A Separator Theorem for Hypergraphs and a CSP-SAT Algorithm

Michal Koucký* Vojtěch Rödl† Navid Talebanfard‡

Abstract

We show that for every $r \geq 2$ there exists $\epsilon_r > 0$ such that any r-uniform hypergraph on m edges with bounded vertex degree has a set of at most $(\frac{1}{2} - \epsilon_r)m$ edges the removal of which breaks the hypergraph into connected components with at most $m/2$ edges. We use this to give a satisfiability algorithm for n-variable (d,k)-CSPs in which every variable appears in at most r constraints in time $d^{(1-\epsilon_r)n}$ where ϵ_r depends only on r, provided that k is small enough as a function of m. We will also show that unsatisfiable $(2,k)$-CSPs with variable frequency r can be refuted in tree-like resolution in size $2^{(1-\epsilon_r)n}$. Furthermore for Tseitin formulas on graphs with degree at most k (which are $(2,k)$-CSPs) we give a deterministic algorithm finding such a refutation.

1 Introduction

The (d,k)-SAT problem which naturally generalizes k-SAT is the problem of deciding whether a system of constraints on variables from an alphabet of size d, where each constraint is on at most k variables can be satisfied. We will call such a system of constraints a (d,k)-CSP and we will assume that when given as input it is represented by the set of truth tables of its constraints. Therefore the satisfiability of a (d,k)-CSP Ψ can be checked by exhaustive search in time $|\Psi|^{O(1)}d^n$, where n is the number of variables. Therefore it makes sense to look for exponential time algorithms beating this trivial running time.

For $d = 2$, namely the usual k-SAT problem, of course there is a plethora of such algorithms (see e.g. [PPZ99, PPSZ05, DGH$^+$02]). When the CSP encodes a certain structured problem we can also find improved algorithms. The notable example here is graph d-coloring problem which is a special case of $(d,2)$-SAT that can be solved in time $O(2^n)$ [BHK09]. More generaly $(d,2)$-SAT also admits non-trivial algorithms [BE05].

For the general (d,k)-SAT we are interested in finding algorithms running in time $d^{(1-\epsilon)n}$ for some $\epsilon > 0$ which we call the *savings* of the algorithm, and we would like these savings to be as large as possible. Note that any k-SAT algorithm can be easily converted to a (d,k)-SAT algorithm. For each of the original variables, introduce $\log d$ boolean variables representing the original value in binary, and then express each constraint as a $k \log d$-CNF. The conjunction of these CNFs is satisfiable if and only if the original CSP is satisfiable. Assuming that we can solve k-SAT in time $2^{(1-\epsilon_k)n}$, this yields an algorithm running in time $d^{(1-\epsilon_{k \log d})n}$. That is any non-trivial savings for k-SAT yields non-trivial savings for (d,k)-SAT. However these savings deteriorate as d grows. This turns out to be case also for algorithms which are directly designed to solve (d,k)-SAT. Schöning's seminal algorithm runns in time $O((\frac{d(k-1)}{k})^n)$. Similary a generalization of PPSZ anlyzed by Hertli

*Charles University. Supported by GAČR grant 19-27871X. koucky@iuuk.mff.cuni.cz

†Emory University. Research supported by NSF grant DMS 1764385. vrodl@emory.edu

‡Corresponding author. The Czech Academy of Sciences, Institute of Mathematics. Supported by GAČR grant 19-27871X. talebanfard@math.cas.cz

et al. [HHM$^+$16] has the same shortcoming. The central question is then whether it is possible to obtain savings independent of d.

Let us be more precise. Define

$$\sigma_{d,k} := \sup\{\delta : (d,k)\text{-SAT can be solved in time } O(d^{(1-\delta)n})\}.$$

The argument above gives $\sigma_{d,k} \geq \sigma_{2,k\log d}$. Furthermore Traxler [Tra08] shows that for all d, $\sigma_{2,k} \geq \sigma_{d,k}$. Therefore it follows that under Strong Exponential Time Hypothesis, for all d, $\lim_{k\to\infty}\sigma_{d,k} = 0$. Our central question can be rephrased as follows.

Question 1. *Is it case that for every k, $\lim_{d\to\infty}\sigma_{d,k} > 0$?*

We do not succeed in answering this question. However we show that if each variable appears in a small number of constraints then it is possible to decide the satisfiability with savings independent of d. This can be considered as an extension of a result of Wahlström [Wah05] who gave such an algorithm for CNF-SAT when variables have bounded occurences. However our argument is entirely different.

Theorem 2 (Main result, informally stated). *There exists an algorithm which decides the satisfiability of an n variable (d,k)-CSP in which every variable appears in at most r constraints in time $d^{(1-\epsilon)n}$ where ϵ depends only on r, provided that k does not grow too fast (as a function of n).*

The algorithm follows a simple randomized branching strategy. At every step we find a small set of variables that once given a value, the CSP breaks into disjoint parts each with at most half of the original variables. For every assignment on these variables we recursively solve the problem on the resulting smaller instances. To prove this we associate a natural hypergraph to the CSP. Then we prove a more general result and show that in every r-uniform hypergraph of small vertex degree on m edges, there exists a small set of hyperedges the removal of which breaks the hypergraph into connected components with at most $m/2$ hyperedges.

Theorem 3 (Hypergraph separator theorem, informally stated). *Let $H = (V,E)$ be an r-uniform hypergraph on m hyperedges and maximum vertex degree k. Provided that k does not grow too fast as a function of m, there exists set of $(1-\epsilon)m/2$ hyperedges the removal of which breaks the hypergraph into connected components with at most $m/2$ hyperedges. Furthermore, ϵ depends only on r.*

Another application of Theorem 3 concerns tree-like resolution refutations of unsatisfiable $(2,k)$-CSPs. We show that we can refute such CSPs with savings depending only on the maximum variable frequency. One extensively studied such class of CSPs are Tseitin formulas on graphs of degree k. Our result implies that these formulas can be refuted with constant savings. It is known that every n-variable unsatisfiable k-CNF has a tree-like resolution refutation of size $2^{(1-\Omega(1/k))n}$ ([BT16]). However the best known lower bounds are only $2^{(1-1/k^{\beta})n}$ for some $\beta < 0$ ([PI00, BI13, BT16]). Closing this gap remains a challenge. Our result sheds light (though limited) on this question suggesting that it might even possible to improve the upper bound.

2 Hypergraph Separator Theorem

A *connected component* in a hypergraph $H = (V,E)$ is a maximal subset of vertices $S \subseteq V$ such that for every pair $u,v \in S$ there exists a sequence of edges $e_1,\ldots,e_t$ only on vertices in S with $u \in e_1$ and $v \in e_t$ and $e_i \cap e_{i+1} \neq \emptyset$ for every $i \in [t-1]$.

Definition 4. *Let $H = (V, E)$ be a hypergraph. A balanced separator for H is a set $R \subseteq E$ such that any connected component in $(V, E \setminus R)$ has at most $|E|/2$ edges.*

Note that any set $R \subseteq E$ of size $|E|/2$ is trivially a balanced separator. Therefore the question is whether it is possible to get balanced separators of size strictly less than $|E|/2$. We show that when the hypergraph has bounded vertex degree this is indeed possible. Let us first fix some notation. For any $S \subseteq V$ we denote by $E(S)$ the set of edges induced on S, i.e., all edges which are entirely on vertices in S. For $S, T \subseteq V$ we define $E(S, T) = \{e \in E : e \cap S \neq \emptyset, e \cap T \neq \emptyset\}$.

We will first show that when the maximum vertex degree of the hypergraph is small enough, then small balanced separators do exist. We then show that our bound is quantitatively tight.

2.1 Small Separators

We need the following concentration bound. This is a consequence of an inequality due to Alon, Kim and Spencer [AKS97] which was used by Dellamonica and Rödl [DR19].

Lemma 5. *Let $X_1, \dots, X_n$ be independent Bernoulli variables with $\Pr[X_i = 1] = p$ for all $i \in [n]$. Let $c > 0$ and assume that there is a function $f : \{0,1\}^n \to \mathbb{R}$ such that for all $(x_1, \dots, x_n) \in \{0,1\}^n$ and every $i \in [n]$,*

$$|f(x_1, \dots, x_n) - f(x_1, \dots, x_{i-1}, 1 - x_i, x_{i+1}, \dots, x_n)| \leq c.$$

Then for $\sigma^2 := nc^2 p(1-p)$ and any $0 < \alpha < 2\sigma/c$, we have

$$\Pr[|X - \mathbb{E}[X]| \geq \alpha\sigma] \leq 2e^{-\alpha^2/4},$$

where $X := f(X_1 \dots, X_n)$.

Theorem 3. *Let $r \geq 2$ be fixed and let $H = (V, E)$ be a hypergraph on m edges and maximum vertex degree k where each hyperedge has size at most r. Assume that $k\sqrt{n} \ll m$. Then H has a balanced separator of size at most $(\frac{1}{2} - \epsilon_r)m + o(m)$, where $\epsilon_r = (1 - 2^{-1/r})^r$. Furthermore such a balanced separator can be found by a randomized algorithm in expected polynomial time and by a deterministic algorithm in time $m^{O(1)} 2^{(1-2\epsilon_r^2)m + o(m)}$.*

Proof. We first observe that with a small modification we may assume without loss of generality that the hypergraph is r-uniform. To make the hypergraph r-uniform we partition the hyperedges of size less than r into blocks of size k. For each of these blocks we introduce r new vertices and we add sufficiently many of them to the hyperedges in the block to make them r-uniform. This guarantees that the maximum degree remains at most k and we have added at most mr/k new vertices. We then remove isolated vertices if any exists. We let n denote the number of vertices.

The idea is to find a set S of vertices such that $|E(S)| = m/2 \pm o(m)$ and $|E(S, \overline{S})| \leq (\frac{1}{2} - \epsilon_r)m \pm o(m)$. Observe that for such S, we also have $|E(\overline{S})| \leq m/2 + o(m)$. Furthermore note that $E(S)$ and $E(\overline{S})$ are separated by $R := E(S, \overline{S})$, i.e., every connected component in $(V, E \setminus R)$ is entirely contained in S or in $\overline{S}$. Then we arbitrarily select two sets $W_1 \subseteq E(S)$ and $W_2 \subseteq E(\overline{S})$ with $|W_1|, |W_2| \leq o(m)$ such that $|E(S) \setminus W_1| \leq m/2$ and $|E(\overline{S}) \setminus W_2| \leq m/2$, and we know that this is possible by the assumption on S. It follows that $R \cup W_1 \cup W_2$ is a balanced separator of size at most $(1 - \epsilon_r)m/2 + o(m)$.

We pick the set $S \subseteq V$ by including each vertex independently with probability $p := 2^{-1/r}$. Let $V = [n]$. For each vertex $i \in [n]$, let X_i be the random variable which takes value 1 if vertex

i is chosen, and it take value 0 otherwise. Thus we have $S = \{i \in [n] : X_i = 1\}$. We define two functions f_1 and f_2 as follows:

$$f_1(X_1, \ldots, X_n) := |E(\{i \in [n] : X_i = 1\})| = |E(S)|$$

and

$$f_2(X_1, \ldots, X_n) := |E(\{i \in [n] : X_i = 1\}, \{i \in [n] : X_i = 0\})| = |E(S, \overline{S})|.$$

Hence observe that

$$\mathbb{E}[f_1(X_1, \ldots, X_n)] = p^r m = \frac{m}{2}$$

and

$$\mathbb{E}[f_2(X_1, \ldots, X_n)] = (1 - p^r - (1-p)^r)m = (\frac{1}{2} - \epsilon_r)m.$$

We would like to apply Lemma 5 on f_1 and f_2. Note that since the maximum degree of H is at most k, for any $b \in [2]$ and $i \in [n]$,

$$|f_b(X_1, \ldots, X_n) - f_b(X_1, \ldots, X_{i-1}, 1 - X_i, X_{i+1}, \ldots, X_n)| \leq k.$$

Setting $\alpha = 4$, $c = k$ and $\sigma^2 = nk^2 p(1-p)$ we apply Lemma 5 on f_1 and f_2 and obtain

$$\Pr\left[|f_1(X_1, \ldots, X_n) - \frac{m}{2}| \geq 4k\sqrt{np(1-p)}\right] \leq 2e^{-4}$$

and

$$\Pr\left[f_2(X_1, \ldots, X_n) - (\frac{1}{2} - \epsilon_r)m \geq 4k\sqrt{np(1-p)}\right] \leq 2e^{-4}.$$

Since r is fixed and $k\sqrt{n} \ll m$ we have $4k\sqrt{np(1-p)} \leq 4k\sqrt{nr} = o(m)$. Therefore with probability at least $1 - 4e^{-4} > 0$, $|E(S)| = m/2 \pm o(m)$ and $|E(S, \overline{S})| \leq (\frac{1}{2} - \epsilon_r)m + o(m)$ and hence there exists a choice of S satisfying these properties. As explained earlier by adding at most $o(m)$ edges to $E(S, \overline{S})$ we obtain a balanced separator of size $(\frac{1}{2} - \epsilon_r)m + o(m)$.

It is clear the above argument also yields a randomized algorithm for finding such a balanced separator. The probability that S satisfies our desired properties is a constant and in polynomial time we can verify whether it indeed satisfies those properties. Thus in expected polynomial time we find our balanced separator.

The deterministic algorithm exhaustively checks all sets of at most $(\frac{1}{2} - \epsilon_r)m + o(m)$ edges to find a balanced separator. This has running time $m^{O(1)}\binom{m}{(\frac{1}{2}-\epsilon_r)m+o(m)} \leq m^{O(1)}2^{h(1/2-\epsilon_r)m+o(m)}$, where $h(\cdot)$ is the binary entropy function. Using $h(1/2 - x/2) \leq 1 - x^2/2$, we can bound the running time by $m^{O(1)}2^{(1-2\epsilon_r^2)m+o(m)}$.

$\square$

2.2 Optimality

We now show that Theorem 3 is tight.

Lemma 6. *For every fixed $\alpha < 1$ and $r \geq 2$ and $k, n \to \infty$ with $k \ll n$, there exists an r-uniform hypergraph $H = (V, E)$ with the following properties:*

1. $|E| = (1 \pm o(1))\frac{nk}{r}$
2. $\Delta(H) = O(k)$

3. *For every $S \subseteq V$ with $|S| \geq \alpha n$, $|E(S)| = (1 \pm o(1))(\frac{|S|}{n})^r |E|$.*

Proof. We first sample G from $G^{(r)}(n,q)$ where $q = \frac{nk/r}{\binom{n}{r}}$. The expected number of edges is $\frac{nk}{r}$. Thus by Chernoff's bound with probability at least $1 - 2^{-\Omega(n)}$, $|E| = (1 \pm o(1))\frac{nk}{r}$.

Next we bound the number of edges which are incident to vertices of degree at least $2ek$. For every vertex the probability that it has degree at least t is at most $\binom{\binom{n-1}{r-1}}{t} q^t$. Note that for $t \geq 2ek$ this probability is at most 2^{-t}. Therefore we can upper bound the expected number of edges incident with some vertex of degree at least $2ek$ by

$$n \sum_{t=2ek}^{\infty} t 2^{-t} = O(nk/2^k).$$

Therefore by Markov's inequality with constant probability the number of edges incident with some vertex of degree at least $2ek$ is at most $O(nk/2^k)$.

Let $S \subseteq V$ be any set of size αn. The expected number of edges in S is $\binom{\alpha n}{r} q$. Let $\delta = k^{-1/3} = o(1)$. By Chernoff's bound

$$\begin{aligned}
\Pr[|E(S)| \neq (1 \pm \delta)\binom{\alpha n}{r} q] &\leq 2\exp(-\delta^2 \binom{\alpha n}{r} q/3) \\
&= 2\exp(-\delta^2 \binom{\alpha n}{r} nk/3r\binom{n}{r}) \\
&\leq 2\exp(-(1 \pm o(1))\alpha^r n k^{1/3}/3r),
\end{aligned}$$

where in the last inequality we use $\frac{\binom{\alpha n}{r}}{\binom{n}{r}} = (1 \pm o(1))\alpha^r n$ since r is fixed. Since there are $\binom{n}{\alpha n}$ choices for S, the probability that there exists $S \subseteq V$ of size αn with $|E(S)| \neq (1 \pm \delta)\binom{\alpha n}{r} q$ is at most

$$\binom{n}{\alpha n} \times 2\exp(-(1 \pm o(1))\alpha^r n k^{1/3}/3r) = o(1).$$

It follows that there exists an r-uniform hypergraph G with $(1 \pm o(1))\frac{n}{r}k$ edges, at most γn of which are incident with some vertex of degree at least $2ek$, where $\gamma = O(k/2^k)$. Furthermore for every $S \subseteq V$ with $|S| = \alpha n$, $|E(S)| = (1 \pm \delta)\binom{\alpha n}{r} q$. We remove at most γn edges to make the maximum vertex degree at most $O(k)$. Call the resulting hypergraph $H = (V, E)$. Note that $|E(H)| = (1 \pm o(1))\frac{n}{r}k$. Let $\delta' = \delta + \gamma$. We have in H for every $S \subseteq V$ with $|S| = \alpha n$, $|E(S)| = (1 \pm \delta')\binom{\alpha n}{r} q$. A simple averaging argument further gives that for every $S \subseteq V$ with $|S| \geq \alpha n$, $E(S) = (1 \pm \delta')\binom{|S|}{r} q = (1 \pm o(1))(\frac{|S|}{n})^r |E|$. □

Theorem 7. *For every fixed $r \geq 2$ and $k, m \to \infty$ with $k \ll m$, there exists an r-uniform hypergraph with vertex degree $O(k)$ and m edges such that any balanced separator of H has size at least $(\frac{1}{2} - \epsilon_r)m(1 \pm o(1))$, where $\epsilon_r := (1 - 2^{-1/r})^r$.*

We will show that the hypergraph $H = (V, E)$ given by Lemma 6 with $\alpha := 1 - 2^{-1/r}$ satisfies this property. Note that $\epsilon_r = \alpha^r$ and $(1 - \alpha)^r = 1/2$. Let $m = |E|$. The following fact proves the result for the case when the balance separator is a bipartition. As it turns out and which we will show later, this is also the core of the argument for the general case.

Fact 1. *Let (A, B) be a bipartition of H with $\min\{|A|, |B|\} \geq \alpha n$. Then $|E(A, B)| \geq (1/2 - \epsilon_r)m(1 \pm o(1))$.*

Proof. Assume $|A| \le |B|$ and $|A| = \gamma n$. Since $\gamma \ge \alpha$, Lemma 6 guarantees that $|E(A)| = \gamma^r m(1 \pm o(1))$ and $|E(B)| = (1-\gamma)^r m(1 \pm o(1))$ and hence

$$\begin{aligned} |E(A,B)| &= (1-\gamma^r-(1-\gamma)^r)m(1\pm o(1)) \\ &\ge (1-\alpha^r-(1-\alpha)^r)m(1\pm o(1)) \\ &= (\frac{1}{2}-\epsilon_r)m(1\pm o(1)), \end{aligned}$$

where the inequality follows since the function $1-x^r-(1-x)^r$ is increasing for $x \in (0, \frac{1}{2})$. □

Proof of Theorem 7. Let $R \subseteq E$ be a balanced separator in H of minimum size. The removal of edges in R breaks H into two or more connected components each with at most $m/2$ edges. By minimality of R these connected components are induced subgraphs. We group these connected components in two parts A and B such that $||A|-|B||$ is minimized. Assume that $|A| \le |B|$. We have two cases. Either B is connected or it contains more than one connected component. Note that $R \supseteq E(A,B)$. Assume first that B is connected. We have $|A| \ge (\alpha + o(1))n$ since otherwise $|B| > (1-\alpha+\mu)n$ for some $\mu > 0$ and hence $|E(B)| \ge (1-\alpha+\mu+o(1))^r m(1\pm o(1)) > m/2$, contradicting that $|E(B)| \le m/2$. Fact 1 then implies that $|R| \ge |E(A,B)| \ge (\frac{1}{2}-\epsilon_r)m(1\pm o(1))$.

Now assume that B contains more than one connected component. Thus we can write $B = B_1 \cup B_2$, where B_1 and B_2 is an arbitrary bipartition of these components. Assume $|B_1| \ge |B_2|$. We show that $|A| \ge |B_1|$. Assume for a contradiction that this is not the case. Then $|A| < \min\{|A \cup B_2|, |B_1|\}$ and further we have $\max\{|A\cup B_2|, |B_1|\} \le |B_1 \cup B_2|$. This means that $A \cup B_2$ and B_1 give a more balanced bipartition, contradicting the minimality of $||A|-|B||$. Since $|A|+|B_1|+|B_2| = n$ and $|A| \ge |B_1| \ge |B_2|$, we have $|A| \ge n/3 \ge \alpha n$ (recall that $\alpha = 1-2^{-1/r}$ and $\alpha \le 1/3$ when $r \ge 2$). Once again we can apply Fact 1 to conclude that $|R| \ge |E(A,B)| \ge (\frac{1}{2}-\epsilon_r)m(1\pm o(1))$. □

3 A CSP-SAT Algorithm

A (d,k)-CSP Ψ is defined by a set of variables X taking values in an alphabet Σ of size d and a set $\mathcal{C}$ of constraints each on most k of these variables. We write $\Psi = (X, \mathcal{C})$ to specify the variables and the constraint set. We will assume that the CSP is represented by the set of truth tables of its constraints. Observe that a $(2,k)$-CSP can be represented as a k-CNF. An assignment to the variables *satisfies* Ψ if it satisfies every constraint. The variable frequency of Ψ is the largest number of contraints that any variable appears in. Given a partial assignment ρ which gives values to a set $D \subseteq X$, the restriction of Ψ is denoted by $\Psi|_\rho$ which is a CSP on $X \setminus D$ and each constraint is restricted by fixing the values of variables in D by ρ.

Let $\Psi = (X, \mathcal{C})$ be a CSP. We construct a hypergraph $H_\Psi = (V, E)$ as follows. We set $V = \mathcal{C}$, that is every constraint is represented by a vertex in H. For every variable $x \in X$ we create a hyperedge $e_x := \{C : x \in C\}$, that is the set of constraints containing x form a hyperedge.

Proposition 8. *Assume that H_Ψ consists of connected components $H_1, \ldots, H_t$. Then Ψ can be expressed as $\wedge_{i=1}^t \Psi_i$ where $H_{\Psi_i} = H_i$ for each $1 \le i \le t$.*

Proof. This is immediate once we observe that for any $C \in H_i$ and $D \in H_j$ for $i \ne j$, C and D do not have any variable in common. □

Proposition 9. *Let $\Psi = (X, \mathcal{C})$ be a CSP and $H_\Psi = (V, E)$ be the corresponding hypergraph. Let ρ be a partial assignment which gives values to a set $D \subseteq X$. Then $H_{\Psi|_\rho}$ is obtained by removing all e_x with $x \in D$ from H_Ψ. Furthermore, if Ψ is unsatisfiable so is $\Psi|_\rho$ and consequently if $\Psi|_\rho$ breaks into $\wedge_{i=1}^t \Psi_i$ as in Proposition 8, then at least one of Ψ_is is unsatisfiable.*

Proof. After restricting some of the variables, those variables disappear and some constraints get simplified. But no new constraint is introduced and hence the hypergraph is obtained by removing the corresponding hyperedges. If a CSP is unsatisfiable, obviously it is unsatisfiable under any partial assignment. If an unsatisfiable CSP is decomposed into disjoint CSPs, at least one of these CSPs is unsatisfiable, since otherwise we can take a satisfying assignment from each part and since they are on disjoint sets of variables together they form a satisfying assignment of the whole CSP. □

We are now ready to describe our algorithm CSP-SAT.

Algorithm 1 CSP-SAT(Ψ)

Construct $H_\Psi = (V, E)$.
Construct a small balanced separator R as in Theorem 3 (either probabilistically or deterministically).
for all $\rho \in \Sigma^R$ **do**
  Let $\Psi|_\rho = \wedge_{i=1}^t \Psi_i$ as in Proposition 9.
  for all $i \in [t]$ **do**
    Exhaustively check if Ψ_i is satisfiable
  end for
  if all Ψ_i's are satisfiable **then return** satisfiable
  end if
end for
  return unsatisfiable

Theorem 2. *Let $r \geq 2$ be a fixed integer, $n, d, k \geq 1$ be integers such that $k \ll \sqrt{n}$. Let $\epsilon_r = (1 - 2^{-1/r})^r$. Let Ψ be an n-variable (d, k)-CSP with variable frequency at most r.* CSP-SAT(Ψ) *correctly decides the satisfiability of Ψ. Moreover, if $d \geq 3$ then it runs deterministically in time $|\Psi|^{O(1)} d^{(1-\epsilon_r)n+o(n)}$, if $d = 2$ then it runs in expected time $|\Psi|^{O(1)} 2^{(1-\epsilon_r)n+o(n)}$ if we find R randomly, and in deterministic time $|\Psi|^{O(1)} 2^{(1-2\epsilon_r^2)n+o(n)}$ if we find R deterministically.*

Proof. The correctness of the algorithm follows immediately from Proposition 9. In polynomial time we can construct $H_\Psi = (V, E)$. Observe that H_Ψ has n edges each of size at most r and it has vertex degree at most k. By Theorem 3, we can find a balanced separator of size $(\frac{1}{2} - \epsilon_r)n + o(n)$ deterministically in time $2^{(1-2\epsilon_r^2)n+o(n)}$ or probabilistically in expected polynomial time in n. After having found the balanced separator R, there are at most $d^{(\frac{1}{2}-\epsilon_r)n+o(n)}$ runs of the for loop over ρ. For each restriction ρ we spend $|\Psi|^{O(1)}$ time to compute the decomposition of Ψ. Then for each of these parts we exhaustively check its satisfiability in time at most $d^{n/2}$. Since Ψ breaks into at most $m := |V|$ parts the total running time after finding R is at most $|\Psi|^{O(1)} d^{(\frac{1}{2}-\epsilon_r)n+o(n)+n/2} = |\Psi|^{O(1)} d^{(1-\epsilon_r)n+o(n)}$. For $d \geq 3$, $2^n < 3^{(1-\epsilon_r)n}$ so the total running time including finding the separator is bounded by $|\Psi|^{O(1)} d^{(1-\epsilon_r)n+o(n)}$. For $d = 2$, if we use the randomized procedure to find R, the total expected running time will be at most $|\Psi|^{O(1)} 2^{(1-\epsilon_r)n+o(n)}$, and if we run the deterministic procedure to find R, the total running time is at most $|\Psi|^{O(1)}(2^{(1-2\epsilon_r^2)n+o(n)} + 2^{(1-\epsilon_r)n+o(n)}) \leq |\Psi|^{O(1)} 2^{(1-2\epsilon_r^2)n+o(n)}$. □

Notice, for $x \in [0,1]$, $2^{-x} \le 1 - \frac{x}{2}$ so $2^{-\frac{1}{r}} \le 1 - \frac{1}{2r}$ so $(1 - 2^{-\frac{1}{r}}) \ge 1 - (1 - \frac{1}{2r}) = \frac{1}{2r}$. Hence, $\epsilon_r \ge \frac{1}{(2r)^r}$.

Remark. We can slightly modify the algorithm and instead of performing exhaustive search on the disjoint parts of the CSP we can make a recursive call to the algorithm. It is easy to verify that this improves the savings by a factor of two.

Corollary 10. *Let $r \ge 1$ be a fixed real, $n, d, k \ge 1$ be integers such that $k \ll \sqrt{n}$. Let $\epsilon_r = (1 - 2^{-1/r})^r$. Let Ψ be an n-variable (d,k)-CSP with average variable frequency at most r.* $\mathsf{CSP\text{-}SAT}(\Psi)$ *correctly decides the satisfiability of Ψ. Moreover, if $d \ge 3$ then it runs deterministically in time $|\Psi|^{O(1)} d^{(1-\epsilon_{\lceil 2r \rceil}/2)n + o(n)}$, if $d = 2$ then it runs in expected time $|\Psi|^{O(1)} 2^{(1-\epsilon_{\lceil 2r \rceil}/2)n+o(n)}$ if we find R randomly, and in deterministic time $|\Psi|^{O(1)} 2^{(1-\epsilon^2_{\lceil 2r \rceil})n+o(n)}$ if we find R deterministically.*

Proof. Consider the case $d \ge 3$. There are at most $n/2$ variables with frequency $\ge \lceil 2r \rceil$. For all possible settings of those variables run the $\mathsf{CSP\text{-}SAT}(\Psi)$ algorithm on Ψ restricted to that setting. In such a restricted formula all variables have frequency at most $\lceil 2r \rceil$. By the previous theorem the running time will be $d^{n/2} \cdot |\Psi|^{O(1)} d^{(1-\epsilon_{\lceil 2r \rceil})n/2+o(n)}$. The case for $d = 2$ is analoguous. □

4 Upper Bounds for Tree-like Resolution

In this section we use our separator theorem to give non-trivial refutations of unsatisfiable $(2,k)$-CSPs (recall that these CSPs can be represented by k-CNFs) with bounded variable frequency. This class of CSPs includes the extensively studied Tseitin formulas which essentially encode that in a simple graph the number of odd degree vertices is even. Here we consider a more general definition for hypergraphs due to Pudlák and Impagliazzo [PI00].

Definition 11. *Let $H = (V, E)$ be a hypergraph and let $\lambda : V \to \{0, 1\}$. The Tseitin tautology on H, $T(H, \lambda)$, has a variable x_e for every edge $e \in E$ and states that for every $v \in V$, $\oplus_{e \ni v} x_e \equiv \lambda(v)$. When $\oplus_{v \in V} \lambda(v) \equiv 1$ (in which case we call λ an odd charge labelling) and each edge has even cardinality, $T(H, \lambda)$ is unsatisfiable. When H has maximum degree k, $T(H, \lambda)$ is a $(2,k)$-CSP.*

Theorem 12. *Let Ψ be an unsatisfiable $(2,k)$-CSP with variable frequency at most r on n variables. If r is fixed and $k \ll \sqrt{n}$, then there exists a tree-like resolution refutation of Ψ of size $2^{(1-\epsilon_r)n+o(n)}$, where $\epsilon_r = (1 - 2^{-1/r})^r$.*

Proof. It is useful to think of a tree-like resolution proof as a Prover-Delayer game (due to Urquhart [Urq11]). The players maintain a partial assignment ρ which is initially empty. At every step Prover queries a variable x which is not given a value by ρ. Delayer then assigns a value $b \in \{0, 1\}$ to x. Then ρ is extended by $(x = b)$ and they continue. The game stops when ρ falsifies some constraint in $\mathcal{C}$. Note that this will eventually happen since Ψ is unsatisfiable. It is easy to see that if there is a strategy for Prover which can always force a contradiction in at most t steps, then there exists a tree-like resolution refutation of size 2^t for Ψ. We therefore give a strategy for Prover which forces contradiction in at most $(1 - \epsilon_r)n + o(n)$ steps.

We will make use of Theorem 3 applied to H_Ψ and the strategy is quite immediate. Prover queries all variables corresponding to the hyperedges in the balanced separator R given by Theorem 3 of size at most $(\frac{1}{2} - \epsilon_r)n + o(n)$. Delayer assigns a value to these variables. Call the resulting partial assignment ρ. By the separator property and Proposition 9 and Proposition 8 we can write $\Psi|_\rho = \wedge_{i=1}^t \Psi_i$ for some t, where Ψ_is are on disjoint sets of at most $n/2$ variables. Furthermore at least one of Ψ_is is unsatisfiable. Prover then queries all the variables in Ψ_i. It is clear that once Delayer assigns to these variables a contradiction is forced. The total number of queried variables is at most $(\frac{1}{2} - \epsilon_r)n + o(n) + n/2 = (1 - \epsilon_r)n + o(n)$. □

Corollary 13. *Let $H = (V, E)$ be an r-uniform hypergraph of maximum degree k where r is even and let $\lambda : V \to \{0, 1\}$ be an odd charge labeling. If r is fixed and $k \ll \sqrt{|E|}$ then $T(H, \lambda)$ can be refuted in tree-like resolution in size $2^{(1-\epsilon_r)|E|+o(|E|)}$, where $\epsilon_r = (1 - 2^{-1/r})^r$.*

Proof. Observe that each variable appears in r constraints. □

The case $r = 2$ corresponds to the the usual Tseitin tautologies on simple graphs. We give a finer analysis for this case which also involves a sharper derandomization.

Theorem 14. *There exists a deterministic algorithm which on input $T(G, \lambda)$ where graph $G = (V, E)$ satisfies $\Delta(G) = o(\sqrt{|E|}/\log |E|)$ and λ is any odd-charge labeling, produces a tree-like resolution refutation of $T(G, \lambda)$ in time $|E|^{O(1)}2^{(1-(1-\frac{1}{\sqrt{2}})^2)|E|+o(|E|)} \leq |E|^{O(1)}2^{0.914|E|}$.*

Proof. We will use a finer way of characterizing tree-like resolution size. A decision tree refuting an unsatisfiable CNF φ, is a decision tree on variables from φ and each leaf is labeled by a clause from φ with the condition every branch viewd as partial assignment falsifies the clause at the corresponding leaf. It is known and easy to show that there is a tree-like resolution refutation of φ of the same size as the decision tree (see e.g. [BGL13]). Thus it is sufficient to construct a small decision tree.

Let $n = |V|$ and $m = |E|$. We first compute the size of a small balanced separator of G. By Theorem 3, G has a balanced separator of size at most $(\frac{1}{2} - \epsilon_2)m$ where $\epsilon_2 = (1 - \frac{1}{\sqrt{2}})^2$. We now give a deterministic procedure for finding such a separator. The derandomization is based on the following observation. If G has minimum degree at least 3 we have $n \leq 2m/3$. We can then exhaustively check all subsets $S \subseteq V$ to find a small balanced separator in time $m^{O(1)}2^{2m/3}$. We need to address the assumption that the minimum degree is at least 3 and how we can afford this exponential running time. We construct a rooted tree T_G as follows. Each node ν_i on this tree is labeled by a pair (G_i, R_i) where G_i is a subgraph of G and R_i is a subset of edges of G_i. For any graph G, let G^* be the result of removing all edges incident with vertices of degree at most 2 in G. Initially we set $G_0 = G$. Then we set R_0 to be the small balanced separator that we deterministically find in G_0^*. For each connected component H in $G_0^* \setminus R_0$ we add a child ν_i to ν_0, set $G_i = H$ and continue recursively. It is clear that T_G has depth at most $\log n$ and that at each level it has at most m nodes (since there are at most n connected components). Therefore T_G has at most $n \log n$ nodes and the label of each node can be computed in time $m^{O(1)}2^{\frac{2}{3}m}$. Therefore T_G can be computed in this time.

We are now ready to construct the decision tree for $T(G, \lambda)$. Starting with $G_0 = G$, we first query all edges incident with vertices of degree at most 2. Note that this can be done by a polynomial size tree. Recall that assigning values to these edges amounts to removing those edges and hence we have obtained G_0^* at the end of a branch which is not falsifying any axiom of $T(G, \lambda)$. We can now query all edges in R_0 (which we have computed already). Each branch along these queries puts the odd charge in some connected component in $G_0 \setminus R_0$. We can then move to the corresponding node in T_G and continue recursively.

Since we query some edges to get rid of small degree vertices, this causes a blow up of polynomial factor in the size. The total running time is thus $m^{O(1)}2^{\frac{2}{3}m} + 2^{(1-\epsilon_2)m} = m^{O(1)}2^{(1-\epsilon_2)m}$.

□

5 Conclusion

We showed that we can remove a small number of edges from r-uniform hypergraphs with bounded vertex degree to break it into connected components each with at most half of the edges. This was used to solve the satsifiability of sparse CSPs. It would be interesting to find other applications of this separator theorem.

References

[AKS97] Noga Alon, Jeong-Han Kim, and Joel Spencer. Nearly perfect matchings in regular simple hypergraphs. *Israel J. Math.*, 100:171–187, 1997.

[BE05] Richard Beigel and David Eppstein. 3-coloring in time $\mathrm{o}(1.3289^{\mathrm{n}})$. *J. Algorithms*, 54(2):168–204, 2005.

[BGL13] Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A characterization of tree-like resolution size. *Inf. Process. Lett.*, 113(18):666–671, 2013.

[BHK09] Andreas Björklund, Thore Husfeldt, and Mikko Koivisto. Set partitioning via inclusion-exclusion. *SIAM J. Comput.*, 39(2):546–563, 2009.

[BI13] Christopher Beck and Russell Impagliazzo. Strong ETH holds for regular resolution. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 487–494, 2013.

[BT16] Ilario Bonacina and Navid Talebanfard. Improving resolution width lower bounds for k-CNFs with applications to the strong exponential time hypothesis. *Inf. Process. Lett.*, 116(2):120–124, 2016.

[DGH+02] Evgeny Dantsin, Andreas Goerdt, Edward A. Hirsch, Ravi Kannan, Jon M. Kleinberg, Christos H. Papadimitriou, Prabhakar Raghavan, and Uwe Schöning. A deterministic $(2 - 2/(k + 1))^n$ algorithm for k-SAT based on local search. *Theor. Comput. Sci.*, 289(1):69–83, 2002.

[DR19] Domingos Dellamonica and Vojtěch Rödl. Packing paths in steiner triple systems. *accepted*, 2019.

[HHM+16] Timon Hertli, Isabelle Hurbain, Sebastian Millius, Robin A. Moser, Dominik Scheder, and May Szedlák. The PPSZ algorithm for constraint satisfaction problems on more than two colors. In *Principles and Practice of Constraint Programming - 22nd International Conference, CP 2016, Toulouse, France, September 5-9, 2016, Proceedings*, pages 421–437, 2016.

[PI00] Pavel Pudlák and Russell Impagliazzo. A lower bound for DLL algorithms for k-SAT (preliminary version). In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA.*, pages 128–136, 2000.

[PPSZ05] Ramamohan Paturi, Pavel Pudlák, Michael E. Saks, and Francis Zane. An improved exponential-time algorithm for k-SAT. *J. ACM*, 52(3):337–364, 2005.

[PPZ99] Ramamohan Paturi, Pavel Pudlák, and Francis Zane. Satisfiability coding lemma. *Chicago J. Theor. Comput. Sci.*, 1999, 1999.

[Tra08] Patrick Traxler. The time complexity of constraint satisfaction. In *Parameterized and Exact Computation, Third International Workshop, IWPEC 2008, Victoria, Canada, May 14-16, 2008. Proceedings*, pages 190–201, 2008.

[Urq11] Alasdair Urquhart. The depth of resolution proofs. *Studia Logica*, 99(1-3):349–364, 2011.

[Wah05] Magnus Wahlström. An algorithm for the SAT problem for formulae of linear length. In *Algorithms - ESA 2005, 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, pages 107–118, 2005.