

INAPPROXIMABILITY OF COUNTING INDEPENDENT SETS IN LINEAR HYPERGRAPHS

GUOLIANG QIU, JIAHENG WANG

ABSTRACT. It is shown in this note that approximating the number of independent sets in a k -uniform linear hypergraph with maximum degree at most Δ is NP-hard if $\Delta \geq 5 \cdot 2^{k-1} + 1$. This confirms that for the relevant sampling and approximate counting problems, the regimes on the maximum degree where the state-of-the-art algorithms work are tight, up to some small factors. These algorithms include: the approximate sampler and randomised approximation scheme by Hermon, Sly and Zhang (RSA, 2019), the perfect sampler by Qiu, Wang and Zhang (ICALP, 2022), and the deterministic approximation scheme by Feng, Guo, Wang, Wang and Yin (FOCS, 2023).

1. INTRODUCTION

This note is concerned with the problem of counting independent sets in hypergraphs. We start with the basic definitions. A *hypergraph* $H = (W, \mathcal{E})$ is specified by a set of vertices W and a set of hyperedges \mathcal{E} , where each hyperedge $e \in \mathcal{E}$ is a subset of W . It is said to be *k -uniform*, if each hyperedge contains exactly k vertices. The degree of a vertex is the number of hyperedges in which it appears, and the *degree* Δ of the hypergraph is the maximum degree of its vertices. A set $I \subseteq W$ is a (weak) *independent set* if $I \cap e \neq e$ holds for all $e \in \mathcal{E}$. We remark that this problem is parameterised by k and Δ .

This problem was first studied by Bordewich, Dyer and Karpinski [BDK06, BDK08], where they showed that the straight-forward Markov chain over all independent sets is rapid mixing when $\Delta \leq k - 2$. This also yields a fully-polynomial randomised approximation scheme (FPRAS) for counting the number of independent sets by the standard counting-to-sampling reductions. For more information on such reductions, see for example [JVV86] or [ŠVV09]. The regime on the maximum degree was further improved by Hermon, Sly and Zhang [HSZ19] to $\Delta \leq c \cdot 2^{k/2}$ for some absolute constant $c > 0$. A later work by Qiu, Wang and Zhang [QWZ22] provided a perfect sampler (i.e., the output distribution is unbiased) which runs in expected polynomial time when $\Delta \leq c \cdot 2^{k/2}/k$ for some absolute constant $c > 0$. Very recently, Feng, Guo, Wang, Wang and Yin [FGW⁺23] further derandomised the Markov chain Monte Carlo approach and provided a fully-polynomial deterministic approximation scheme (FPTAS) when $\Delta \leq c \cdot 2^{k/2}/k^2$ for some absolute constant $c > 0$. On the other hand, Bezáková, Galanis, Goldberg, Guo and Štefankovič [BGG⁺19] proved that approximating the number of independent sets is intractable when $\Delta \geq 5 \cdot 2^{k/2}$ unless $\text{NP} = \text{RP}$. These results established a sharp *computational phase transition* for this problem.

The notion of *linear* hypergraphs (aka. *simple* hypergraphs) has also attracted some attention. A hypergraph is considered linear if the intersection of any two hyperedges contains at most one vertex. When the input hypergraph is restricted to be linear, the tractability regime of the aforementioned algorithms can be extended. Specifically, the same work mentioned above also established the following regimes: $\Delta \leq c \cdot 2^k/k^2$ for both the FPRAS and the perfect sampler, and $\Delta \leq 2^{(1-o(1))k}$ for the FPTAS. Under this setting, can we also establish a computational phase transition up to some low-order factors? As the main claim of this note, we answer the question affirmatively. Our result holds for a larger class of hypergraphs with overlap b , where the intersection of any two hyperedges contains at most b vertices.

(Guoliang Qiu) JOHN HOPCROFT CENTER FOR COMPUTER SCIENCE, SHANGHAI JIAO TONG UNIVERSITY, 800 DONGCHUAN ROAD, MINHANG DISTRICT, SHANGHAI, CHINA. E-mail: guoliang.qiu@sjtu.edu.cn

(Jiaheng Wang) SCHOOL OF INFORMATICS, UNIVERSITY OF EDINBURGH, INFORMATICS FORUM, EDINBURGH, EH8 9AB, UNITED KINGDOM. E-mail: jiaheng.wang@ed.ac.uk

Theorem 1. For any $k \geq 2$, $1 \leq b \leq k/2$ and $\Delta \geq 5 \cdot 2^{k-b} + 1$, it is NP-hard to approximate the number of independent sets in k -uniform hypergraphs with maximum degree at most Δ and overlap at most b .

The hardness for the linear hypergraph can be obtained by setting $b = 1$ in the theorem. Additionally, it subsumes the general case presented in [BGG⁺19] by setting $b := \lfloor k/2 \rfloor$.

In the general case [BGG⁺19], the hardness can be shown by reducing from the hard-core model (counting weighted independent sets) on *graphs*, where each vertex in the graph is replaced by $k/2$ copies in the resulting hypergraph, and each edge in the graph is associated with the hyperedge containing these copies. However, this approach naturally requires large overlaps in the resulting hypergraph. In our case, we ensure linearity (or small overlaps) by creating b copies for each vertex in the graph, and then filling up each hyperedge to k vertices. This boils down to a general anti-ferromagnetic 2-spin system, instead of the hard-core model. The main complication is to establish the so-called “non-uniqueness” property, which allows us to utilize a classical result by Sly and Sun [SS14] to derive inapproximability.

An open problem is to determine the exact computational phase transition for the linear case, as there remains a $\Theta(k^2)$ gap in-between. Notably, for independent sets on the infinite linear hypertree, the uniqueness of the Gibbs distribution holds when $\Delta \leq (1 + o(1)) 2^k/k$ [BGG⁺19]. It is still unclear which of these three thresholds, if any, would be the ground truth for the computational phase transition point.

The phenomenon of a more relaxed algorithmic regime (and thus a more restricted hardness regime) existing for linear hypergraphs is also observed in the context of hypergraph q -colouring problem. Currently, the best known algorithmic bound in the general hypergraph is $\Delta \lesssim q^{k/3}$ as established in [JPV21, HSW21, FGW⁺23].¹ In the case of linear hypergraphs, the bound can be improved to $q^{k/2-o(k)}$, as shown in [FGW22]. From the hardness side, the approximate counting problem is known to be NP-hard when $\Delta \geq 5 \cdot q^{k/2}$ [GGW22] in the general case when q is even, but $\Delta \geq 2kq^k \log q + 2q$ in the linear case [GGW22].

The problem of sampling hypergraph independent sets and q -colourings are both special cases of sampling solutions of constraint satisfaction problems (CSPs), and the computational results for these problems can be interpreted in the *local lemma* regime. In this framework, each hyperedge is treated as a constraint associated with a “bad event” that this constraint fails when all vertices are independently assigned values uniformly at random. Suppose that each bad event occurs with probability at most p and depends on at most D other bad events. By noticing that $D = k\Delta - 1$ in the hypergraph setting, the inapproximability conditions for counting both general hypergraph independent sets and colourings are unified as the form of $pD^2 \gtrsim 1$. In contrast, the well-known Lovász local lemma [EL75] provides a solution to *any* CSP when $pD \lesssim 1$, and such a solution can be efficiently constructed using the Moser-Tardos Algorithm [MT10]. This confirms the presence of a quadratic gap between the complexity of approximate counting/sampling and searching. For general CSPs, the conjectured tractable regime is $pD^2 \lesssim 1$, and there is a bunch of supporting evidence, see for example [WY23]. However, in light of the recent developments concerning linear hypergraphs, the tractable regime for CSPs with bounded overlap might extend further to $pD \lesssim 1$.

2. REDUCTION FROM 2-SPIN SYSTEMS

Our reduction is based on the hardness of approximating the partition function of the *2-spin system* on graphs. A 2-spin system on a graph $G = (V, E)$ is specified by an interaction matrix \mathbf{B} and a vector \mathbf{h} for the external field:

$$(1) \quad \mathbf{B} = \begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix}, \quad \mathbf{h} = \begin{bmatrix} \lambda \\ 1 \end{bmatrix},$$

where $\beta, \gamma, \lambda \geq 0$. The system is called *anti-ferromagnetic*, if $\beta\gamma < 1$. A configuration $\sigma : V \rightarrow \{0, 1\}$ assigns each vertex $v \in V$ with a spin either 0 or 1. The *weight* of a configuration σ is defined by

$$\text{wt}(\sigma) := \lambda^{n_0(\sigma)} \beta^{m_{00}(\sigma)} \gamma^{m_{11}(\sigma)}$$

¹The notations \lesssim and \gtrsim hide low-order factors.

where $n_0(\sigma)$ is the number of vertices assigned 0 under σ , and $m_{00}(\sigma)$ (resp. $m_{11}(\sigma)$) is the number of edges whose both endpoints are assigned 0 (resp. 1) under σ . The *partition function* is defined by

$$Z_{\beta,\gamma,\lambda}(G) := \sum_{\sigma} \text{wt}(\sigma).$$

The 2-spin system we are interested in is specified by the following choices of parameters:


$$(2) \quad \beta = 1, \quad \gamma = 1 - \frac{1}{2^{k-2b}}, \quad \lambda = 2^b - 1.$$

The subscription in $Z_{\beta,\gamma,\lambda}$ is thus omitted as the parameters are now fixed.

We now state the reduction. For any given Δ -regular graph $G = (V, E)$, construct the hypergraph H_G according to the following steps.

- (T1) Interpret the graph as a 2-uniform hypergraph.
- (T2) Replace each vertex with b vertices.
- (T3) For each hyperedge, insert another $k - 2b$ vertices independently.

Below is an example illustrating the reduction where $k = 7$, $b = 3$ and $\Delta = 3$.



It is immediate to verify that H_G is k -uniform, has overlap b and maximum degree Δ . Let $\mathcal{I}(H_G)$ be the set of independent sets of H_G .

Lemma 2. *For any Δ -regular graph $G = (V, E)$ and the constructed hypergraph H_G , it holds that $|\mathcal{I}(H_G)| = 2^{|E|(k-2b)} Z(G)$.*

Proof. We define the following partition over all the independent sets $\mathcal{I}(H_G) = \bigsqcup_{\sigma} \mathcal{S}(\sigma)$ in the hypergraph H_G , where σ ranges over all configurations of the 2-spin system. For any vertex $v \in V$, let B_v be the set of constructed vertices in H_G corresponding to v as in step (T2) of the construction. Given an independent set $I \in \mathcal{I}(H_G)$, the part $\mathcal{S}(\sigma)$ that I falls into is given by, for any $v \in V$,

- $\sigma(v) = 0$, if $|B_v \cap I| \leq b - 1$;
- $\sigma(v) = 1$, if $|B_v \cap I| = b$ (namely, $B_v \subseteq I$).

This is a partition because each $I \in \mathcal{I}(H_G)$ falls into exactly one part. Next, we demonstrate that $|\mathcal{S}(\sigma)| = 2^{|E|(k-2b)} \text{wt}(\sigma)$, which can be immediately used to establish the lemma. In subsequent discussions, a partial configuration of a subset $S \subseteq V(H_G)$ refers to a subset of vertices $S' \subseteq S$ that is included in the (weak) independent set of H_G . A partial configuration S' of S is considered feasible if there exists an independent set $I \in \mathcal{I}(H_G)$ such that $I \cap S = S'$.

- Consider the vertices constructed in (T2).
 - For each $v \in V$ such that $\sigma(v) = 0$, there are $2^b - 1$ feasible partial configurations of B_v .
 - For each $v \in V$ such that $\sigma(v) = 1$, there is just one feasible partial configuration of B_v .
- Consider the vertices constructed in (T3).
 - For each edge e such that both its endpoints take spin 1, the rest $k - 2b$ vertices of the corresponding hyperedge cannot be in an independent set together, so there are $2^{k-2b} - 1$ feasible partial configurations.
 - For any other edge, the corresponding $k - 2b$ vertices are free to be included in an independent set, so there are 2^{k-2b} feasible partial configurations.

In all, this gives

$$|\mathcal{S}(\sigma)| = \left(2^{k-2b} - 1\right)^{m_{11}(\sigma)} \left(2^{k-2b}\right)^{|E| - m_{11}(\sigma)} \left(2^b - 1\right)^{n_0(\sigma)} = 2^{|E|(k-2b)} \text{wt}(\sigma). \quad \square$$

Our goal then boils down to showing the inapproximability of the constructed 2-spin system. To establish this, we invoke the following celebrated result by Sly and Sun [SS14], which connects the so-called non-uniqueness property of any general anti-ferromagnetic 2-spin system with computational hardness. Denote by \mathbb{T}_Δ the infinite Δ -regular tree, and by $\hat{\mathbb{T}}_\Delta$ the infinite $(\Delta - 1)$ -ary tree.

Theorem 3 ([SS14]). *For any nondegenerate homogeneous anti-ferromagnetic 2-spin system specified by the interaction matrix \mathbf{B} and the external field \mathbf{h} on Δ -regular graphs that lies in the \mathbb{T}_Δ non-uniqueness region, the partition function is NP-hard to approximate, even within a factor of 2^{cn} for some constant $c > 0$ depending on Δ and the spin system.*

We remark that uniqueness/non-uniqueness regions for \mathbb{T}_Δ coincide with those for $\hat{\mathbb{T}}_\Delta$. However, $(\Delta - 1)$ -ary trees are more convenient to handle, so we move to $\hat{\mathbb{T}}_\Delta$ onwards. It is known that the $\hat{\mathbb{T}}_\Delta$ (non-)uniqueness region can be characterised by the solutions of the standard tree recursion formula on the ratio of the marginal probability of the root induced by the Gibbs measure. The following lemma, originally due to Martinelli, Sinclair and Weitz [MSW07, Section 6.2], characterises these solutions.

Lemma 4 ([GŠV16, Lemma 7]). *For $\Delta \geq 3$ and an anti-ferromagnetic 2-spin system specified by (1), consider the system of equations*

$$(3) \quad x = R(y), \quad y = R(x), \quad \text{where } R(z) := \lambda \left(\frac{\beta z + 1}{z + \gamma} \right)^{\Delta-1} \text{ and } x, y \geq 0.$$

Then,

- in the $\hat{\mathbb{T}}_\Delta$ uniqueness region, the system has a unique solution (Q^\times, Q^\times) ;
- in the $\hat{\mathbb{T}}_\Delta$ non-uniqueness region, the system has three solutions (Q^+, Q^-) , (Q^\times, Q^\times) , (Q^-, Q^+) where $Q^+ > Q^\times > Q^-$.

$R(z)$ in (3) is called the *one-step recursion*. For any pair of solutions (x, y) to (3), we have that $R(R(x)) = x$ and $R(R(y)) = y$. Therefore, both x and y are fixed points of the *two-step recursion* $R(R(z))$. Let $d := \Delta - 1$. Using the above lemma, it suffices to show that the two-step recursion has three fixed points $Q^+ > Q^\times > Q^-$ in order to establish non-uniqueness. Equivalently, in the regime of parameters specified in Theorem 1, we need to show that the following function, obtained by plugging the parameters specified in (2) into the two-step recursion, has three distinct zeros in $(0, +\infty)$:

$$(4) \quad f(z) := (2^b - 1) \left(1 + \frac{1}{2^{k-2b} (2^b - 1) \left(1 + \frac{1}{2^{k-2b} z + 2^{k-2b} - 1} \right)^d + 2^{k-2b} - 1} \right)^d - z.$$

Since the solution Q^\times is the unique fixed point of the one-step recursion, that is, $Q^\times = R(Q^\times)$, it is also helpful to consider the function

$$(5) \quad g(z) := (2^b - 1) \left(1 + \frac{1}{2^{k-2b} z + 2^{k-2b} - 1} \right)^d - z$$

obtained by plugging the parameters specified in (2) into the one-step recursion. The following lemma is sufficient to derive our main theorem.

Lemma 5. *Assume integers $k \geq 2$, $1 \leq b \leq k/2$ and $d = 5 \cdot 2^{k-b}$. Define $z^* := d/2^{k-2b} = 5 \cdot 2^b$. Then $f(z^*) > 0$ and $g(z^*) < 0$.*

Proof of Theorem 1. Note that $g(0) > 0$ and $\lim_{z \rightarrow +\infty} g(z) = -\infty$. By $g(z^*) < 0$, we know that the unique zero of g , which is Q^\times , is smaller than z^* . On the other hand, by $f(z^*) > 0$ and $\lim_{z \rightarrow +\infty} f(z) = -\infty$, there is a zero of f on $(z^*, +\infty)$, and it cannot be Q^\times . This establishes non-uniqueness due to Lemma 4. NP-hardness then follows after Theorem 3. \square

In the proof of Lemma 5, the following standard inequality is useful.

$$(6) \quad \exp\{x\} > \left(1 + \frac{x}{y} \right)^y > \exp\left\{ \frac{xy}{x+y} \right\} \quad \text{for all } x, y > 0.$$

Proof of Lemma 5. The $g(z^*)$ part is due to a straightforward estimation:

$$g(z^*) \leq (2^b - 1) \left(1 + \frac{1}{5 \cdot 2^{k-b}}\right)^{5 \cdot 2^{k-b}} - 5 \cdot 2^b < (2^b - 1)e - 5 \cdot 2^b < 0.$$

When $k = 2$, b can only take 1, and a direct calculation gives $f(z^*) > 16.0$ in this case. For $k \geq 3$, we have

$$\begin{aligned} f(z^*) &= (2^b - 1) \left(1 + \frac{1}{2^{k-2b}(2^b - 1) \left(1 + \frac{1}{d+2^{k-2b}-1}\right)^d + 2^{k-2b} - 1}\right)^d - 5 \cdot 2^b \\ &\geq (2^b - 1) \left(1 + \frac{1}{2^{k-2b}(2^b - 1) \left(1 + \frac{1}{d}\right)^d + 2^{k-2b} - 1}\right)^d - 5 \cdot 2^b \\ &> (2^b - 1) \left(1 + \frac{1}{2^{k-2b}(2^b - 1)e + 2^{k-2b} - 1}\right)^d - 5 \cdot 2^b \\ &> (2^b - 1) \left(1 + \frac{1}{2^{k-2b}(2^b - 1)e + 2^{k-2b}}\right)^d - 5 \cdot 2^b \end{aligned}$$

$$\text{(By (6) and } d = 5 \cdot 2^{k-b}\text{)} \quad > (2^b - 1) \exp\left\{\frac{5 \cdot 2^{b+k}}{2^k + 2^{2b} + 2^k(2^b - 1)e}\right\} - 5 \cdot 2^b.$$

Note that the fraction in $\exp\{\cdot\}$ is monotone increasing with respect to k . We can further analyse this by considering two cases.

- Fix $b \geq 2$. The whole term is minimised at $k = 2b$. Plugging this in, we further get

$$f(z^*) > (2^b - 1) \exp\left\{\frac{5 \cdot 2^b}{2 + (2^b - 1)e}\right\} - 5 \cdot 2^b =: h(b).$$

Now it suffices to show $h(b) > 0$ for $b \geq 2$. If $b = 2$, we have $h(2) > 1.5$. If $b \geq 3$, then

$$h(b) > (2^b - 1) \exp\left\{\frac{5 \cdot 2^b}{e \cdot 2^b}\right\} - 5 \cdot 2^b > 1.29 \cdot 2^b - 6.29 > 0.$$

- Fix $b = 1$. The whole term is minimised at $k = 3$ (we assume $k \geq 3$), and the value is at least 0.7.

Combining all these facts, the proof is complete. \square

ACKNOWLEDGEMENT

Jiaheng Wang has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 947778), and an Informatics Global PhD Scholarship at The University of Edinburgh.

REFERENCES

- [BDK06] Magnus Bordewich, Martin E. Dyer, and Marek Karpinski. Stopping times, metrics and approximate counting. In *ICALP*, volume 4051 of *Lecture Notes in Computer Science*, pages 108–119. Springer, 2006. [1](#)
- [BDK08] Magnus Bordewich, Martin E. Dyer, and Marek Karpinski. Path coupling using stopping times and counting independent sets and colorings in hypergraphs. *Random Struct. Algorithms*, 32(3):375–399, 2008. [1](#)
- [BGG⁺19] Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, Heng Guo, and Daniel Štefankovič. Approximation via correlation decay when strong spatial mixing fails. *SIAM J. Comput.*, 48(2):279–349, 2019. [1](#), [2](#)
- [EL75] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. II, pages 609–627. Colloq. Math. Soc. János Bolyai, Vol. 10. North-Holland, Amsterdam, 1975. [2](#)
- [FGW22] Weiming Feng, Heng Guo, and Jiaheng Wang. Improved bounds for randomly colouring simple hypergraphs. In *RANDOM*, volume 245 of *LIPICs*, pages 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. (full version in arXiv:2202.05554). [2](#)
- [FGW⁺23] Weiming Feng, Heng Guo, Chunyang Wang, Jiaheng Wang, and Yitong Yin. Towards derandomising Markov chain Monte Carlo. In *FOCS. IEEE*, 2023. To appear. [1](#), [2](#)

- [GGW22] Andreas Galanis, Heng Guo, and Jiaheng Wang. Inapproximability of counting hypergraph colourings. *ACM Trans. Comput. Theory*, 14(3–4), 2022. [2](#)
- [GŠV16] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combin. Probab. Comput.*, 25(4):500–559, 2016. [4](#)
- [HSW21] Kun He, Xiaoming Sun, and Kewen Wu. Perfect sampling for (atomic) Lovász local lemma. *arXiv*, abs/2107.03932, 2021. [2](#)
- [HSZ19] Jonathan Hermon, Allan Sly, and Yumeng Zhang. Rapid mixing of hypergraph independent sets. *Random Struct. Algorithms*, 54(4):730–767, 2019. [1](#)
- [JPV21] Vishesh Jain, Huy Tuan Pham, and Thuy Duong Vuong. On the sampling Lovász local lemma for atomic constraint satisfaction problems. *arXiv*, abs/2102.08342, 2021. [2](#)
- [JVV86] Mark Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theor. Comput. Sci.*, 43:169–188, 1986. [1](#)
- [MSW07] Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Struct. Algorithms*, 31(2):134–172, 2007. [4](#)
- [MT10] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2):11, 2010. [2](#)
- [QWZ22] Guoliang Qiu, Yanheng Wang, and Chihao Zhang. A perfect sampler for hypergraph independent sets. In *ICALP*, volume 229 of *LIPICs*, pages 103:1–103:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. [1](#)
- [SS14] Allan Sly and Nike Sun. Counting in two-spin models on d -regular graphs. *Ann. Probab.*, 42(6):2383–2416, 2014. [2](#), [4](#)
- [ŠVV09] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. Adaptive simulated annealing: A near-optimal connection between sampling and counting. *J. ACM*, 56(3):18, 2009. [1](#)
- [WY23] Chunyang Wang and Yitong Yin. Correlation decay up to the sampling threshold in the local lemma regime. *arXiv*, abs/2307.14872, 2023. [2](#)