upharpoonright \rho \text{ has decision-tree depth } \geq s] = (40pr2^r)^s.$$

With this “weak switching lemma” in hand, the standard argument that is used to get a lower bound for PARITY from a switching lemma yields the following:

Corollary 8.9 (Parity Lower Bound). *Depth- d circuits for PARITY $_n$ require size $\exp(\Omega((\log n)/d)^2)$. In particular, poly-size circuits for PARITY $_n$ require depth $\Omega(\sqrt{\log n})$.*

Proof sketch. Apply the weak switching lemma d times, with random restrictions $\rho_1, \dots, \rho_d \sim \mathcal{R}_p$, taking $p = n^{-1/2d}$ and $r = s = (\log n)/(4d)$, so that $(40pr2^r)^s = \exp(-\Omega((\log n)/d)^2)$, and using a union bound over the gates of a depth- d circuit computing PARITY $_n$ as in the usual argument. \square

8.2 A quantity $r_{\mathcal{H}_{n'',n'}}(m, s)$ that we need to bound

We define a function $r_{\mathcal{H}_{n'',n'}} : \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$ which will play a crucial role in the proof of our switching lemma:

$$r_{\mathcal{H}_{n'',n'}}(m, s) = \max_{X \subseteq E(\mathcal{G}_{n'}) : |X| \leq m} \Pr_{\mathbf{H} \sim \mathcal{H}_{n'',n'}}[X \text{ intersects at least } s \text{ super-edges in } \mathbf{H}]. \quad (7)$$

To prove the switching lemma we require an upper bound on $r_{\mathcal{H}_{n'',n'}}(m, s)$. The following theorem provides such a bound:

Theorem 3.

$$r_{\mathcal{H}_{n'',n'}}(m, s) \leq \binom{3n''/2}{s/29} \cdot \left(\frac{C(\ln n')^3}{n'} \right)^{s/29} \cdot \binom{m}{s}$$

where $C > 0$ is a universal constant.

The rest of Section 8.2 is devoted to proving Theorem 3.

Recall that an outcome of (H, H') from $\mathcal{H}_{n'',n'}$ specifies a set of bundles $B(i)$, $i \in [3n'/2]$, and that each such bundle corresponds to an edge in $\mathcal{G}_{n''}$ and vice versa. Let us say that a set of bundles is a *bundle matching* if the corresponding set of edges in $\mathcal{G}_{n''}$ is a matching. Write $\text{mbun}(H', X)$ to denote the size of the largest bundle matching in H' each of whose constituent bundles contains at least one edge of X , and let us define

$$\begin{aligned} \star'_{\mathcal{H}_{n'',n'}}(m, s) &:= \max_{X \subseteq E(\mathcal{G}_{n'}), |X|=m} \Pr_{(\mathbf{H}, \mathbf{H}') \sim \mathcal{H}_{n'',n'}}[\text{mbun}(\mathbf{H}', X) \geq s] \\ &= \max_{X \subseteq E(\mathcal{G}_{n'}), |X|=m} \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'',n'}}[\text{mbun}(\mathbf{H}', X) \geq s]. \end{aligned} \quad (8)$$

Now suppose that a set $X \subseteq E(\mathcal{G}_{n'})$ intersects at least s super-edges in H . Since $\mathcal{G}_{n''}$ is 3-regular, there must be a set of at least $s/29$ super-edges in H that X intersects, such that for any two of these super-edges, the corresponding edges in $\mathcal{G}_{n''}$ all have pairwise distance at least three from each other. (The “29” here comes from the fact that if e is an edge in a 3-regular graph, then there are at most 29 edges (including e itself) within distance at most two from e , where we say that two edges are at distance 0 from each other if they share a vertex.) It follows that the set of bundles that X intersects must contain a bundle matching of size at least $s/29$.² Consequently, we have that

Claim 8.10. $r_{\mathcal{H}_{n'',n'}}(m, s) \leq \star'_{\mathcal{H}_{n'',n'}}(m, s/29)$.

So it suffices to upper bound $\star'_{\mathcal{H}_{n'',n'}}(m, s)$; we do this below. For conciseness henceforth we simply write \star' for $\star'_{\mathcal{H}_{n'',n'}}$.

8.2.1 Bounding $\star'(m, s)$

Fix an $X \subseteq E(\mathcal{G}_{n'})$ with $|X| = m$ that achieves the maximum in (8). There are $\binom{m}{s}$ ways to choose a particular subset $X' \subseteq X$, $|X'| = s$; fix one such X' and denote its elements e_1, \dots, e_s . We may upper bound $\star'(m, s)$ by

$$\star'(m, s) \leq \binom{m}{s} \cdot \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'',n'}}[\text{a bundle matching in } \mathbf{H}' \text{ hits all of } e_1, \dots, e_s].$$

²Note that we could not conclude this if we only had a pairwise distance lower bound of two.

Order all the edges $\{a_{1,1}, a_{1,2}\}, \dots, \{a_{3n''/2,1}, a_{3n''/2,2}\}$ of $\mathcal{G}_{n''}$ in some fixed canonical way. Recall that there is a bundle in H' corresponding to each such edge. There are at most $\binom{3n''/2}{s}$ ways to choose s edges in $\mathcal{G}_{n''}$ that could form a matching. Fix a subset $1 \leq i_1 < \dots < i_s \leq 3n''/2$ of size s . We may upper bound $\star'(m, s)$ as

$$\star'(m, s) \leq \binom{m}{s} \cdot \binom{3n''/2}{s} \cdot \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'', n'}}[\text{bundle } \mathbf{B}(i_j) \text{ hits } e_j \text{ for all } j = 1, \dots, s]. \quad (9)$$

Below we shall focus on on

$$\Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'', n'}}[\text{bundle } \mathbf{B}(i_j) \text{ hits } e_j \text{ for all } j = 1, \dots, s]. \quad (10)$$

Remark 4. *It is certainly not the case that we can treat each bundle independently of the others. This is the case for at least two reasons: first, the bundles share endpoints and this destroys independence. Beyond this, even the choice of the bundle endpoints is not independent because of the conditioning in step (1) of the draw of $\mathbf{H} \sim \mathcal{H}$, which stipulates that no two vertices v_i, v_j are too close to one another. Indeed, if we were to draw n'' vertices independently and uniformly at random, the probability that they would satisfy all $\binom{n''}{2}$ of the desired pairwise distance lower bounds would be at most*

$$\left(1 - \frac{(\ln n')^{\Theta(1)}}{n'}\right)^{\Theta(n''^2)},$$

which is extremely small (recall that $n'' = n'/2^{(\log n')^c}$).

Nevertheless, regarding the first reason, intuition might suggest that the fact that bundles can share endpoints should not have a major effect; and regarding the second reason, intuitively the conditioning based on pairwise distance between the v_i 's is fairly "mild" since in a uniform draw of the n'' vertices most pairs would indeed satisfy the desired distance condition. Thus one might hope to achieve a bound similar to that which would be obtained if the bundles could be handled independently (and indeed we will rigorously establish such a bound below).

To gain intuition for what "should happen", let us briefly digress and analyze (10) pretending that the bundles were all mutually independent. In this case the analysis would be very simple: we would have that (10) equals

$$\Pr[\text{bundle } \mathbf{B}(1) \text{ hits edge } e_1]^s.$$

There are at most $\ell := \kappa_2(\ln n')^3$ many edges in bundle \mathbf{B}_1 , so this is at most

$$\left(\sum_{j=1}^{\ell} \Pr[\text{the } j\text{-th edge in bundle } \mathbf{B}(i) \text{ hits edge } e_1]\right)^s.$$

Each edge in $\mathcal{G}_{n'}$ is equally likely to be the j -th edge in bundle $\mathbf{B}(i)$, so the above is

$$\left(\frac{2\kappa_2(\ln n')^3}{3n'}\right)^s.$$

This would give an upper bound on (9) of

$$\binom{m}{s} \cdot \binom{3n''/2}{s} \cdot \left(\frac{2\kappa_2(\ln n')^3}{3n'}\right)^s. \quad (11)$$

So a bound of this form is what we will shoot for (and achieve) in the actual analysis below.

Let us return to the actual analysis of (10). Recall that each $a_{i_j,b}$ (where $j \in [s]$ and $b \in \{1,2\}$) is an element of $[n'']$ (an element of $V(\mathcal{G}_{n''})$) and that all $2s$ such elements are distinct (since we have a matching). We rewrite the quantity we want to bound,

$$(10) = \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'',n'}}[\text{bundle } \mathbf{B}(i_j) \text{ hits } e_j \text{ for all } j = 1, \dots, s],$$

as $\prod_{j=1}^s p_j$, where

$$p_j = \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'',n'}}[\text{bundle } \mathbf{B}(i_j) \text{ hits edge } e_j \mid \text{bundle } \mathbf{B}(i_{j'}) \text{ hits edge } e_{j'} \text{ for } j' = 1, \dots, j-1]. \quad (12)$$

We take an alternate view of the draw of $\mathbf{H}' \sim \mathcal{H}'_{n'',n'}$ as follows:

- 1'. For $j = a_{i_1,1}, a_{i_1,2}, a_{i_2,1}, a_{i_2,2}, \dots, a_{i_s,1}, a_{i_s,2}$ (in that specific order) do the following:
 - Pick a uniform random vertex $\mathbf{v}_j \in [n']$ conditioned on \mathbf{v}_j having distance at least $3\kappa_1 \ln \ln n'$ in $\mathcal{G}_{n'}$ from all of the \mathbf{v}_i 's that have already been picked. (This vertex \mathbf{v}_j corresponds to vertex j of $\mathcal{G}_{n''}$.)
 - For each pair $\{\mathbf{v}_i, \mathbf{v}_j\}$ such that $\{i, j\}$ is an edge in $E(\mathcal{G}_{n''})$ and \mathbf{v}_i has already been picked, construct the bundle of paths between \mathbf{v}_i and \mathbf{v}_j as described earlier. Denote this bundle of paths by $\mathbf{B}(i, j)$.
- 2'. Finish the draw for the remaining possibilities of $j \in [n'']$ (in an arbitrary order). I.e. for the remaining $j \in [n'']$, in an arbitrary order,
 - Pick a uniform random vertex $\mathbf{v}_j \in [n']$ conditioned on \mathbf{v}_j having distance at least $3\kappa_1 \ln \ln n'$ in $\mathcal{G}_{n'}$ from all of the \mathbf{v}_i 's that have already been picked. (This vertex \mathbf{v}_j corresponds to vertex j of $\mathcal{G}_{n''}$.)
 - For each pair $\{\mathbf{v}_i, \mathbf{v}_j\}$ such that $\{i, j\}$ is an edge in $E(\mathcal{G}_{n''})$ and \mathbf{v}_i has already been picked, construct the bundle of paths between \mathbf{v}_i and \mathbf{v}_j as described earlier. Denote this bundle of paths by $\mathbf{B}(i, j)$.

Remark 5. *The difference between this description of a draw of $\mathbf{H}' \sim \mathcal{H}'_{n'',n'}$ and our original one is that in the original description the draw of \mathbf{H}' is generated by first selecting all the locations of the vertices (going in the ordering $1, \dots, n''$), and only then selecting all the bundles. In the new description, we go through the vertices $\{1, \dots, n''\}$ selecting their locations in a different ordering (prioritizing the vertices that are pertinent for (10)), and when we select a vertex we immediately select all the bundles joining that vertex to previously-selected vertices.*

Fix any $\ell \in [s]$; our goal is to bound p_ℓ . We do this by analyzing the conditional probability that bundle $\mathbf{B}(a_{i_\ell,1}, a_{i_\ell,2})$ hits edge e_ℓ where we condition on a more restrictive event than the one specified in (12). Let S denote the set $\{a_{i_1,1}, a_{i_1,2}, \dots, a_{i_{\ell-1},1}, a_{i_{\ell-1},2}\}$ (note that this is a subset of $[n'']$). Fix any distinct elements $(r_t)_{t \in S}$ in $[n']$ and let $\Phi((r_t)_{t \in S})$ denote the event “ $(\mathbf{v}_t = r_t)_{t \in S}$, and bundle $\mathbf{B}(a_{i_{j'},1}, a_{i_{j'},2})$ hits edge $e_{j'}$ for $j' = 1, \dots, \ell - 1$.” Our goal is to give a uniform upper bound on the probability

$$q_\ell((r_t)_{t \in S}) := \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'',n'}}[\text{bundle } \mathbf{B}(a_{i_\ell,1}, a_{i_\ell,2}) \text{ hits edge } e_\ell \mid \Phi((r_t)_{t \in S})] \quad (13)$$

that holds for all possible $(r_t)_{t \in S}$. Since the event conditioned on in (12), “bundle $\mathbf{B}(i_{j'})$ hits edge $e_{j'}$ for $j' = 1, \dots, j - 1$,” can be partitioned into disjoint events $\Phi((r_t)_{t \in S})$ across all possibilities for $(r_t)_{t \in S}$, such an upper bound implies the same upper bound on p_ℓ .

Recall that we are working with a matching, so all $2s$ elements $a_{i_j, b}$ ($j \in [s], b \in \{1, 2\}$) are distinct. Moreover, for any $k < \ell$, given the endpoints $v_{a_{i_k, 1}}$ and $v_{a_{i_k, 2}}$, the actual draw of the bundle $\mathbf{B}(a_{i_k, 1}, a_{i_k, 2})$ is independent from everything else that happens in the draw of $\mathbf{H}' \sim \mathcal{H}'_{n'', n'}$. This implies that conditioning on $\Phi'((r_t)_{t \in S})$ is equivalent to conditioning on $\Phi((r_t)_{t \in S})$ for the purpose of analyzing $q_\ell((r_t)_{t \in S})$, where $\Phi'((r_t)_{t \in S})$ is the event “ $(v_t = r_t)_{t \in S}$.” So we have that $q'_\ell((r_t)_{t \in S}) = q_\ell((r_t)_{t \in S})$, where

$$q'_\ell((r_t)_{t \in S}) := \Pr_{\mathbf{H}' \sim \mathcal{H}'_{n'', n'}}[\text{bundle } \mathbf{B}(a_{i_\ell, 1}, a_{i_\ell, 2}) \text{ hits edge } e_\ell \mid \Phi'((r_t)_{t \in S})].$$

(Intuitively, in passing from $q_\ell((r_t)_{t \in S})$ to $q'_\ell((r_t)_{t \in S})$, we are only discarding the conditioning on whether bundles joining up some of the earlier vertices hit the e_j edges, but we are keeping the conditioning on where the endpoints of those earlier bundles were located. This endpoint location information is the only relevant information for the ℓ -th stage, because the location of the endpoints of the earlier bundles is the only thing affecting the distribution of the location of the endpoints of the ℓ -th stage bundles, and the distribution of the paths comprising the ℓ -th stage bundles depends only on the location of their endpoints.)

We have reduced the problem to that of bounding $q'_\ell((r_t)_{t \in S})$. We first observe that

$$\begin{aligned} q'_\ell((r_t)_{t \in S}) &\leq (\ln n')^2 \cdot \Pr[\text{the first path } P \text{ in bundle } \mathbf{B}(a_{i_\ell, 1}, a_{i_\ell, 2}) \text{ hits edge } e_\ell \mid \Phi'((r_t)_{t \in S})] \\ &\leq (\ln n')^2 \cdot \Pr[\text{the walk } \mathbf{W} \text{ (recall Step 2(b)) corresponding to the first path } P \\ &\quad \text{in bundle } \mathbf{B}(a_{i_\ell, 1}, a_{i_\ell, 2}) \text{ hits edge } e_\ell \mid \Phi'((r_t)_{t \in S})] \\ &\leq (\ln n')^2 \cdot \sum_{u \in [2\kappa_2 \ln n']} \Pr[\text{the } u\text{th edge in the first walk } \mathbf{W} \text{ in bundle } \mathbf{B}(a_{i_\ell, 1}, a_{i_\ell, 2}) \\ &\quad \text{hits edge } e_\ell \mid \Phi'((r_t)_{t \in S})] \end{aligned} \tag{14}$$

Now fix any $u \in [2\kappa_2 \ln n']$. (It suffices to deal only with the first half of the walk \mathbf{W} because we can deal with the second half, starting from the other endpoint, symmetrically.) We would like to upper bound the quantity

$$\Pr[\text{the } u\text{th edge in the first walk } \mathbf{W} \text{ in bundle } \mathbf{B}(a_{i_\ell, 1}, a_{i_\ell, 2}) \text{ hits edge } e_\ell \mid \Phi'((r_t)_{t \in S})]. \tag{15}$$

Observe that the randomness here is over

- the draw of $a_{i_\ell, 1}$ and $a_{i_\ell, 2}$ given $\Phi'((r_t)_{t \in S})$;
- the (uniform over $[n']$) draw of the “midpoint” \mathbf{x} of the walk \mathbf{W} (recall Step 2(b) of the description of how \mathbf{H}' is drawn from $\mathcal{H}'_{n'', n'}$); and
- the draw of the first half-walk (call it \mathbf{W}') of \mathbf{W} in bundle $\mathbf{B}(a_{i_\ell, 1}, a_{i_\ell, 2})$.

Similar to the proof of Lemma 2.7, we note that given the outcome of $a_{i_\ell, 1}$, the combined draw of \mathbf{x} and of the half-walk \mathbf{W}' ending at \mathbf{x} does *not* result in \mathbf{W}' being perfectly distributed as a uniform $(\kappa_2 \ln n')$ -step random walk starting from $a_{i_\ell, 1}$, because \mathbf{W}' is constrained to end at the uniform random vertex \mathbf{x} and the distribution of the endpoint of a $(\kappa_2 \ln n')$ -step random walk

starting from $a_{i_\ell,1}$ will not be perfectly uniform random. However, as in that earlier argument, the variation distance between the distribution of \mathbf{W}' and the distribution of a truly random walk is extremely small — at most $(n')^{-100}$ for a suitable choice of the constant κ_2 , because the variation distance between a truly uniform random $\mathbf{x} \in [n']$ and the endpoint of a $(\kappa_2 \ln n')$ -length random walk is at most $(n')^{-100}$. So in what follows, we shall analyze (15) under the assumption that \mathbf{W}' is a truly uniform random $(\kappa_2 \ln n')$ -step random walk starting from $a_{i_\ell,1}$ (we refer to the corresponding probability as (15')); adding $(n')^{-100}$ to the resulting upper bound obtained under this assumption, we get a legitimate upper bound on the actual value of (15).

We can upper bound (15') by upper bounding the probability that the $(u-1)$ -th vertex reached (after the $(u-1)$ -th step of the random walk) is one of the two endpoints of e_ℓ . For $u=1$, we observe that the conditioning $\Phi'((r_t)_{t \in S})$ has the effect of ruling out at most $(2s-2) \cdot 3^{3\kappa_1 \ln \ln n'} = s \cdot \text{polylog}(n')$ many possibilities for $a_{i_\ell,1}$ and $a_{i_\ell,2}$ out of the n' possibilities, so there are still at least $0.99n'$ many possibilities for $a_{i_\ell,1}$. Hence for $u=1$ we have that $\alpha_u := \Pr[\text{the } (u-1)\text{-th vertex reached after the } (u-1)\text{-th step of the random walk is one of the two endpoints of } e_\ell] \leq 2/(0.99n')$. In fact, we claim that for every $u \in \{0, 1, \dots, \kappa_2 \ln n'\}$, the value α_u is at most $2/(0.99n')$. This follows from the following elementary lemma:

Lemma 8.11. *Let G be a regular graph and let $S \subset V(G)$. Let $p_{S,k}(v)$ denote the probability that a random walk that starts at a uniform random vertex of S reaches vertex v at the k -th step. Then $p_{S,k}(v) \leq 1/|S|$ for all $v \in V$.*

Proof. We will prove this by induction on k . The base case $k=0$ is clearly true. For the induction, we will use the following self-evident fact (note that regularity is essential for this fact to be true):

Fact 8.12. *Consider a random walk on a d -regular graph with any distribution on initial vertices. For any $v \in V(G)$, let $p_k(v)$ denote the probability that the walk reaches vertex v at the k -th step. Then*

$$p_k(v) = \frac{1}{d} \sum_{u \sim v} p_{k-1}(u) \leq \max_{u \sim v} \{p_{k-1}(u)\}.$$

By induction, the max on the RHS of Fact 8.12 is always $1/|S|$, and so $p_{S,k}(v) \leq 1/|S|$ as desired. \square

All that remains is to retrace our steps and combine the various bounds. We have established that (15') is at most $2/(0.99n')$, so as discussed above (15) $\leq 2/(0.99n') + (n')^{-100} < 4/n'$. Hence for (14) we get that

$$q'_\ell((r_t)_{t \in S}) \leq 8\kappa_2(\ln n')^3/n'$$

for all possible outcomes of $((r_t)_{t \in S})$; as discussed earlier this gives

$$p_\ell \leq 8\kappa_2(\ln n')^3/n'.$$

Recalling (9), (10) and that (10) $= \prod_{j=1}^s p_j$, we get that $\star'(m, s)$ is at most

$$\star'(m, s) \leq \binom{m}{s} \cdot \binom{3n''/2}{s} \cdot \left(\frac{8\kappa_2(\ln n')^3}{n'} \right)^s$$

(note that this is essentially as good as the idealized bound (11) from Remark 4). Putting the final piece in place by recalling Claim 8.10, we have established Theorem 3.

8.3 The Tseitin switching lemma

Having proved Theorem 3 upper bounding $r_{\mathcal{H}_{n'',n'}}$, we are almost at the statement of 8.15. We need the following notion of how setting a single edge on an H -super-edge “forces” an assignment to the entire super-edge. (Existence is immediate in the following definition, and uniqueness follows easily from the fact that an H -super-edge is simply a path of degree-2 vertices.)

Definition 8.13 (Super-edge inference). *Fix $(\rho, H) \in \text{supp}(\mathcal{F}_{n'',n'})$ and let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. For $e \in E(H)$ and $b \in \mathbb{Z}_2$ we write $\Psi_H(\alpha' \upharpoonright \rho, e \rightarrow b) \in \mathbb{Z}_2^{\text{super}_H(e)}$ to denote the unique assignment to $\text{super}_H(e)$ extending $(e \rightarrow b)$ so that $\alpha' \upharpoonright (\rho \circ \Psi_H(\alpha' \upharpoonright \rho, e \rightarrow b))$ is α' -consistent.*

Next we introduce a notion of “super-edge pruning” which is crucial for our switching lemma. Super-edge pruning can be viewed as a relaxed form of the pruning that was defined in Section 5.3. Roughly speaking, an internal node e of a tree T is shortcut in “regular” pruning if it lies in the closure of its ancestors. In contrast, an internal node e of T is shortcut in super-edge pruning if and only one of its ancestors e' is contained in the same super-edge as e (note that in this event certainly e is in the closure of its ancestors). The idea is that any assignment to e' “forces” a unique assignment along the entire super-edge containing e' so as not to have a contradiction at any path node in that super-edge.

Definition 8.14 (H -super-edge pruning). *Fix $(\rho, H) \in \text{supp}(\mathcal{F}_{n'',n'})$ and let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. Let T be a decision tree over $E(H)$. The H -super-edge pruning of T under $\alpha' \upharpoonright \rho$, denoted $\text{SuperEdgePrune}_{H,\alpha' \upharpoonright \rho}(T)$, is the decision tree obtained from T by shortcutting internal nodes e such that $\text{super}_H(e) = \text{super}_H(e')$ for some $e' \in E(H)$ which is an ancestor of e in T . The shortcutting of e is done according to $\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow b)(e) \in \mathbb{Z}_2$ where $b \in \mathbb{Z}_2$ denotes the subtree of e' in T that e belongs to.*

Lemma 8.15 (Switching lemma). *Let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge and T be an r -clipped $\mathcal{G}_{n'}$ -independent decision tree over $E(\mathcal{G}_{n'})$. Then for all s ,*

$$\Pr_{(\rho, \mathbf{H})} [\text{depth}(\text{SuperEdgePrune}_{H,\alpha' \upharpoonright \rho}(T \upharpoonright \rho)) \geq s] \leq \binom{3n''/2}{s/29} \cdot \left(\frac{8\kappa_2(\ln n')^3}{n'} \right)^{s/29} \cdot (20r2^r)^s,$$

where $(\rho, \mathbf{H}) \sim \mathcal{F}_{n'',n'}$.

Proof. We have

$$\begin{aligned}
& \Pr_{(\rho, \mathbf{H})} [\text{depth}(\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(T \upharpoonright \rho) \geq s)] \\
& \leq \Pr_{(\rho, \mathbf{H})} \left[\Pr_{\sigma \sim \mathcal{W}(\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho})} [|\sigma| \geq s] \geq 2^{-s} \right] && \text{(Definition of } \mathcal{W}(\cdot) \text{)} \\
& \leq 2^s \mathbf{E}_{(\rho, \mathbf{H})} \left[\Pr_{\sigma \sim \mathcal{W}(\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho})} [|\sigma| \geq s] \right] && \text{(Markov's inequality)} \\
& = 2^s \mathbf{E}_{\mathbf{H}} \left[\Pr_{\pi \sim \mathcal{W}(T)} [\pi \text{ intersects at least } s \text{ super-edges in } \mathbf{H}] \right] && \text{(key equivalence lemma)} \\
& = 2^s \mathbf{E}_{\pi \sim \mathcal{W}(T)} \left[\Pr_{\mathbf{H}} [\pi \text{ intersects at least } s \text{ super-edges in } \mathbf{H}] \right] && \text{(trivial)} \\
& \leq 2^s \mathbf{E}_{\pi \sim \mathcal{W}(T)} r_{\mathcal{H}_{n'', n'}}(|\pi|, s) && \text{(Definition of } r(\cdot, \cdot) \text{)} \\
& \leq \binom{3n''/2}{s/29} \cdot \left(\frac{8\kappa_2(\ln n')^3}{n'} \right)^{s/29} \cdot \mathbf{E}_{\pi \sim \mathcal{W}(T)} \binom{|\pi|}{s} && \text{(Theorem 3)} \\
& \leq \binom{3n''/2}{s/29} \cdot \left(\frac{8\kappa_2(\ln n')^3}{n'} \right)^{s/29} \cdot (20r2^r)^s. && \text{(Lemma 8.4)}
\end{aligned}$$

□

It remains to state and prove the key equivalence lemma (this is analogous to Lemma 8.7, but much more involved) underlying the first equality step above. Note that this equivalence is a “pointwise” one which does not involve an expectation over the choice of \mathbf{H} — it holds for every possible outcome of \mathbf{H} .

Definition 8.16. Fix $H \in \text{supp}(\mathcal{H}_{n'', n'})$ and let $\pi = \langle \pi_1, \dots, \pi_k \rangle \in (E(H) \times \mathbb{Z}_2)^k$. We define $\text{SuperEdgePrune}_H(\pi)$ to be the sublist of π obtained by greedily removing all $\pi_j = (e, b)$ such that $\pi_i = (e', b')$ for some $i < j$ and $\text{super}_H(e) = \text{super}_H(e')$.

Here is the key equivalence lemma:

Lemma 8.17 (Key equivalence lemma). Let T be a $\mathcal{G}_{n'}$ -independent decision tree over $E(\mathcal{G}_{n'})$ and $\alpha' \in \mathbb{Z}_2^{\mathcal{V}_{n'}}$ be an odd charge. Fix $H \in \text{supp}(\mathcal{H}_{n'', n'})$. The following two distributions \mathcal{D}_1 and \mathcal{D}_2 are equivalent:

1. $\mathcal{D}_1(T)$: Let (ρ, \mathbf{H}) be distributed as a draw from $\mathcal{F}_{n'', n'}$ conditioned on $\mathbf{H} = H$. Output $\sigma \sim \mathcal{W}(\text{SuperEdgePrune}_{H, \alpha' \upharpoonright \rho}(T \upharpoonright \rho))$.
2. $\mathcal{D}_2(T)$: Draw $\pi = \langle \pi_1, \dots, \pi_k \rangle \sim \mathcal{W}(T)$. Let $\tilde{\pi}$ be the sublist of π where we discard all $\pi_i = (e, b)$ where $e \in E(\mathcal{G}_{n'}) - E(H)$. Output $\text{SuperEdgePrune}_H(\tilde{\pi})$.

It is clear that Lemma 8.17 gives the equality that is labeled by “key equivalence lemma” above; that equality simply randomizes over \mathbf{H} . Thus, to complete the proof of the Switching Lemma it remains only to prove the “key equivalence lemma”; the rest of Section 8 is devoted to its proof.

8.4 Auxiliary lemma for the proof of Lemma 8.17

Below is the main auxiliary lemma we will need for Lemma 8.17. Intuitively, when we apply this lemma we are at some internal labeled by variable e of a $\mathcal{G}_{n'}$ -independent tree T . In both cases in Lemma 8.18 below, J corresponds to the set of “pioneering” $E(H)$ -ancestors of e in T , where a “pioneering” ancestor is one which, when it was queried, was the first edge in its super-edge to be queried along that branch. In Part 1, I corresponds to e along with the set of all ancestors of its in T that are in $E(\mathcal{G}_{n'}) - E(H)$ (i.e. I is fixed by ρ). In Part 2, I corresponds to the set of all ancestors of e in T (but now not including e) that are in $E(\mathcal{G}_{n'}) - E(H)$, and the set L corresponds to e along with its all its ancestors that belong to the same super-edge but are not pioneering.

(Note that Part (3) of Lemma 8.18 is not required for the proof of Lemma 8.17; however, we will use it in Section 9 later.)

Lemma 8.18. *Let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge and fix $H \in \text{supp}(\mathcal{H}_{n'',n'})$. Let ρ be distributed as $(\rho, \mathbf{H}) \sim \mathcal{F}_{n'',n'}$ conditioned on $\mathbf{H} = H$. Fix $I \subseteq E(\mathcal{G}_{n'}) - E(H)$, and $J \subseteq E(H)$ where J contains at most one edge from each super-edge of H . Suppose $I \cup J$ is $\mathcal{G}_{n'}$ -independent.*

1. For all $y \in \mathbb{Z}_2^I$,

$$\Pr[\rho(I) = y] = 2^{-|I|}.$$

2. Let $e' \in J$ and L be an ordered list of distinct elements of $\text{super}_H(e') \setminus \{e'\}$, where $I \cup J \cup L$ is $\mathcal{G}_{n'}$ -independent. Then for all $y \in \mathbb{Z}_2^I$, $z \in \mathbb{Z}_2^J$, and $u \in \mathbb{Z}_2^L$, we have that

$$\Pr[\rho(I) = y \text{ and } (\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow z(e')))(L) = u] = 2^{-(|I|+|L|)}.$$

3. For each $e' \in J$ let $L_{e'}$ be an ordered list of distinct elements of $\text{super}_H(e') \setminus \{e'\}$, where $I \cup J \cup \bigcup_{e' \in J} L_{e'}$ is $\mathcal{G}_{n'}$ -independent. Then for all $y \in \mathbb{Z}_2^I$, $z \in \mathbb{Z}_2^J$, and $(u_{e'} \in \mathbb{Z}_2^{L_{e'}})_{e' \in J}$, we have that

$$\Pr[\rho(I) = y \text{ and } (\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow z(e')))((L_{e'})_{e' \in J}) = (u_{e'} \in \mathbb{Z}_2^{L_{e'}})_{e' \in J}] = 2^{-|I|} \cdot \prod_{e' \in J} 2^{-|L_{e'}|}.$$

Remark 6. *Note that (1) and (2) above correspond to two kinds of shortcutting that go on in $\text{SuperEdgePrune}_{H, \alpha' \upharpoonright \rho}(T \upharpoonright \rho)$. The first corresponds to the straightforward shortcutting that occurs when a tree T is hit by a restriction ρ to yield $T \upharpoonright \rho$, and the second corresponds to the shortcutting that happens in $\text{SuperEdgePrune}_{H, \alpha' \upharpoonright \rho}(\cdot)$.*

We give some setup for the proof of Lemma 8.18. Recall that ρ is draw uniformly at random among all α' -consistent assignments to $E(\mathcal{G}_{n'}) - E(H)$. Let C_1, \dots, C_ℓ be the connected components of the graph $\mathcal{G}_{n'} \setminus H := (V(\mathcal{G}_{n'}) \setminus V(H), E(\mathcal{G}_{n'}) \setminus E(H))$. Let $\mathcal{G}_{n',i}$ be the connected subgraph of $\mathcal{G}_{n'}$ obtained by taking all edges incident to $V(C_i)$ (and the vertices on both endpoints of these edges of course). For each $i \in [\ell]$ let $\alpha'_i \in \mathbb{Z}_2^{V(C_i)}$ be α' restricted to the vertices in $V(C_i)$. Note that:

- $V(C_i) \subsetneq V(\mathcal{G}_{n',i})$, and so $\mathcal{G}_{n',i}$ and α'_i satisfy the conditions of Fact 2.4: the set of all α'_i -consistent assignments to $E(\mathcal{G}_{n',i})$ form an affine subspace. Moreover, every vertex $v \in V(\mathcal{G}_{n',i}) \setminus V(C_i)$ lies in $V(H)$ and has degree 2 in H (i.e. v is a non-real-vertex of H).
- $E(\mathcal{G}_{n'})$ is the disjoint union of $E(\mathcal{G}_{n',1}), \dots, E(\mathcal{G}_{n',\ell})$ and $E(H)$.

- Furthermore, for every $i \neq j$ we have that $V(\mathcal{G}_{n',i})$ and $V(\mathcal{G}_{n',j})$ are disjoint (this is a consequence of $\mathcal{G}_{n'}$ having degree 3). So $V(\mathcal{G}_{n'})$ is the disjoint union of $V(\mathcal{G}_{n',1}), \dots, V(\mathcal{G}_{n',\ell})$ and the “real vertices” of $V(H)$, and $\mathcal{G}_{n',1}, \dots, \mathcal{G}_{n',\ell}$ are the connected components of $\mathcal{G}_{n'} - E(H)$.

By the third bullet above, we get the following:

Fact 8.19. *A uniform random α' -consistent assignment $\rho \in \mathbb{Z}_2^{E(\mathcal{G}_{n'}) - E(H)}$ can be generated by independently generating $\rho_i \in \mathbb{Z}_2^{E(\mathcal{G}_{n',i})}$, a uniformly random α'_i -consistent assignment to $E(\mathcal{G}_{n',i})$, for each $i \in [\ell]$.*

We recall the following elementary linear-algebraic fact concerning affine subspaces:

Fact 8.20. *Let $S \subseteq \mathbb{Z}_2^n$ be an affine subspace. Fix $T \subseteq [n]$, and suppose that for every $y \in \mathbb{Z}_2^T$ there exists $z \in S$ such that $y_i = z_i$ for all $i \in T$. Then in fact for all $y \in \mathbb{Z}_2^T$ we have that:*

$$\Pr_{\mathbf{x} \in S}[\mathbf{x}_i = y_i \text{ for all } i \in T] = 2^{-|T|}.$$

Proof of Lemma 8.18. We begin with Part 1 of the claim. For each $i \in [\ell]$ let $I_i = I \cap E(\mathcal{G}_{n',i})$ and note that I is the disjoint union of I_1, \dots, I_ℓ . By Fact 8.19, we have that for all $y \in \mathbb{Z}_2^I$,

$$\Pr[\rho(I) = y] = \prod_{i=1}^{\ell} \Pr[\rho_i(I_i) = y_{I_i}].$$

By Facts 2.4 (the set of all α'_i -consistent assignments to $E(\mathcal{G}_{n',i})$ form an affine subspace) and Fact 8.20 it then suffices to show that for every $w \in \mathbb{Z}_2^{I_1}$ there exists an extension $\tilde{w} \in \mathbb{Z}_2^{E(\mathcal{G}_{n',1})}$ of w that is α'_1 -consistent, and likewise for $w' \in \mathbb{Z}_2^{I_2}, w'' \in \mathbb{Z}_2^{I_3}$ and so on. Equivalently, for every $w \in \mathbb{Z}_2^{I_1}$ there is an $(\alpha'_1 \upharpoonright w)$ -consistent assignment to $E(\mathcal{G}_{n',1}) \setminus I_1$. We consider two cases depending on whether or not $\mathcal{G}_{n',1} - I_1$ is connected:

- If $\mathcal{G}_{n',1} - I_1$ is connected, then by Fact 2.3 there is an $(\alpha'_1 \upharpoonright w)$ -consistent assignment to $E(\mathcal{G}_{n',1}) \setminus I_1$ (since $(\alpha'_1 \upharpoonright w) \in \mathbb{Z}_2^{V(C_1)}$ and $V(C_1) \subsetneq V(\mathcal{G}_{n',1} - I_1) = V(\mathcal{G}_{n',1})$).
- Otherwise, if $\mathcal{G}_{n',1} - I_1$ is not connected we let its components be $\mathcal{G}_{n',1,1}, \dots, \mathcal{G}_{n',1,k}$ for some $k \geq 2$.
 - If every such component contains a vertex in $V(\mathcal{G}_{n',1}) - V(C_1)$, then again we are done by Fact 2.3.
 - The remaining case is that some component (without loss of generality $\mathcal{G}_{n',1,1}$) has $V(\mathcal{G}_{n',1,1}) \subseteq V(C_1)$. We have that $\mathcal{G}_{n',1,1}$ is a connected component of $\mathcal{G}_{n'} - (E(H) \cup I_1)$ and that this graph contains at least one other connected component $\mathcal{G}_{n',1,2}$. Recalling that C_1 is incident to no edges in $E(H)$ and that $V(\mathcal{G}_{n',1,1}) \subseteq V(C_1)$, it follows that $\mathcal{G}_{n',1,1}$ is a connected component in $\mathcal{G}_{n'} - I_1$. But this contradicts the fact that I_1 is a $\mathcal{G}_{n'}$ -independent set (since its removal disconnects $\mathcal{G}_{n',1,1}$ from the rest of $\mathcal{G}_{n'}$), so this case cannot occur.

An identical argument applies to $w' \in \mathbb{Z}_2^{I_2}, w'' \in \mathbb{Z}_2^{I_3}$ and so on. This completes the proof of Part 1 of Lemma 8.18.

Next we turn to Part 2. Let H^* be H with the non-real vertices of $\text{super}(e')$ and edges of $\text{super}(e')$ removed. Again consider the connected components of $\mathcal{G}_{n'} - H^*$. Note that $\{e'\} \cup L$ all belong to the same connected component. Define $I^* = I \cup \{e'\} \cup L$, and observe that I^* is $\mathcal{G}_{n'}$ -independent. Recall that the statement to be proved is that:

$$\Pr[\rho(I) = y \text{ and } (\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow z(e')))(L) = u] = 2^{-(|I|+|L|)}.$$

Let y^* be y augmented by fixing e' to $z(e')$ and fixing L according to $\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow z(e'))$. Let ρ^* be a uniform random α' -consistent assignment to $E(\mathcal{G}_{n'}) - E(H^*)$. Note that

$$\Pr_{\alpha'\text{-consistent } \rho} [\rho(I) = y \text{ and } (\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow z(e')))(L) = u] = \Pr_{\alpha'\text{-consistent } \rho^*} [\rho^*(I^*) = y^*],$$

and so we can apply the argument from Part 1 with H^* in place of H , I^* in place of I , y^* in place of y , and ρ^* in place of ρ to obtain the desired result.

Part 3 follows from the same argument as Part 2 with trivial modifications. \square

8.5 Proof of Lemma 8.17

We obtain Lemma 8.17 as a special case of a more general statement which is better suited to an inductive proof. To make this more general statement we need an extension of the notion of H -super-edge pruning of T under ρ .

Definition 8.21 (H -super-edge (J, z) -pruning of T under ρ). *Fix $(\rho, H) \in \text{supp}(\mathcal{F}_{n'', n'})$ and let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. Let T be a decision tree over $E(H)$. Let $J \subseteq E(H)$ such that J contains at most one edge from each super-edge of H , and J does not contain any edges queried in T . Let $z \in \mathbb{Z}_2^J$.*

The H -super-edge (J, z) -pruning of T under ρ , denoted $\text{SuperEdgePrune}_{H, \alpha' \upharpoonright \rho}(T, J \rightarrow z)$, is the decision tree obtained from T by shortcutting internal nodes e such that either of the following (mutually exclusive) conditions holds:

1. e' is an ancestor of e in T for some $e' \in E(H)$ where $\text{super}_H(e) = \text{super}_H(e')$. Or,
2. $e' \in J$ for some $e' \in E(H)$ where $\text{super}_H(e) = \text{super}_H(e')$.

In Case 1 the shortcutting of e is done according to $\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow b)(e) \in \mathbb{Z}_2$ where $b \in \mathbb{Z}_2$ denotes the subtree of e' in T that e belongs to (exactly as in Definition 8.14), and in Case 2 it is done according to $\Psi_H(\alpha' \upharpoonright \rho, e' \rightarrow z(e'))(e)$.

Lemma 8.17 follows directly from the following Lemma 8.22 taking $I = J = \emptyset$.

(For intuition on the following lemma, it may be helpful to think of T as a subtree of a larger $\mathcal{G}_{n'}$ -independent tree T' , and $I \sqcup J \sqcup K$ as a partition of the edges queried so far along the partial branch leading from the root of T' to the root of T . All the $(E(\mathcal{G}_{n'}) - E(H))$ -nodes of T' encountered on this partial branch are stored in I , and all the pioneering $E(H)$ -nodes encountered are stored in J ; K contains the remaining nodes (non-pioneering nodes labeled by elements of $E(H)$). The settings of the I, J and K variables given by the partial branch leading from the root of T' to T are captured by y, z and u respectively. The setting of I by y and J by z will affect how we simplify T .)

Lemma 8.22. *Let T be a $\mathcal{G}_{n'}$ -independent decision tree and $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$. Fix $H \in \text{supp}(\mathcal{H}_{n'',n'})$, and let $J \subseteq E(H)$ contain at most one edge from each super-edge of H . Then the following two distributions \mathcal{D}_1 and \mathcal{D}_2 are equivalent:*

- $\mathcal{D}_1(T)$: *Let $I, K \subseteq (E(\mathcal{G}_{n'}) \setminus J)$ be such that*
 - $I \subseteq E(\mathcal{G}_{n'}) \setminus E(H)$;
 - $K \subseteq E(H)$ and for each $e \in K$ there is some $e' \in J$ such that $\text{super}_H(e) = \text{super}_H(e')$;
 - For every branch B in T , $I \cup J \cup K \cup B$ is $\mathcal{G}_{n'}$ -independent and $(I \cup J \cup K) \cap B = \emptyset$.

Let $y \in \mathbb{Z}_2^I$, $z \in \mathbb{Z}_2^J$ and $u \in \mathbb{Z}_2^K$.

A draw from $\mathcal{D}_1(T)$ is generated as follows: Let $(\boldsymbol{\rho}, \mathbf{H})$ be distributed as a draw from $\mathcal{F}_{n'',n'}$ conditioned on (i) $\mathbf{H} = H$ (ii) $\boldsymbol{\rho}(I) = y$, and (iii) for every $e' \in J$, it is the case that u restricted to the coordinates $e \in K$ that have $\text{super}_H(e) = \text{super}_H(e')$ agrees with $\Psi_H(\alpha' \upharpoonright \boldsymbol{\rho}, e' \rightarrow z(e'))$. Output $\boldsymbol{\sigma} \sim \mathcal{W}(\text{SuperEdgePrune}_{H,\alpha' \upharpoonright \boldsymbol{\rho}}(T \upharpoonright \boldsymbol{\rho}, J \rightarrow z))$.

- $\mathcal{D}_2(T)$: *Draw $\boldsymbol{\pi} = \langle \boldsymbol{\pi}_1, \dots, \boldsymbol{\pi}_k \rangle \sim \mathcal{W}(T)$. Let $\tilde{\boldsymbol{\pi}}$ be the sublist of $\boldsymbol{\pi}$ where we discard all $\boldsymbol{\pi}_i = (e, \mathbf{b})$ where either (a) $e \in E(\mathcal{G}_{n'}) - E(H)$ or (b) $e' \in J$ for some $e' \in E(H)$ where $\text{super}_H(e) = \text{super}_H(e')$. Output $\text{SuperEdgePrune}_H(\tilde{\boldsymbol{\pi}})$.*

The proof is by induction on the depth of T . The base case is that T is a depth-0 tree. In this case both $\mathcal{D}_1(T)$ and $\mathcal{D}_2(T)$ output the empty set of super-edges with probability 1.

For the inductive step, let $T = (e; T_0, T_1)$ be a $\mathcal{G}_{n'}$ -independent decision tree, and let I, J, K, y, z, u be as in the statement of the lemma.

We first analyze the distribution $\mathcal{D}_1(T)$. For $(\boldsymbol{\rho}, \mathbf{H})$ distributed as in the statement of the lemma, $T \upharpoonright \boldsymbol{\rho}$ is distributed as follows:

- If $e \in E(\mathcal{G}_{n'}) \setminus E(H)$ (and hence e is fixed by $\boldsymbol{\rho}$), then $T \upharpoonright \boldsymbol{\rho} = T_0 \upharpoonright \boldsymbol{\rho}$ with probability 1/2, and $T \upharpoonright \boldsymbol{\rho} = T_1 \upharpoonright \boldsymbol{\rho}$ with probability 1/2. (This uses Part 1 of Lemma 8.18 taking its “ I ” to be “the current I ” $\cup \{e\}$.)
- If $e \in E(H)$ (hence e is left unfixed by $\boldsymbol{\rho}$), then $T \upharpoonright \boldsymbol{\rho}$ is distributed as $(e; T_0 \upharpoonright \boldsymbol{\rho}, T_1 \upharpoonright \boldsymbol{\rho})$.

Building on this, a draw from $\mathcal{D}_1(T)$ is distributed as follows:

1. Suppose $e \in E(\mathcal{G}_{n'}) \setminus E(H)$ (and hence e is fixed by $\boldsymbol{\rho}$). As above, by Part 1 of Lemma 8.18, taking its “ I ” to be “the current I ” $\cup \{e\}$, we get that for each $b \in \mathbb{Z}_2$, with probability 1/2 a draw from $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha' \upharpoonright \boldsymbol{\rho}}(T \upharpoonright \boldsymbol{\rho}, J \rightarrow z))$ is distributed according to $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha' \upharpoonright \boldsymbol{\rho}'}(T_b \upharpoonright \boldsymbol{\rho}', J \rightarrow z))$ where $(\boldsymbol{\rho}', \mathbf{H})$ is distributed as a draw from $\mathcal{F}_{n'',n'}$ conditioned on (i) $\mathbf{H} = H$, (ii) $\boldsymbol{\rho}'(I) \circ \boldsymbol{\rho}'(e) = y \circ b$, and (iii) for every $e' \in J$, it is the case that u restricted to the coordinates $e'' \in K$ that have $\text{super}_H(e'') = \text{super}_H(e')$ agrees with $\Psi_H(\alpha' \upharpoonright \boldsymbol{\rho}', e' \rightarrow z(e'))$.

For each $b \in \mathbb{Z}_2$, we apply the inductive hypothesis to T_b , taking the “ I ” in the inductive hypothesis to be “the current I ” $\cup \{e\}$, taking the “ J ” (respectively, “ K ”) in the inductive hypothesis to be “the current J ” (respectively, “the current K ”), taking the “ y ” in the inductive hypothesis to be “the current y , extended by mapping e to b ”, and taking

the “ z ” (respectively, u) to be “the current z ” (respectively, “the current u ”). By the inductive hypothesis applied in this way, for each $b \in \mathbb{Z}_2$ with probability $1/2$ a draw from $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha'|\rho}(T \upharpoonright \rho, J \rightarrow z))$ is distributed according to $\mathcal{D}_2(T_b)$ with its “ J ” being “the current J ”.

2. Suppose $e \in E(H)$ has a “cohort” edge in J (i.e. $\text{super}_H(e) = \text{super}_H(e')$ for some $e' \in J$). Now by Part 2 of Lemma 8.18, taking its “ I ” to be “the current I ”, its “ J ” to be “the current J ”, and taking its “ L ” to be “the current $(K \cap \text{super}_H(e)) \cup \{e\}$ ”, we get that for each $b \in \mathbb{Z}_2$ with probability $1/2$ a draw from $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha'|\rho}(T \upharpoonright \rho, J \rightarrow z))$ is distributed according to $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha|\rho'}(T_b \upharpoonright \rho', J \rightarrow z))$ where (ρ', H) is distributed as a draw from $\mathcal{F}_{n'',n'}$ conditioned on (i) $\mathbf{H} = H$, (ii) $\rho'(I) = y$, and (iii) for every $e' \in J$, it is the case that u restricted to the coordinates $e'' \in K$ that have $\text{super}_H(e'') = \text{super}_H(e')$ agrees with $\Psi_H(\alpha' \upharpoonright \rho', e' \rightarrow z(e'))$ and moreover $\Psi_H(\alpha' \upharpoonright \rho', e' \rightarrow z(e'))(e) = b$.

For each $b \in \mathbb{Z}_2$ we apply the inductive hypothesis to T_b , taking the “ I ” (respectively, “ J ”) in the inductive hypothesis to be “the current I ” (respectively, “the current J ”), taking the “ K ” in the inductive hypothesis to be “the current K ” $\cup \{e\}$, taking the “ y ” (respectively, “ z ”) in the inductive hypothesis to be “the current y ” (respectively, “the current z ”) and taking the “ u ” in the inductive hypothesis to be “the current u , extended by mapping e to b ”. By the inductive hypothesis applied in this way, for each $b \in \mathbb{Z}_2$ with probability $1/2$ a draw from $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha'|\rho}(T \upharpoonright \rho, J \rightarrow z))$ is distributed according to $\mathcal{D}_2(T_b)$ with its “ J ” being “the current J ”.

3. The remaining case is that $e \in E(H)$ and e does not have any cohort edges in J . (Note that this is the only case in which an internal node survives in the simplified tree.) In this case, simply by the unbiasedness of the random walk, for each $b \in \mathbb{Z}_2$ with probability $1/2$ a draw from $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha'|\rho}(T \upharpoonright \rho, J \rightarrow z))$ is distributed according to $(e, b) \circ \mathcal{W}(\text{SuperEdgePrune}_{H,\alpha'|\rho'}(T_b \upharpoonright \rho', J \circ e \rightarrow z \circ b))$ where (ρ', \mathbf{H}) is distributed as a draw from $\mathcal{F}_{n'',n'}$ conditioned on (i) $\mathbf{H} = H$, (ii) $\rho'(I) = y$, and (iii) for every $e' \in J$, it is the case that u restricted to the coordinates $e'' \in K$ that have $\text{super}_H(e'') = \text{super}_H(e')$ agrees with $\Psi_H(\alpha' \upharpoonright \rho', e' \rightarrow z(e'))$.

For each $b \in \mathbb{Z}_2$ we apply the inductive hypothesis to T_b , taking the “ I ” (respectively, “ K ”) in the inductive hypothesis to be “the current I ” (respectively, “the current K ”), taking the “ J ” in the inductive hypothesis to be “the current J ” $\cup \{e\}$, taking the “ y ” (respectively, “ u ”) in the inductive hypothesis to be “the current y ” (respectively, “the current u ”) and taking the “ z ” in the inductive hypothesis to be “the current $z \circ b$ ”. By the inductive hypothesis applied in this way, for each $b \in \mathbb{Z}_2$ with probability $1/2$ a draw from $\mathcal{W}(\text{SuperEdgePrune}_{H,\alpha'|\rho}(T \upharpoonright \rho, J \rightarrow z))$ is distributed according to $(e, b) \circ \hat{\pi}_b$ where $\hat{\pi}_b$ is drawn from $\mathcal{D}_2(T_b)$ with its “ J ” being “the current J ” $\cup \{e\}$.

Now we consider $\mathcal{D}_2(T)$. The distribution of $\text{SuperEdgePrune}_H(\tilde{\pi})$ is as follows:

- A. If $e \in E(\mathcal{G}_{n'}) \setminus E(H)$: since a random walk proceeds left and right with equal probability, the distribution of $\text{SuperEdgePrune}_H(\tilde{\pi})$ is that for each $b \in \mathbb{Z}_2$ with probability $1/2$ it is distributed according to $\mathcal{D}_2(T_b)$ with its “ J ” being “the current J ”. This matches exactly (1.) above.

- B. If $e \in E(H)$ has a cohort edge in J : the distribution of $\text{SuperEdgePrune}_H(\tilde{\pi})$ is that for each $b \in \mathbb{Z}_2$ with probability $1/2$ it is distributed according to $\mathcal{D}_2(T_b)$ with its “ J ” being “the current J ”. This matches exactly (2.) above.
- C. If $e \in E(H)$ has no cohort edge in J : the distribution of $\text{SuperEdgePrune}_H(\tilde{\pi})$ is that for each $b \in \mathbb{Z}_2$ with probability $1/2$ it is distributed according to $(e, b) \circ \hat{\pi}_b$ where $\hat{\pi}_b$ is drawn from $\mathcal{D}_2(T_b)$ with its “ J ” being “the current J ” $\cup \{e\}$. This matches exactly (3.) above.

This concludes the proof of Lemma 8.17. □

9 Safe trees become total under KR random restrictions

We start with the following observation:

Lemma 9.1. *For any tree T over $E(\mathcal{G}_{n'})$ and any odd charge $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$, the tree $\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)$ is $(n'/(3\lambda))$ -safe, i.e. it has no occurrences of \perp at any depth $\leq n'/(3\lambda)$.*

Proof. Fix $\pi \in \text{Branches}_\perp(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T))$. By Lemma 5.16, $\mathcal{G}_{n'} - \text{closure}_{\mathcal{G}_{n'}}(\text{supp}(\pi))$ does not have a giant component, and hence $|\text{supp}(\pi)| \geq n'/(3\lambda)$ by Fact 5.6. □

Thus the structure of a tree $\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)$ is as follows: At depths $< n'/(3\lambda)$ it may have leaves labeled by 0 or 1 but not by \perp , and at depths $\geq n'/(3\lambda)$ it may have leaves labeled by any element of $\mathbb{Z}_2 \cup \{\perp\}$. However, very deep branches (of depth $\geq n'/(3\lambda)$) that are labeled by 0 or 1 will be problematic for us later (intuitively, because $\mathcal{G}_n -$ (the closure of such branches) may not have a giant component). To circumvent such problems we introduce the following “lopping” operator on decision trees (which is useful later for technical reasons).

Definition 9.2. *Let T be a decision tree (total or partial) and $m \geq 0$. $\text{Lop}_m(T)$ is the decision tree of depth at most m that is obtained by removing each node at depth m in T (and of course all its children) and replacing the node by \perp .*

With this definition in hand, recall that Lemma 8.15 guarantees that if an s -clipped, $\mathcal{G}_{n'}$ -independent decision tree T over $E(\mathcal{G}_{n'})$ is hit with a random restriction $(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}$, with high probability we have that $\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(T \upharpoonright \rho)$ does not have large depth (hence likewise $\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho)$ does not have large depth). In this section we show that with high probability additionally $\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho)$ is a total decision tree — even though all nodes at depth m were (i.e. it does not have any leaves labeled \perp). The following lemma (the main result of this section) gives us this:

Lemma 9.3. *Let T be an s -clipped, m -safe, $\mathcal{G}_{n'}$ -independent partial decision tree over $E(\mathcal{G}_{n'})$ where $m \leq n'/(3\lambda)$, and let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. Then*

$$\Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho) \text{ is not total}] \leq \left(1 - \frac{1}{2^{s+1}}\right)^m, \quad (16)$$

provided that

$$\left(\frac{C_1 n'' (\ln n')^3}{sn'}\right)^{1/29} < \frac{1}{2^{s+1}},$$

where C_1 is a universal constant.

We use the following technical lemma:

Lemma 9.4. *If T is an s -clipped decision tree then T has at most $O(2^s \cdot \nu^m)$ nodes at depth m , where $\nu < 2 - 1/2^s$.*

Proof. Fix s . The Fibonacci numbers of degree s are given by the sequence

$$\begin{aligned} F_0^{(s)} &= 0, F_1^{(s)} = 1, F_j^{(s)} = 2^{j-2}, \quad j \in \{2, \dots, s+1\}; \\ F_j^{(s)} &= F_{j-1}^{(s)} + F_{j-2}^{(s)} + \dots + F_{j-s}^{(s)}, \quad j \geq s+2. \end{aligned}$$

By Lemma 1 of [Cap90], the number of nodes at depth m in an r -clipped decision tree is at most $2F_{m+1}^{(s)}$ which is known [Con] to be $O(2^s \cdot \nu^m)$, where $\nu < 2 - 1/2^s$ is the unique positive root of $x + x^{-s} = 2$ that is real and greater than 1. \square

Proof of Lemma 9.3. Fix any branch π in $\text{Branches}(\text{Lop}_m(T))$ whose leaf node ℓ is at depth m and hence is labeled by \perp (these are the only \perp leaves in $\text{Lop}_m(T)$ by Lemma 9.1). By Lemma 9.4, to prove (16), it suffices to bound

$$\Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\ell \text{ survives in } \text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho)]. \quad (17)$$

We have that

$$\begin{aligned} & \Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\ell \text{ survives in } \text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho)] \\ &= \sum_{s=0}^m \Pr_{\mathbf{H} \sim \mathcal{H}_{n'', n'}} [\pi \text{ intersects exactly } s \text{ super-edges of } \mathbf{H}] \\ & \cdot \Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\ell \text{ survives in } \text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho) \mid \pi \text{ intersects exactly } s \text{ super-edges of } \mathbf{H}]. \end{aligned}$$

Fix $s \in \{0, \dots, m\}$. For the first probability, recall that

$$\Pr_{\mathbf{H} \sim \mathcal{H}_{n'', n'}} [\pi \text{ intersects exactly } s \text{ super-edges of } \mathbf{H}] \leq r_{\mathcal{H}_{n'', n'}}(m, s).$$

For the second probability fix any $H \in \text{supp}(\mathcal{H}_{n'', n'})$ such that π intersects exactly s super-edges of H . Define I to be the subset of $E(\mathcal{G}_{n'}) \setminus E(H)$ that occur in π , and $y \in \mathbb{Z}_2^I$ to be their assignments under π . Let $J \subseteq E(H)$ denote the set of pioneering $E(H)$ -edges in π and $z \in \mathbb{Z}_2^J$ denote their assignments under π . For each $e' \in J$ let $L_{e'}$ denote the edges in $\text{super}_H(e')$ that occur in π , and let $u_{e'} \in \mathbb{Z}_2^{L_{e'}}$ denote the assignment of those edges under π . Observe that $|I \cup J \cup \bigcup_{e' \in J} L_{e'}| = m$ and $|J| = s$. By part (3) of Lemma 8.18, we have that

$$\Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\ell \text{ survives in } \text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(T) \upharpoonright \rho) \mid \mathbf{H} = H] = 2^{-(m-s)}.$$

Combining the above bounds we get that

$$(17) \leq \sum_{s=0}^m r(m, s)(1/2)^{m-s},$$

and hence for a suitable absolute constant c_1 , we have

$$\begin{aligned}
(\text{LHS of (16)}) &\leq O(2^s) \cdot (2 - 1/2^s)^m \cdot \sum_{s=0}^m r(m, s) (1/2)^{m-s} \\
&= O(1) \cdot (1 - 1/2^s)^m \cdot \sum_{s=0}^m 2^{2s} r(m, s) \\
&\leq O(1) \cdot (1 - 1/2^s)^m \cdot \sum_{s=0}^m \left(2^{2s} \cdot \binom{3n''/2}{s/29} \cdot \left(\frac{C(\ln n')^3}{n'} \right)^{s/29} \right) \cdot \binom{m}{s} \quad (\text{Theorem 3}) \\
&\leq (1 - 1/2^s)^m \cdot \sum_{s=0}^m \left(\left(\frac{C_1 n'' (\ln n')^3}{sn'} \right)^{1/29} \right)^s \cdot \binom{m}{s} \\
&= \left(\left(1 - \frac{1}{2^s} \right) \cdot \left(1 + \left(\frac{C_1 n'' (\ln n')^3}{sn'} \right)^{1/29} \right) \right)^m \\
&\leq \left(1 - \frac{1}{2^{s+1}} \right)^m,
\end{aligned}$$

provided that

$$\left(\frac{C_1 n'' (\ln n')^3}{sn'} \right)^{1/29} < \frac{1}{2^{s+1}},$$

where C_1 is a universal constant. \square

10 Final Tseitin switching lemma

We are now ready to prove our final switching lemma, restated from Section 7 for the reader's convenience:

Theorem 2. (Final switching lemma) *Let $0 \leq i \leq d < d^*$. Fix $(\rho^{(i)}, H^{(i)}) \in \text{supp}(\mathcal{A}^{(i)})$. Let T_1, \dots, T_M be $(k, H^{(i)})$ -good decision trees over $E(H^{(i)})$ where $k = (\log n)/200d$ and let $\alpha' = \alpha \upharpoonright \rho^{(i)}$. Then except with failure probability at most $n^{-\varepsilon(\log n)/d^2}$ over a draw of $(\rho^{(i+1)}, \mathbf{H}^{(i+1)}) \sim \mathcal{A}^{(i+1)}(\rho^{(i)}, H^{(i)})$, there is a $(k, \mathbf{H}^{(i+1)})$ -good decision tree over $E(\mathbf{H}^{(i+1)})$ that $(\mathbf{H}^{(i+1)}, \alpha' \upharpoonright \rho^{(i+1)})$ -represents $\vee_j \text{Prune}_{\mathbf{H}^{(i+1)}, \alpha' \upharpoonright \rho^{(i+1)}}(T_j \upharpoonright \rho^{(i+1)})$.*

We establish Theorem 2 by proving the following:

Theorem 7. *Let $0 \leq i \leq d < d^*$, and define $n' = \tau(i, n)$ and $n'' = \tau(i+1, n)$. Let T_1, \dots, T_M be $(k, \mathcal{G}_{n'})$ -good decision trees over $E(\mathcal{G}_{n'})$ where $k = (\log n)/200d$ and let $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be any odd charge. Then except with failure probability at most $n^{-\varepsilon(\log n)/d^2}$ over a draw of (ρ, \mathbf{H}) from $\mathcal{F}_{n'', n'}$, there is a (k, \mathbf{H}) -good decision tree over $E(\mathbf{H})$ that $(\mathbf{H}, \alpha' \upharpoonright \rho)$ -represents $\vee_j \text{Prune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho)$.*

The proof of Theorem 2 proceeds along essentially identical lines to that of Theorem 7 by the correspondence (established in Section 4) between $H^{(i)}$ and $\mathcal{G}_{n'}$ ($H^{(i)}$ is a topological embedding of $\mathcal{G}_{n'}$ in \mathcal{G}_n where $n' = \tau(i, n)$) and between $\mathcal{A}^{(i)}$ and $\mathcal{F}_{n'', n'}$ (a draw $(\rho^{(i+1)}, \mathbf{H}^{(i+1)}) \sim \mathcal{A}^{(i+1)}$ is defined in terms of $\mathcal{F}_{n'', n'}$ where $n' = \tau(i, n)$ and $n'' = \tau(i+1, n)$).³

³In a bit more detail, a full proof of Theorem 2 would use analogues of Lemmas 9.1 and 9.3 for graphs $H^{(i)}$ instead of $\mathcal{G}_{n'} = \mathcal{G}_{\tau(i, n)}$; these analogues would have λ^* in place of λ since they would employ Fact 5.7 in place of Fact 5.6. It would also have $H^{(i)}$ -analogues of Sections 10.1 and 10.2 below; these analogues are straightforward.

As stated in Theorem 7, we emphasize that throughout the rest of this section k is fixed to be $\log n/(200d)$.

10.1 Useful tools for Theorem 7

We will use the following straightforward fact:

Fact 10.1. *Let T_1, \dots, T_M be depth- k decision trees over $E(\mathcal{G}_{n'})$. There is a k -clipped decision tree T over $E(\mathcal{G}_{n'})$ satisfying the following: for every $\rho \in \mathbb{Z}_2^S$ where $S \subseteq E(\mathcal{G}_{n'})$, $b \in \mathbb{Z}_2$, and $\pi \in \mathbb{Z}_2^{S'}$ where $S' \subseteq E(\mathcal{G}_{n'}) - S$,*

$$(T \upharpoonright \rho) \upharpoonright \pi = b \implies (\bigvee_j \text{Disj}(T_j \upharpoonright \rho)) \upharpoonright \pi = b.$$

Proof. Observe that $\bigvee_j \text{Disj}(T_j)$ is a k -DNF; T is the k -clipped decision tree obtained from this k -DNF as in the proof of Lemma 8.3. Fix any S, S', ρ, π as in the statement of the fact. To prove the desired implication suppose first that $b = 1$, so we have $(T \upharpoonright \rho) \upharpoonright \pi = 1$, or equivalently $T \upharpoonright \rho \circ \pi = 1$. This means that there is some $j \in [M]$ such that $T_j \upharpoonright \rho \circ \pi = 1$, or equivalently $(T_j \upharpoonright \rho) \upharpoonright \pi = 1$. Hence we have that $(\bigvee_j \text{Disj}(T_j \upharpoonright \rho)) \upharpoonright \pi = 1$ as desired.

Next suppose that $(T \upharpoonright \rho) \upharpoonright \pi = 0$, i.e. $T \upharpoonright \rho \circ \pi = 0$. It follows that for each $j \in [M]$ we have that either $T_j \upharpoonright \rho \circ \pi = 0$, or $T_j \upharpoonright \rho \circ \pi$ is a 0-tree. Consequently for each $j \in [M]$ we have that either $(T_j \upharpoonright \rho) \upharpoonright \pi = 0$, or $(T_j \upharpoonright \rho) \upharpoonright \pi$ is a 0-tree. As a result $(\bigvee_j \text{Disj}(T_j \upharpoonright \rho)) \upharpoonright \pi = 0$, and the proof is complete. \square

For the rest of this section, fix $m = n^{1/3}$. The following lemma combines the main results of Sections 8 and 9, recalling that $n' = \tau(i, n)$ and $n'' = \tau(i + 1, n)$:

Lemma 10.2. *Let T be any k -clipped decision tree over $E(\mathcal{G}_{n'})$ and $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. Then we have*

$$\begin{aligned} & \Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\text{SuperEdgePrune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)) \upharpoonright \rho) \text{ is total and has depth } \leq k] \\ & \geq 1 - \left(C \left(\frac{(\log n)^3}{k n^{1/(2d)}} \right)^{1/29} \cdot k 2^k \right)^k - \left(1 - \frac{1}{2^{k+1}} \right)^m \\ & \geq 1 - 2 \left(C \left(\frac{(\log n)^3}{k n^{1/(2d)}} \right)^{1/29} \cdot k 2^k \right)^k \\ & \geq 1 - n^{-\varepsilon(\log n)/d^2} \end{aligned}$$

for a universal constant $\varepsilon > 0$.

Proof. Observe that $\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)$ is $\mathcal{G}_{n'}$ -independent, and since T is k -clipped it is easy to see that $\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T))$ is also k -clipped for any m . Thus we may apply Lemma 8.15, recalling that $n''/n' = n^{-1/(2d)}$ and that $\log n' = \Theta(\log n)$ to get the first subtrahend in the first inequality. We further have that $\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)$ is m -safe (by Lemma 9.1, with room to spare, recalling that $n^{1/3} \ll n'/(3\lambda)$),⁴ and by our choice of $k = (\log n)/200d$ we may apply Lemma 9.3 to get the second subtrahend. The second inequality follows by observing that the second subtrahend is extremely tiny. The third inequality holds by our choice of k . \square

⁴For the Theorem 2 analogue, observe that we also have $n^{1/3} \ll n'/(3\lambda^*)$.

A simple consequence of this lemma is the following corollary:

Corollary 10.3 (Prune \prec SuperEdgePrune). *Let T be any k -clipped decision tree over $E(\mathcal{G}_{n'})$ and $\alpha' \in \mathbb{Z}_2^{V(\mathcal{G}_{n'})}$ be an odd charge. Then we have*

$$\begin{aligned} & \Pr_{(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}} [\text{Prune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)) \upharpoonright \rho) \text{ is total and has depth } \leq k] \\ & \geq 1 - n^{-\varepsilon(\log n)/d^2} \end{aligned}$$

for a universal constant $\varepsilon > 0$.

The corollary follows from the fact that for any $H \in \text{supp}(\mathcal{H}_{n'', n'})$, any odd charge α'' , and any tree T' over $E(H)$, $\text{Prune}_{H, \alpha''}(T')$ is a sub-tree of $\text{SuperEdgePrune}_{H, \alpha''}(T')$. A formal proof of this fact is by a straightforward induction on the depth of T' ; intuitively, SuperEdgePrune shortcuts a subset of the tree nodes that Prune shortcuts. The key observation is that whenever SuperEdgePrune shortcuts an internal node e in the tree T' , it must be the case that e is a bridge in the current graph (because of the removal of the earlier edge e' in that same super-edge).

We are now ready to prove Theorem 7.

10.2 Proof of Theorem 7

Let T_1, \dots, T_M be as in the theorem statement, and consider the k -clipped decision tree T given by Fact 10.1. Since $\text{Prune}_{\mathcal{G}_{n'}, \alpha}(T)$ remains k -clipped (as is easily verified), we may apply Corollary 10.3 to get that with probability at least $1 - n^{-\varepsilon(\log n)/d^2}$ over $(\rho, \mathbf{H}) \sim \mathcal{F}_{n'', n'}$, $\text{Prune}_{\mathbf{H}, \alpha' \upharpoonright \rho}(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)) \upharpoonright \rho)$ is \mathbf{H} -independent, total, and has depth at most k ; in other words, it is (k, \mathbf{H}) -good. Fix any such outcome (ρ, H) and let

$$T^* := \text{Prune}_{H, \alpha' \upharpoonright \rho}(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)) \upharpoonright \rho);$$

we claim that $T^*(H, \alpha' \upharpoonright \rho)$ -represents $\bigvee_j \text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho)$. (Recall that “ $T^*(H, \alpha' \upharpoonright \rho)$ -represents $\bigvee_j \text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho)$ ” means:

$$\pi \in \text{Branches}_b(T^*) \implies (\bigvee_j \text{Disj}(\text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho))) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = b.$$

Fix $b \in \mathbb{Z}_2$ and $\pi \in \text{Branches}_b(T^*)$. By Lemma 5.17, we have that

$$(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha}(T)) \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = b \tag{18}$$

We claim that in fact

$$(T \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = b \tag{19}$$

Since $(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha}(T)) \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = b$, it follows that there exists

$$\pi' \in \text{Branches}_b(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T)))$$

such that $\rho \circ \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi)$ extends π' . Furthermore, since $\text{supp}(\pi')$ is a $\mathcal{G}_{n'}$ -independent set of size at most m (and hence $\mathcal{G}_{n'} - \text{closure}_{\mathcal{G}_{n'}}(\text{supp}(\pi'))$ has a giant component by Fact 5.6) we may apply Lemma 5.17 (since $\text{Branches}_b(\text{Lop}_m(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T))) \subseteq \text{Branches}_b(\text{Prune}_{\mathcal{G}_{n'}, \alpha'}(T))$) to get that

$$T \upharpoonright \text{closure}_{\mathcal{G}_{n'}, \alpha'}(\pi') = b. \tag{20}$$

Next, applying the $i = 0$ case of Lemma 7.4 (checking that ρ, H, π and π' satisfy its conditions), we get that $\rho \circ \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi)$ extends $\text{closure}_{\mathcal{G}_{n', \alpha}}(\pi')$ and so indeed (19) holds.

Given (19), we may apply Fact 10.1 to get that

$$((\bigvee_j \text{Disj}(T_j \upharpoonright \rho)) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi)) = b. \quad (21)$$

We claim that this implies that

$$(\bigvee_j \text{Disj}(\text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho))) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = b, \quad (22)$$

noting that this would complete the proof of Theorem 7. To establish (22), we consider two cases depending on whether $b = 1$ or 0 . If $b = 1$ there must exist some j such that $(T_j \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = 1$. Observe that $T_j \upharpoonright \rho$ is total and has depth at most k . It follows from Lemma 5.16, Fact 5.6, and inspection of the Prune procedure that $\text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho)$ is total as well. Hence we may apply Fact 5.15 to obtain that $\text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = 1$, and hence (22) holds with $b = 1$. It remains to deal with the $b = 0$ case. In this case, for all j we have that either $(T_j \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = 0$, or $(T_j \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi)$ is a 0-tree. Again by Fact 5.15, it follows that either $\text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi) = 0$, or $\text{Prune}_{H, \alpha' \upharpoonright \rho}(T_j \upharpoonright \rho) \upharpoonright \text{closure}_{H, \alpha' \upharpoonright \rho}(\pi)$ is a 0-tree. We conclude that (22) holds with $b = 0$, and the proof is complete.

Acknowledgements

We thank Noga Alon and Julia Chuzhoy for helpful discussions. We also thank Mladen Mikša for pointing out an error in an earlier version of this paper and suggesting a fix.

References

- [Ajt94] Miklós Ajtai. The complexity of the pigeonhole principle. *Combinatorica*, 14(4):417–433, 1994. [1](#), [1](#)
- [Ale11] Michael Alekhovich. Lower Bounds for k-DNF Resolution on Random 3-CNFs. *Computational Complexity*, 20(4):597–614, 2011. [1](#)
- [BFU94] Andrei Broder, Alan Frieze, and Eli Upfal. Existence and construction of edge-disjoint paths on expander graphs. *SIAM Journal on Computing*, 23(5):976–989, 1994. [3.2.1](#), [A](#)
- [BI99] Eli Ben-Sasson and Russell Impagliazzo. Random cnf’s are hard for the polynomial calculus. In *40th Annual Symposium on Foundations of Computer Science, FOCS ’99*, pages 415–421, 1999. [1](#)
- [BKPS98] Paul Beame, Richard Karp, Toniann Pitassi, and Michael Saks. On the complexity of unsatisfiability of random k -CNF formulas. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, Dallas, TX, May 1998. [1](#)
- [BPU92] Stephen Bellantoni, Toniann Pitassi, and Alasdair Urquhart. Approximation and small depth Frege proofs. *SIAM J. on Comput.*, 21(6):1161–1179, 1992. [1](#)

- [BS02] Eli Ben-Sasson. Hard examples for bounded depth frege. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC-02)*, pages 563–572, New York, May 19–21 2002. ACM Press. (document), 1, 1
- [BSW99] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow – resolution made simple. In *STOC 1999*, pages 517–526, Atlanta, GA, May 1999. 1
- [Bus87] Samuel Buss. Polynomial size proofs of the propositional pigeonhole principle. *The Journal of Symbolic Logic*, 52(04):916–927, 1987. 1
- [Cap90] Renato Capocelli. A generalization of fibonacci trees. In *Applications of Fibonacci Numbers*, pages 37–56. Springer, 1990. 9
- [Con] Wikipedia Contributors. Generalizations of Fibonacci numbers. Posted at https://en.wikipedia.org/wiki/Generalizations_of_Fibonacci_numbers , accessed October 27, 2015. 9
- [CR79] Stephen A Cook and Robert A Reckhow. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic*, 44(01):36–50, 1979. 1, 2.1
- [CS88] V. Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, 1988. 1, 1
- [Die10] Reinhard Diestel. *Graph Theory*. Springer-Verlag, 2010. 3.2.1
- [Hak85] A. Haken. The intractability of resolution. *Theor. Comp. Sci.*, 39:297–305, 1985. 1
- [Hås86] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing*, pages 6–20, 1986. 1, 1
- [Juk12] Stasys Jukna. *Boolean Function Complexity*. Springer, 2012. 2.4
- [KPW95] Jan Krajíček, Pavel Pudlák, and Alan Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures & Algorithms*, 7(1):15–39, 1995. (document), 1, 1
- [KR96] Jon Kleinberg and Ronitt Rubinfeld. Short paths in expander graphs. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 86–95. IEEE, 1996. 1, 3.2.1, 3.2.1
- [PBI93] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational complexity*, 3(2):97–140, 1993. (document), 1, 1
- [Raz02] Alexander Razborov. Proof complexity of pigeonhole principles. In Werner Kuich, Grzegorz Rozenberg, and Arto Salomaa, editors, *Developments in Language Theory*, volume 2295 of *Lecture Notes in Computer Science*, pages 100–116. Springer Berlin Heidelberg, 2002. 1
- [RST15] Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. An average-case depth hierarchy theorem for boolean circuits. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS)*, 2015. To appear. 1

- [SBI04] Nathan Segerlind, Sam Buss, and Russell Impagliazzo. A switching lemma for small restrictions and lower bounds for k -DNF resolution. *SIAM Journal on Computing*, 33(5):1171–1200, 2004. [1](#)
- [SJ89] Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing markov chains. *Inf. Comput.*, 82(1):93–133, 1989. [A](#)
- [Tao11] T. Tao. 254B, Notes 1: Basic theory of expander graphs. posted at <https://terrytao.wordpress.com/2011/12/02/245b-notes-1-basic-theory-of-expander-graphs/>, 2011. [2.5](#)
- [Tre11] L. Trevisan. Lecture Notes on Expansion, Sparsest Cut, and Spectral Graph Theory. posted at <http://www.eecs.berkeley.edu/luca/books/expanders.pdf>, 2011. [2.5](#)
- [Tse68] G. S. Tseitin. On the complexity of derivation in the propositional calculus. In A. O. Slisenko, editor, *Studies in Constructive Mathematics and Mathematical Logic, Part II*. 1968. [1](#)
- [UF96] Alasdair Urquhart and Xudong Fu. Simplified lower bounds for propositional proofs. *Notre Dame Journal of Formal Logic*, 37(4):523–544, 1996. [1](#), [2.1](#)
- [Urq87] Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, 1987. [1](#)

A Proof of Lemma [2.7](#)

Recall the statement of Lemma [??](#): For every $S \subseteq E(\mathcal{G}_n)$, the graph $\mathcal{G}_n - S$ contains at most $C_1 \cdot |S| \cdot \log^2 n$ bridges, where C_1 is a constant depending only on the expansion parameter γ .

Proof. First some notation and terminology: for $v \in V(\mathcal{G}_n)$ define $\mathbf{W}(v) = (v, \mathbf{v}_1, \dots, \mathbf{v}_M)$ to be an M -step random walk on \mathcal{G}_n starting from v , where $M = K \log n$ for a suitable large constant K (to be specified later, but depending only on the expansion parameter γ). We say that two walks $W_1 = (a_0, \dots)$ and $W_2 = (b_0, \dots)$ *cross* if $a_i = b_j$ for some pair $(i, j) \neq (0, 0)$ (so if two walks start at the same vertex but do not intersect at any later step, they do not cross).

At the heart of Lemma [2.7](#) is the following statement, which says that random walks on an expander are not too likely to cross.

Claim A.1. *Fix any two vertices $u, v \in V(\mathcal{G}_n)$ (which may coincide). Let $\mathbf{W}_1(u) = (u, \mathbf{u}_1, \dots, \mathbf{u}_M)$, $\mathbf{W}_2(v) = (v, \mathbf{v}_1, \dots, \mathbf{v}_M)$ be two independent random walks as described above. Then*

$$\Pr[\mathbf{W}_1(u) \text{ and } \mathbf{W}_2(v) \text{ do not cross}] \geq \tau$$

for a constant $\tau > 0$ depending only on γ .

Proof. (The following argument is due to Noga Alon.) Since \mathcal{G}_n is a γ -expander for some constant γ , its girth is at least $c \log n$. Hence for any constant C we have that with probability at least 9^{-C} the length- C prefixes of $\mathbf{W}_1(u)$ and $\mathbf{W}_2(v)$ do not cross in the first C steps and have $\mathbf{u}_C, \mathbf{v}_C$ at

least distance C apart from each other in \mathcal{G}_n ; in what follows we condition on this taking place. Given this, we have

$$\begin{aligned}
\Pr[\mathbf{W}_1(u) \text{ and } \mathbf{W}_2(v) \text{ cross}] &\leq \sum_{i,j \leq M, i+j \geq C} \Pr[\mathbf{u}_i = \mathbf{v}_j] \\
&= \sum_{i,j \leq M, i+j \geq C} \sum_{v \in V(\mathcal{G}_n)} \Pr[\mathbf{u}_i = \mathbf{v}_j = v] \\
&= \sum_{i,j \leq M, i+j \geq C} \sum_{v \in V(\mathcal{G}_n)} \Pr[\mathbf{u}_i = v] \cdot \Pr[\mathbf{v}_j = v]
\end{aligned}$$

where the second equality is by independence of $\mathbf{W}_1(u)$ and $\mathbf{W}_2(v)$. By the eigenvalue characterization of expansion, though, we have that $\Pr[\mathbf{u}_i = v] \leq \frac{1}{n} + (1 - c')^i$ for a constant c' depending only on γ . Hence for any fixed i, j , we have that $\sum_{v \in V(\mathcal{G}_n)} \Pr[\mathbf{u}_i = v] \cdot \Pr[\mathbf{v}_j = v]$ is an inner product of two probability vectors (non-negative vectors whose entries sum to 1) where the maximum entry in the first vector is at most $1/n + (1 - c')^i$ and the maximum entry in the second is at most $1/j + (1 - c')^j$; hence $\sum_{v \in V(\mathcal{G}_n)} \Pr[\mathbf{u}_i = v] \cdot \Pr[\mathbf{v}_j = v]$ is at most $1/n + (1 - c')^{\max\{i,j\}}$. We thus have

$$\Pr[\mathbf{W}_1(u) \text{ and } \mathbf{W}_2(v) \text{ cross}] \leq \sum_{i,j \leq M, i+j \geq C} 1/n + (1 - c')^{\max\{i,j\}} \leq \frac{M^2}{n} + \sum_{i,j \leq M, i+j \geq C} (1 - c')^{\max\{i,j\}}.$$

Recalling that $M = O(\log n)$, we see that choosing C to be a suitably large absolute constant compared with $1/c'$, the preceding quantity is at most (say) $1/2$, and the claim is proved. \square

Armed with Claim A.1 we prove Lemma 2.7 as follows. Let $T > 0$ be an integer. Given two distinct nodes $u, v \in V(\mathcal{G}_n)$, we consider the following random experiment, denoted $\mathbf{EX}(u, v)$, which generates T pairs of random simple paths $\mathbf{P}_i(u, v), \mathbf{P}'_i(u, v)$ each of which joins u and v . (Note that the following is very similar in spirit to the generation of the walks that comprise ‘‘bundles’’ in Step 2(b) of the distribution $\mathcal{H}_{n'', n'}$ described in Section 3.2.1). The experiment is as follows: for $i = 1, \dots, T$,

1. Draw a uniform vertex \mathbf{x}_i from $V(\mathcal{G}_n)$ and let $\mathbf{W}_i(u, \mathbf{x}_i)$ denote a uniform random trajectory of length M starting at u and ending at \mathbf{x}_i (i.e. $\mathbf{W}_i(u, \mathbf{x}_i)$ is a draw from $\mathbf{W}_i(u)$ conditioned on ending at \mathbf{x}_i). Similarly let $\mathbf{W}_i(v, \mathbf{x}_i)$ denote an independent such random trajectory starting at v and ending at \mathbf{x}_i . Let $\overline{\mathbf{W}}_i(u, v)$ be the concatenation of $\mathbf{W}_i(u, \mathbf{x}_i)$ and the reversal of $\mathbf{W}_i(v, \mathbf{x}_i)$ and let $\mathbf{P}_i(u, v)$ be $\overline{\mathbf{W}}_i(u, v)$ with any cycles removed.
2. Repeat the previous step but using a fresh independent draw of \mathbf{x}'_i in place of \mathbf{x}_i , obtaining $\mathbf{W}'_i(u, \mathbf{x}'_i), \mathbf{W}'_i(v, \mathbf{x}'_i), \overline{\mathbf{W}}'_i(u, v)$, and $\mathbf{P}'_i(u, v)$.

We observe that for a given $i \in [T]$, if $\overline{\mathbf{W}}_i(u, v)$ and $\overline{\mathbf{W}}'_i(u, v)$ are vertex-disjoint from each other except at their endpoints, then $\mathbf{P}_i(u, v)$ and $\mathbf{P}'_i(u, v)$ together give a simple cycle containing both u and v . If

- (i) $\mathbf{W}_i(u, \mathbf{x}_i)$ and $\mathbf{W}'_i(u, \mathbf{x}'_i)$ do not cross,
- (ii) $\mathbf{W}_i(u, \mathbf{x}_i)$ and $\mathbf{W}'_i(v, \mathbf{x}'_i)$ do not cross,

(iii) $\mathbf{W}_i(v, \mathbf{x}_i)$ and $\mathbf{W}'_i(u, \mathbf{x}'_i)$ do not cross,

(ii) $\mathbf{W}_i(v, \mathbf{x}_i)$ and $\mathbf{W}'_i(v, \mathbf{x}'_i)$ do not cross,

then $\overline{\mathbf{W}}_i(u, v)$ and $\overline{\mathbf{W}}'_i(u, v)$ are vertex-disjoint from each other except at their endpoints. We now observe that for a suitable absolute constant choice of the constant K , the distribution of $\mathbf{W}_i(u, \mathbf{x}_i)$ (note that this includes the uniform random choice of \mathbf{x}_i) has total variation distance at most (say) n^{-100} from the distribution of $\mathbf{W}_i(u)$, a random walk of length M starting from u . (Here we are using the well known fact that random walks on non-bipartite expander graphs mix rapidly; see e.g. equation (12) of [BFU94] or [SJ89].) So up to a (negligible) $O(n^{-100})$ additive factor, we have that

$$\Pr[\mathbf{P}_i(u, v) \text{ and } \mathbf{P}'_i(u, v) \text{ form a simple cycle containing } u, v]$$

is at most

$$\Pr \left[\begin{aligned} &\mathbf{W}_i(u) \text{ and } \mathbf{W}'_i(u) \text{ do not cross, } \mathbf{W}_i(u) \text{ and } \mathbf{W}'_i(v) \text{ do not cross,} \\ &\mathbf{W}_i(v) \text{ and } \mathbf{W}'_i(u) \text{ do not cross, and, } \mathbf{W}_i(v) \text{ and } \mathbf{W}'_i(v) \text{ do not cross} \end{aligned} \right]. \quad (23)$$

By independence and Claim A.1 we have that (23) is at least τ^4 , so $\mathbf{P}_i(u, v)$ and $\mathbf{P}'_i(u, v)$ together give a simple cycle containing both u and v with probability at least $\tau^4/2$ (accounting for the additive n^{-100} factor). By independence across different choices of $i \in [T]$, we have that with probability at least $1 - (1 - \tau^4/2)^T$, at least one of $(\mathbf{P}_1(u, v), \mathbf{P}'_1(u, v)), \dots, (\mathbf{P}_T(u, v), \mathbf{P}'_T(u, v))$ is a simple cycle containing both u and v .

Now we consider a uniform random independent draw of \mathbf{u}, \mathbf{v} from $V(\mathcal{G}_n)$ conditioned on $\mathbf{u} \neq \mathbf{v}$. Since \mathbf{u} and \mathbf{v} are each uniform, it follows that for each i , each individual edge of $\mathbf{W}_i(\mathbf{u})$ is uniform random over $E(\mathcal{G}_n)$, as is each individual edge of $\mathbf{W}'_i(\mathbf{u})$, and likewise for each individual edge of $\mathbf{W}_i(\mathbf{v}), \mathbf{W}'_i(\mathbf{v})$. It follows from a union bound that the probability that any of these edges (across all $i \in T$) belongs to $S \subseteq E(\mathcal{G}_n)$ is at most $(C \cdot T \cdot |S| \cdot \log n)/(3n/2)$ for a constant C . By a suitable choice of $T = \Theta(\log n)$, we additionally get that $\Pr[\text{no } (\mathbf{P}_i(\mathbf{u}, \mathbf{v}), \mathbf{P}'_i(\mathbf{u}, \mathbf{v})) \text{ pair is a simple cycle}] \leq (C \cdot T \cdot |S| \cdot \log n)/(3n/2)$. Hence the probability (over the random choice of (\mathbf{u}, \mathbf{v})) that \mathbf{u} and \mathbf{v} do not both lie on a simple cycle in $\mathcal{G}_n - S$, is at most $(C \cdot T \cdot |S| \cdot \log n)/(3n/2)$. Consequently, there exists some $u \in V(\mathcal{G}_n)$ such that at most $(n-1)(C \cdot T \cdot |S| \cdot \log n)/(3n/2)$ many vertices $v \neq u, v \in V(\mathcal{G}_n)$ do *not* lie on a simple cycle that (i) contains u , and (ii) moreover misses all edges of S . Lemma 2.7 follows on observing that any bridge e in $\mathcal{G}_n - S$ must have such a vertex v as one of its endpoints. \square