SIAM J. COMPUT.
Vol. 31, No. 5, pp. 1527–1541

ON COUNTING INDEPENDENT SETS IN SPARSE GRAPHS*

MARTIN DYER†, ALAN FRIEZE‡, AND MARK JERRUM§

Abstract. We prove two results concerning approximate counting of independent sets in graphs with constant maximum degree Δ. The first implies that the Markov chain Monte Carlo technique is likely to fail if $\Delta \geq 6$. The second shows that no fully polynomial randomized approximation scheme can exist for $\Delta \geq 25$, unless RP = NP.

Key words. Primary, 68Q17; Secondary, 05C69, 60J10, 68E10, 68Q25, 68W40

AMS subject classifications. approximation algorithms, computational complexity, independent sets, mixing times of Markov chains, randomized algorithms

PII. S0097539701383844

1. Introduction.

Counting independent sets in graphs is one of several combinatorial counting problems which have received recent attention. The problem is known to be #P-complete, even for low-degree graphs [5]. On the other hand, it has been shown that, for graphs of maximum degree $\Delta = 4$, randomized approximate counting is possible [9, 5]. This success has been achieved using the *Markov chain Monte Carlo* method [8] to construct a *fully polynomial randomized approximation scheme (fpras)*. This has led to a natural question of how far this success might extend.

Here we consider in more detail this question of counting independent sets in graphs with constant maximum degree. We prove two results. The first, in section 2, shows that the Monte Carlo Markov chain method is likely to fail for graphs with $\Delta = 6$. This leaves open only the case $\Delta = 5$.

Our second result gives an explicit value of Δ above which approximate counting, using any kind of polynomial-time algorithm, is impossible unless RP = NP. The bound we obtain is $\Delta = 25$, though we suspect that the true value is in single figures, probably 6.

We note that Berman and Karpinski [2] have recently given new explicit bounds for the approximation ratio for the maximum independent set and other problems in low-degree graphs. These directly imply an inapproximability result for counting. (See Luby and Vigoda [9], specifically the proof of their Theorem 4.) However, the bound on Δ obtained this way is larger than ours by at least two orders of magnitude.

The questions we address in this article could also be studied in a more general setting in which vertices included in an independent set have weights or "fugacities" other than 1. In this setting, the weight of an independent set of size k is deemed to be λ^k for some constant k, and the goal is to compute the sum of the weights of all independent sets. One could then ask, for each Δ, at what exact λ an fpras ceases to

*Received by the editors January 20, 2001; accepted for publication (in revised form) March 15, 2002; published electronically August 5, 2002. A preliminary version of this article appeared in *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science*, New York, 1999, pp. 210–217.

http://www.siam.org/journals/sicomp/31-5/38384.html

†School of Computer Studies, University of Leeds, LS2 9JT, United Kingdom (dyer@scs.leeds.ac.uk). This author's research was supported by EC ESPRIT Working Group RAND2.

‡Department of Mathematics, Carnegie Mellon University, Pittsburgh, PA 15213 (alan@random.math.cmu.edu). This author's research was supported by NSF grant CCR 9520974.

§Department of Computer Science, University of Edinburgh, EH9 3JZ, United Kingdom (mrj@dcs.ed.ac.uk).

exist (assuming such a λ exists). This question is a more precise version of the one we ask: for $\lambda = 1$, what is the largest Δ for which an fpras exists?

A reasonable guess is that the critical λ just identified is greater than 1 when $\Delta \leq 5$, and less than 1 when $\Delta \geq 6$. One might even rashly conjecture (though we shall not do so) that this critical λ is the same as that marking the boundary between unique and multiple Gibbs measures in the independent set (hard core gas) model in the regular infinite tree of degree Δ (the so-called Bethe lattice). Brightwell and Winkler have computed the fugacity λ at which multiple Gibbs measures appear in the Bethe lattice [3]. The only observation we offer here is that our results are consistent with the critical λ's being the same in both situations.

2. Markov chain Monte Carlo. For a graph G, let $\mathcal{I}(G)$ denote the collection of independent sets of G. Let $\mathcal{M}(G)$ be any Markov chain, asymptotically uniform on $\mathcal{I}(G)$, with transition matrix P_G. In this section, G will be a bipartite graph with a vertex bipartition into classes of equal size n. Let $b(n) \leq n$ be any function of n, and suppose we have $P_G(\sigma_1, \sigma_2) = 0$ whenever $|\sigma_1 \oplus \sigma_2| > b(n)$, where $\oplus$ denotes symmetric difference. We will say that $\mathcal{M}(G)$ is $b(n)$-*cautious*. Thus a $b(n)$-cautious chain is not permitted to change the status of more than $b(n)$ vertices in G at any step. Ideally, for ease of implementation, we would wish to have $b(n)$ a constant (as in [9, 5]). However, we will show that no $b(n)$-cautious chain on $\mathcal{I}(G)$ can mix rapidly unless $b(n) = \Omega(n)$. Thus any chain which does mix rapidly on $\mathcal{M}(G)$ must change the status of a sizable proportion of the vertices at each step.

Before stating our result, we need to formalize what we mean by mixing, rapid or otherwise. Let $\mathcal{M}$ be an ergodic Markov chain with state space Ω, distribution p_t at time t, and asymptotic distribution $p_\infty = \pi$. Let $x_0 \in \Omega$ be the initial state of $\mathcal{M}$, so that p_0 assigns unit mass to state x_0. Define the *mixing time* $\tau(x_0)$ *of* $\mathcal{M}$*, with initial state* $x_0 \in \Omega$, as the first t for which $\mathrm{d_{TV}}(p_t, \pi) \stackrel{\text{def}}{=} \frac{1}{2}||p_t - \pi||_1 \leq e^{-1}$; then define the *mixing time* τ as the maximum of $\tau(x_0)$ over choices of initial state x_0. We are able to show the following.

THEOREM 2.1. *Suppose* $\Delta \geq 6$ *and* $b(n) \leq 0.35n$*. Then there exists a constant* $\gamma > 0$ *and a bipartite graph* G_0*, regular of degree* Δ*, on* $n+n$ *vertices (more precisely a sequence of such graphs parameterized by* n*) with the following property: any* $b(n)$*-cautious Markov chain on* $\mathcal{I}(G_0)$ *has mixing time* $\tau = \Omega(e^{\gamma n})$*.*

Since, of course, there does exist a $2n$-cautious chain which mixes rapidly, our result cannot be strengthened much further. Although we do not identify a specific initial state x_0 satisfying $\tau(x_0) = \Omega(e^{\gamma n})$, our proof does provide a definite (and natural) initial distribution p_0 from which $\tau = \Omega(e^{\gamma n})$ steps are required to achieve $\mathrm{d_{TV}}(p_\tau, \pi) \leq e^{-1}$. The remainder of this section is devoted to the proof of Theorem 2.1.

The counterexample graph G_0 is just a random regular graph of degree Δ. Specifically, let $K_{n,n}$ denote the complete bipartite graph with vertex bipartition V_1, V_2, where $|V_1| = |V_2| = n$, and let G be the union of Δ perfect matchings selected independently and uniformly at random in $K_{n,n}$. (Since the perfect matchings are independent, they may well share some edges.) Denote by $\mathcal{G}(n, n, \Delta)$ the probability space of bipartite graphs G so defined. Where no confusion can arise, we simply write $\mathcal{G}$ for this class below. Note that $\mathcal{G}$ is a class of graphs with degree bound Δ. It is well known (see [1]) that, provided Δ is taken as constant, Δ-regular graphs occur in $\mathcal{G}(n, n, \Delta)$ with probability bounded away from 0. Since we prove that almost every graph $G \in \mathcal{G}(n, n, \Delta)$, for $\Delta \geq 6$, has the property we seek, it will follow that almost every Δ-regular graph (in the induced probability space) has the property too.

Let $0 < \alpha, \beta < 1$ be chosen values. For $G \in \mathcal{G}$, we consider the collection $\mathcal{I}_G(\alpha, \beta)$ of $\sigma \in \mathcal{I}(G)$ such that $|\sigma \cap V_1| = \alpha n$ and $|\sigma \cap V_2| = \beta n$. We will call $\sigma \in \mathcal{I}_G(\alpha, \beta)$ an (α, β)-*set*. Using linearity of expectation, we may easily compute the expected number $\mathcal{E}(\alpha, \beta) = \mathbf{E}(|\mathcal{I}_G(\alpha, \beta)|)$ of (α, β)-sets in G: it is just the number of ways of choosing an αn-subset from V_1 and a βn-subset from V_2, multiplied by the probability that all Δ perfect matchings avoid connecting the αn-subset to the βn-subset. Thus, using Stirling's formula,

$$\begin{aligned}\mathcal{E}(\alpha, \beta) &= \binom{n}{\alpha n}\binom{n}{\beta n}\left[\binom{(1-\beta)n}{\alpha n}\Big/\binom{n}{\alpha n}\right]^{\Delta} \\ &= \left(\frac{(1-\beta)^{(\Delta-1)(1-\beta)}(1-\alpha)^{(\Delta-1)(1-\alpha)}}{\alpha^{\alpha}\beta^{\beta}(1-\alpha-\beta)^{\Delta(1-\alpha-\beta)}}\right)^{n(1+o(1))} \\ &= e^{\varphi(\alpha,\beta)n(1+o(1))}, \qquad (1)\end{aligned}$$

where

$$\begin{aligned}\varphi(\alpha, \beta) = \varphi_{\Delta}(\alpha, \beta) &= -\alpha \ln \alpha - \beta \ln \beta - \Delta(1-\alpha-\beta)\ln(1-\alpha-\beta) \\ &\quad + (\Delta-1)\big((1-\alpha)\ln(1-\alpha) + (1-\beta)\ln(1-\beta)\big). \qquad (2)\end{aligned}$$

Mostly, Δ will be treated as a constant, and we shall suppress the subscript of φ except when we want to emphasize the dependence on Δ.

We shall treat φ as a function of real arguments α and β, even though a combinatorial interpretation is possible only when αn and βn are integers. Then φ is defined on the triangle

$$\mathcal{T} = \{(\alpha, \beta) : \alpha, \beta \geq 0 \text{ and } \alpha + \beta \leq 1\}$$

and is clearly symmetrical in α, β. (The function φ is defined by (2) on the interior of $\mathcal{T}$ and can be extended to the boundary by taking limits.) Moreover, the following facts are established in the appendix about the stationary points of φ on $\mathcal{T}$.

CLAIM 2.2.

(i) *The function φ has no local minima in the interior of $\mathcal{T}$, and no local maxima on the boundary of $\mathcal{T}$.*

(ii) *All local maxima of φ satisfy $\alpha + \beta + \Delta(\Delta-2)\alpha\beta \leq 1$.*

(iii) *If $\Delta \leq 5$, φ has only a single local maximum, which is on the line $\alpha = \beta$.*

(iv) *If $\Delta \geq 6$, φ has exactly two local maxima, symmetrical in α, β, and a single saddle-point on $\alpha = \beta$. The maximum with $\alpha < \beta$ occurs at $(\alpha, \beta) \approx (0.03546955, 0.40831988)$ when $\Delta = 6$ and at $(\alpha, \beta) \approx (0.01231507, 0.45973533)$ when $\Delta = 7$.*

Suppose, for the sake of discussion, we had the additional information that the number $|\mathcal{I}_G(\alpha, \beta)|$ of (α, β)-sets is reasonably well concentrated about its expectation $\mathcal{E}(\alpha, \beta)$. Then it would follow from (iii) and (iv) that a "typical" independent set in a random graph $G \in \mathcal{G}(n, n, \Delta)$ undergoes a dramatic change in passing from $\Delta = 5$ to $\Delta = 6$. For $\Delta \leq 5$, a typical independent set σ would be balanced, i.e., the sets $|\sigma \cap V_1|$ and $|\sigma \cap V_2|$ would be of nearly equal size, whereas for $\Delta \geq 6$ it would be unbalanced.

Unfortunately, we have not been able to prove a concentration result, and it is unclear whether such a result should be expected. Therefore, in examining the first (apparently) unbalanced case, $\Delta = 6$, we must make a slight detour. First, observe

that a knowledge of φ does at least provide an *upper* bound on $|\mathcal{I}_G(\alpha,\beta)|$ via Markov's inequality. In this way we can bound from above the number of (α,β)-sets that lie in the strip $|\alpha-\beta|\le\eta$ for some $\eta>0$. Then, we use a quite crude lower bound to show that the number of (α^*,β^*)-sets—for some chosen α^*,β^* with $\beta^*-\alpha^*>\eta$—is much greater that this.

We shall first deal with the boundary case $\Delta=6$. Once this has been done, it will be easy to dispense with the remaining cases, i.e., $\Delta\ge 7$, which are less finely balanced. So suppose for the time being that $\Delta=6$. Consider the function φ restricted to the region $\mathcal{D}=\mathcal{T}\cap\{(\alpha,\beta):|\alpha-\beta|\le\eta\}$, where $\eta=0.18$. Since the two local maxima of φ on $\mathcal{T}$ lie outside $\mathcal{D}$ (see Claim 2.2(iv)), it must be the case that the maxima of φ on $\mathcal{D}$ all lie on one or the other (and hence, by symmetry, both) of the lines $|\alpha-\beta|=\eta$. Numerically, the (unique) maximum with $\beta-\alpha=\eta$, achieved at $(\alpha,\beta)\approx(0.10021227,0.28021227)$, is a little less than $c=0.70824602$. (The uniqueness of the maximum may be verified by calculus; then the location of the maximum may be found to arbitrary precision by repeated evaluation of the derivative of $\varphi(\alpha,\alpha+0.18)$ with respect to α. Only simple function evaluations are required.)

Now define

$$\theta(\alpha)=-\alpha\ln\alpha-(1-\alpha)\ln(1-\alpha)+(\ln 2)(1-\Delta\alpha)$$

for $\Delta\alpha<1$. Then, for *any* graph $G\in\mathcal{G}$, the total number of independent sets σ with $|\sigma\cap V_1|=\alpha n$ is (crudely) at least

$$|\mathcal{I}_G(\alpha,*)|\ge e^{\theta(\alpha)n(1-o(1))}.$$

(Choose αn vertices from V_1; then choose any subset of vertices from the at least $(1-\Delta\alpha)n$ unblocked vertices in V_2.) Set $\alpha^*=0.015$. Then, by numerical computation, $\theta(\alpha^*)$ is a little greater than $0.70864644>c$. Thus, with high probability, the number of (α,β)-sets in $G\in\mathcal{G}$ lying in either connected component of $\mathcal{T}\setminus\mathcal{D}$ is greater than the number lying within $\mathcal{D}$ by an exponential factor, specifically $e^{\gamma n}$, where $\gamma=0.0004$. The graph G_0 of Theorem 2.1 is any graph $G_0\in\mathcal{G}$ that exhibits the exponential gap just described. (A randomly chosen graph will do with high probability.) The remainder of our argument concerns G_0.

Now consider a $0.35n$-cautious chain $\mathcal{M}(G_0)=\mathcal{M}_0$ on $\mathcal{I}(G_0)$. Let A comprise all (α,β)-sets with $\alpha\ge\beta$, i.e.,

$$A=\big\{\sigma\in\mathcal{I}(G_0):|\sigma\cap V_1|\ge|\sigma\cap V_2|\big\},$$

and assume without loss of generality that A is no larger than its complement $\overline{A}=\mathcal{I}\setminus A$. Denote by M the set of (α,β)-sets with $(\alpha,\beta)\in\mathcal{D}$. Since $\mathcal{M}_0$ is $0.35n$-cautious, it cannot make a transition from A to $\overline{A}$ except by using a state in M. Now, we have already seen that

$$|A|\ge e^{\gamma n}\,|M|. \tag{3}$$

Intuitively, since M is very small in relation to A, the mixing time of $\mathcal{M}_0$ must be very large. This intuition is captured in the following claim, which is implicit in a line of argument used by Jerrum [7].

Claim 2.3. *Let $\mathcal{M}$ be a Markov chain with state space Ω, transition matrix P, and stationary distribution π. Let $A\subset\Omega$ be a set of states such that $\pi(A)\le\frac{1}{2}$, and*

$M \subset \Omega$ be a set of states that form a "barrier" in the sense that $P_{ij} = 0$ whenever $i \in A \setminus M$ and $j \in \overline{A} \setminus M$. Then the mixing time of $\mathcal{M}$ is at least $\pi(A)/8\pi(M)$.

For completeness, a proof using "conductance" is provided in the appendix. Theorem 2.1, in the boundary case $\Delta = 6$, follows from Claim 2.3 and inequality (3) because the sets A and M that we defined earlier satisfy the conditions of the claim. Note that the proof of Claim 2.3 actually provides an explicit initial distribution p_0 from which the mixing time is large, namely, the uniform distribution on A.

Finally, suppose $\Delta \geq 7$. We shall see presently that

$$\varphi_\Delta(\alpha, \beta) < 0.6763 < \ln 2 \quad \text{for all } \Delta \geq 7 \text{ and } (\alpha, \beta) \in \mathcal{D}. \tag{4}$$

On the other hand, there are at least 2^n (α, β)-sets in either connected component of $\mathcal{T} \setminus \mathcal{D}$: this comes simply from considering independent sets with $\alpha = 0$ or $\beta = 0$. Once again, with high probability, the number of (α, β)-sets in $G \in \mathcal{G}$ lying in either connected component of $\mathcal{T} \setminus \mathcal{D}$ is greater than the number lying within $\mathcal{D}$ by an exponential factor, specifically $e^{\gamma n}$, where $\gamma = 0.015$. Theorem 2.1, in the general case $\Delta \geq 7$, follows as before.

It remains only to verify (4). By calculus, $\varphi_\Delta(\alpha, \beta)$ as a function of Δ is monotonically decreasing over the whole region $\mathcal{T}$; thus we need check only the case $\Delta = 7$. (The partial derivative $\partial \varphi_\Delta(\alpha, \beta)/\partial \Delta$ is a function of α and β only; it is zero on $\alpha = 0$ and monotonically decreasing as a function of β.) We now argue, as before, that the maxima of φ on $\mathcal{D}$ all lie on the lines $|\alpha - \beta| = 0.18$. (Here we again use Claim 2.2(iv).) Once again, by calculus, φ has a unique maximum on each of these lines, and direct calculation yields (4).

3. Hardness of approximate counting. The result of the previous section implies that the usual approach to approximating the number of independent sets in a low-degree graph must fail when $\Delta \geq 6$, at least in the worst case. Here we show that, if the degree bound is somewhat larger, then *any* approach to approximating the number of independent sets is doomed to failure, under a reasonable complexity assumption. Precisely, the remainder of this section is devoted to proving the following theorem.

THEOREM 3.1. *Unless* RP = NP*, there is no polynomial-time algorithm that estimates the* logarithm *of the number of independent sets in a Δ-regular graph ($\Delta \geq 25$) within relative error at most $\varepsilon = 10^{-4}$.*

We give a randomized reduction from the following problem E2LIN2, analyzed by Håstad. The input is a system $\mathcal{A}$ of m equations over $\mathbb{Z}_2$ in n variables $x_1, x_2, \ldots, x_n$, such that each equation has exactly two variables. (Thus each equation is of the form $x_i + x_j = 0$ or $x_i + x_j = 1$.) The objective is to find a maximum cardinality consistent subset of equations in $\mathcal{A}$, i.e., to assign values to the variables so as to maximize the number m_C of satisfied equations. Håstad [10] showed, using the powerful theory of probabilistically checkable proofs (PCPs), that it is NP-hard to estimate m_C within any constant factor smaller than $12/11$.[1] Therefore consider an instance $\mathcal{A}$ of E2LIN2, as above. We will construct (by a randomized algorithm) a graph $G = (V, E)$, regular of degree Δ. We then show that, if we can approximate the logarithm of the number of independent sets in G to within the required relative error, we can (with high probability) approximate the size of m_C in $\mathcal{A}$ to within a factor $12/11 - \varepsilon$. Theorem 3.1 will then follow.

[1] In other words, determining a number in the range $[(11/12+\varepsilon)m_C, m_C]$ is as hard as determining m_C exactly. Following convention, Håstad normalizes approximation ratios to be greater than 1, taking the reciprocal in the case of a maximization problem.

Let us write $[n] = \{1, 2, \ldots, n\}$. We construct the graph $G = G(\mathcal{A})$ as follows. We assume $m \geq n$; otherwise, $\mathcal{A}$ is decomposable or consistent. Let $M = m^6$ and, for each $i \in [n]$, let A_i be the multiset of equations containing x_i, with (multiset) cardinality d_i. We represent x_i by a regular bipartite graph H_i of degree $\delta = \Delta - 1$, with vertex partition (L_i, R_i) and edge set F_i. Here $L_i = \bigcup_{a \in A_i} L_{i,a}$, $R_i = \bigcup_{a \in A_i} R_{i,a}$, where the sets $L_{i,a}$, $R_{i,a}$ partition L_i and R_i, respectively, and for all i, a, $|L_{i,a}| = |R_{i,a}| = M$. Thus H_i is bipartite with both its vertex sets of size Md_i. Later, we will associate L_i with the assignment $x_i = 0$, and R_i with $x_i = 1$.

The graph $H_i = (L_i, R_i, F_i)$ will be sampled from $\mathcal{G}(Md_i, Md_i, \delta)$, where $\mathcal{G}$ is the class of random graphs defined in section 2. Just as in that section, and for the same reason, we are at liberty to reject graphs which are not δ-regular. Clearly, the property of being δ-regular can be checked in polynomial time.

The equations a of $\mathcal{A}$ determine the edges connecting the H_i in G, as follows. If a is the equation $x_i + x_j = 1$ ($x_i + x_j = 0$, resp.), we add an arbitrary perfect matching between $L_{i,a}$ and $L_{j,a}$ ($R_{j,a}$, resp.) and another between $R_{i,a}$ and $R_{j,a}$ ($L_{j,a}$, resp.). Thus G is a regular graph of degree Δ. We will show that approximating the logarithm of the number of independent sets in G to within a factor $(1 + 10^{-4})$ will allow us to approximate the E2LIN2 instance within the Håstad bound.

Before returning to the issue of approximation, we will need to establish some crucial properties of the "typical" independent set in G. For this purpose, let I be sampled uniformly from $\mathcal{I}(G)$, the set of all independent sets in G. First we show that I "occupies about half the available space" in each $L_{i,a}$ or $R_{i,a}$.

Let $\mathcal{L}_{i,a}$ be the set of vertices in $L_{i,a}$ with no neighbor in I, and let $\mathcal{L}_i = \bigcup_{a \in A_i} \mathcal{L}_{i,a}$.

LEMMA 3.2. *Suppose that I is sampled uniformly at random (u.a.r.) from $\mathcal{I}(G)$. Then, except for probability $e^{-\Omega(m^2)}$, either $|\mathcal{L}_{i,a}| < m^4$ or $|\mathcal{L}_{i,a}| = (2 \pm O(1/m)) \times |I \cap L_{i,a}|$.*

Proof. If we condition on $I \cap (V \setminus L_{i,a})$, then $I \cap L_{i,a}$ is a random subset of $\mathcal{L}_{i,a}$. If $|\mathcal{L}_{i,a}| \geq m^4$, then Chernoff's bound implies that

$$\mathbf{Pr}\left[|I \cap L_{i,a}| \notin \frac{1}{2}\Big(1 \pm \frac{1}{m}\Big)\,|\mathcal{L}_{i,a}|\right] \leq 2\exp\Big(-\frac{1}{3}m^2\Big),$$

from which the lemma follows. ∎

Clearly, we may define $\mathcal{R}_{i,a}$ and $\mathcal{R}_i$ symmetrically and prove an analogous result. It is also clear that we may claim Lemma 3.2 for all i, a simultaneously, since there are fewer than m^2 such pairs. Now imagine that we choose an independent set $I \in \mathcal{I}(G)$ u.a.r. in two steps: first the part of I that lies outside H_i, followed by the restriction of I to H_i. We now deduce from Lemma 3.2 that, with high probability, at least around half of L_i is "available" to I in the second step.

Let $\mathcal{L}'_i$ be the set of vertices in L_i with no neighbor in I outside of H_i.

LEMMA 3.3. *Suppose that I is sampled u.a.r. from $\mathcal{I}(G)$. Except for probability $e^{-\Omega(m^2)}$,*

$$|\mathcal{L}'_i| \geq \left(\frac{1}{2} - O\left(\frac{1}{m}\right)\right)|L_i|. \tag{5}$$

Proof. If $L_{i,a}$ is joined by a matching to $V_{j,a}$ ($V \in \{L, R\}$), then, from Lemma 3.2, $M \geq (2 - O(1/m))|I \cap V_{j,a}|$. Hence

$$\big|\{v \in L_{i,a} : \{v, w\} \in E \setminus F_i \text{ implies } w \notin I\}\big| \geq \left(\frac{1}{2} - O\left(\frac{1}{m}\right)\right)|L_{i,a}|.$$

Summing this over all $a \in A_i$ gives the lemma. ∎

Again, we may define $\mathcal{R}_i'$ and prove a corresponding result. We now show that for each i either $|\mathcal{L}_i|$ or $|\mathcal{R}_i|$ is "small." We will temporarily drop the suffix i and write H rather than H_i, etc. Let $N = |L| = dM \le m^7$, $a = |\mathcal{L}'|/N$, $b = |\mathcal{R}'|/N$. Write $\sigma = I \cap H$, where I is a uniformly chosen independent set in G. We will say that σ is an (α, β)-set if $|\sigma \cap L| = \alpha a N$, $|\sigma \cap R| = \beta b N$.

LEMMA 3.4. *Let $\delta \ge 24$. If I is a uniformly chosen independent set in G, then, except for probability $e^{-\Omega(m^2)}$,*

$$\min(|\mathcal{L}_i|, |\mathcal{R}_i|) \le \lambda N, \tag{6}$$

where $\lambda = 0.009$.

Proof. We focus attention on a particular H in G (corresponding to a particular variable in the E2LIN2 instance). Suppose that the whole of G aside from the edges within H has been fixed (i.e., the random choices have already been made), except that we have not chosen the edges of H itself. Ultimately, we want to argue about a random independent set I. However, for the time being, suppose that we simply fix the portion of I that lies outside of H; doing this fixes the sets $\mathcal{L}'$ and $\mathcal{R}'$ of vertices in H that have no neighbor in I. About I, we assume only that it satisfies inequality (5) of Lemma 3.3 so that $a \ge b \ge \frac{1}{2} - O(\frac{1}{m})$, where, without loss of generality, we have taken $a \ge b$.

We now reveal H and examine the number of extensions of I to H as a function of α and β. It is easy to see that there are at least 2^{aN} independent sets in H in total. We will show that, for α, β not satisfying the condition of the lemma, the number of (α, β)-sets is so much smaller than this that they appear with probability $e^{-\Omega(m^2)}$. It will be sufficient to show that the expected number of (α, β)-sets in such a case is $2^{aN - \Omega(m^2)}$, because Markov's inequality will then imply the required inequality for the actual number. Now the expected number of (α, β)-sets in H is

$$\begin{aligned}
\mathcal{E}(\alpha,\beta) &= \binom{aN}{\alpha a N}\binom{bN}{\beta b N}\left[\binom{(1-b\beta)N}{\alpha a N}\Big/\binom{N}{\alpha a N}\right]^{\delta} \\
&\le \binom{aN}{\alpha a N}\binom{bN}{\beta b N}\left[\frac{[(1-b\beta)N]^{\underline{\alpha a N}}}{(\alpha a N)!} \times \frac{(\alpha a N)!}{N^{\underline{\alpha a N}}}\right]^{\delta} \\
&\le \binom{aN}{\alpha a N}\binom{bN}{\beta b N}(1-b\beta)^{\alpha a \delta N} \\
&\le \left[\left(\alpha^{\alpha}(1-\alpha)^{(1-\alpha)}\right)^{-a}\left(\beta^{\beta}(1-\beta)^{(1-\beta)}\right)^{-b} e^{-\alpha\beta a b \delta}\right]^{N(1+o(1))} \\
&= e^{\psi(\alpha,\beta)N(1+o(1))}, \qquad (7)
\end{aligned}$$

where an underlined superscript denotes "falling factorial power," and

$$\begin{aligned}
\psi(\alpha,\beta) = &-a(\alpha\ln\alpha + (1-\alpha)\ln(1-\alpha)) \\
&- b(\beta\ln\beta + (1-\beta)\ln(1-\beta)) - \alpha\beta a b\delta. \qquad (8)
\end{aligned}$$

Note that ψ is defined in the unit square $\mathcal{U} = \{(\alpha, \beta) : 0 \le \alpha, \beta \le 1\}$. As before, we shall treat α and β (and indeed a and b) as real variables, even though a combinatorial interpretation requires aN, bN, $\alpha a N$, and $\beta b N$ to be integers. The key property of ψ is captured in the following claim, whose proof can be found in the appendix.

CLAIM 3.5. *Let $\delta = 24$, $\eta > 0$ be sufficiently small, and suppose $\frac{1}{2} - \eta \le b \le a \le 1$. For any $(\alpha, \beta) \in \mathcal{U}$, the inequality $\psi(\alpha, \beta) \ge a \ln 2 - \eta$ entails $\min\{\alpha a, \beta b\} \le 0.004$.*

Recall the crude lower bound 2^{aN} on the total number of independent sets σ extending I to H. The claim tells us that only very unbalanced independent sets—those with either $|\sigma \cap L| \le 0.004$ or $|\sigma \cap R| \le 0.004$—make a significant contribution to that total. All of the above argument was for an independent set I that is fixed outside H, so we have not yet proved Lemma 3.4. Nevertheless, all the key calculations are out of the way, and we can complete the proof with a little algebra.

Let $\mathcal{I}$ be the set of all independent sets on $V(G) \setminus V(H)$. Let $\mathcal{I}_{\text{good}} \subseteq \mathcal{I}$ be the independent sets I that satisfy inequality (5) of Lemma 3.3, and $\mathcal{I}_{\text{bad}} = \mathcal{I} \setminus \mathcal{I}_{\text{good}}$. Let $N(I, H)$ be the number of independent sets in H consistent with I, and let $N^*(I, H)$ be the number of such that do not satisfy inequality (6) of Lemma 3.4. Denote by $\mathcal{H}$ the (multi)set of all possible choices for the graph H viewed as a disjoint union of δ perfect matchings. (Thus each possible graph H will occur with multiplicity μ, where μ is the number of 1-factorizations of H—i.e., decompositions into disjoint perfect matchings—of H. Note that our reduction requires us to select *uniformly* from $\mathcal{H}$.) For convenience, set $\varepsilon = e^{-\Omega(m^2)}$. We have shown in Lemma 3.3 that

$$\sum_{I \in \mathcal{I}_{\text{bad}}} N(I, H) \le \varepsilon \sum_{I \in \mathcal{I}} N(I, H) \qquad \text{for all } H \in \mathcal{H}. \tag{9}$$

(Note that the sum on the right-hand side is the total number of independent sets in G, while that on the left-hand side is the number violating inequality (5).) We will show below that a random H satisfies

$$\sum_{I \in \mathcal{I}_{\text{good}}} N^*(I, H) \le \varepsilon \sum_{I \in \mathcal{I}_{\text{good}}} N(I, H) \tag{10}$$

with high probability, specifically, with probability at least $1 - \varepsilon$. Putting (9) and (10) together, a random H satisfies

$$\begin{aligned} \frac{\sum_{I \in \mathcal{I}} N^*(I, H)}{\sum_{I \in \mathcal{I}} N(I, H)} &\le \frac{\varepsilon \sum_{I \in \mathcal{I}_{\text{good}}} N(I, H) + \sum_{I \in \mathcal{I}_{\text{bad}}} N(I, H)}{\sum_{I \in \mathcal{I}} N(I, H)} \\ &\le \varepsilon + \varepsilon = 2\varepsilon \end{aligned}$$

with high probability, which is what we require.

We now prove (10). Claim 3.5 taken in conjunction with Lemma 3.2 shows that

$$\frac{\sum_{H \in \mathcal{H}} N^*(I, H)}{|\mathcal{H}|} \le \varepsilon^2 \widehat{N}(I) \qquad (I \in \mathcal{I}_{\text{good}})$$

for some $\widehat{N}(I)$ satisfying $\widehat{N}(I) \le N(I, H)$ for all $H \in \mathcal{H}$. (Specifically, $\widehat{N} = 2^{aN}$ will do here.) Summing this over $I \in \mathcal{I}_{\text{good}}$ gives

$$\frac{1}{|\mathcal{H}|} \sum_{H \in \mathcal{H}} \sum_{I \in \mathcal{I}_{\text{good}}} N^*(I, H) \le \varepsilon^2 \sum_{I \in \mathcal{I}_{\text{good}}} \widehat{N}(I),$$

giving

$$\frac{1}{|\mathcal{H}|} \sum_{H \in \mathcal{H}} \frac{\sum_{I \in \mathcal{I}_{\text{good}}} N^*(I, H)}{\sum_{I \in \mathcal{I}_{\text{good}}} \widehat{N}(I)} \le \varepsilon^2,$$

which implies that

$$\frac{1}{|\mathcal{H}|}\sum_{H\in\mathcal{H}}\frac{\sum_{I\in\mathcal{I}_{\text{good}}}N^*(I,H)}{\sum_{I\in\mathcal{I}_{\text{good}}}N(I,H)}\le\varepsilon^2. \tag{11}$$

Let

$$\mathcal{H}^*=\left\{H\in\mathcal{H}:\sum_{I\in\mathcal{I}_{\text{good}}}N^*(I,H)\ge\varepsilon\sum_{I\in\mathcal{I}_{\text{good}}}N(I,H)\right\}.$$

Then, from (11),

$$\frac{1}{|\mathcal{H}|}\,\varepsilon\,|\mathcal{H}^*|\le\varepsilon^2,$$

so

$$\frac{|\mathcal{H}^*|}{|\mathcal{H}|}\le\varepsilon,$$

as is required to establish (10) and complete the proof. □

We now establish the relationship between the number of independent sets in G and the maximum size of a consistent subset of $\mathcal{A}$. Let $\mathcal{I}=\mathcal{I}(G)$. For $\sigma\in\mathcal{I}$ let $S_\sigma\subseteq[n]$ be defined by

$$S_\sigma=\{i:\ |L_i\cap\sigma|>|R_i\cap\sigma|,\ i\in[n]\}.$$

For $S\subseteq[n]$ let $\mathcal{I}_S=\{\sigma\in\mathcal{I}:\ S_\sigma=S\}$ and let $\mu_S=|\mathcal{I}_S|$. Recall that m is the number of equations in $\mathcal{A}$.

LEMMA 3.6. *For $S\subseteq[n]$ let $\theta(S)$ be the number of equations in $\mathcal{A}$ satisfied by the assignment $x_i=1$ $(i\in S)$, $x_i=0$ $(i\notin S)$. Then*

$$4^{M\theta(S)}3^{M(m-\theta(S))}\le\mu_S\le4^{M\theta(S)}3^{M(m-\theta(S))}2^{2\lambda mM}(1+o(1)), \tag{12}$$

where λ is as in Lemma 3.4.

Proof. Fix $S\subseteq[n]$, and for $\sigma\in\mathcal{I}_S$ let $J_\sigma=\sigma\cap\left(\bigcup_{i\in S}L_i\cup\bigcup_{i\notin S}R_i\right)$. Informally, J_σ restricts σ to the left or right of each subgraph H_i, according to which side contains the larger part of σ. Let

$$\hat{\mu}_S=|\{J_\sigma:\ \sigma\in\mathcal{I}_S\}|\le\mu_S.$$

We show that

$$\hat{\mu}_S=4^{M\theta(S)}3^{M(m-\theta(S))}. \tag{13}$$

This immediately proves the lower bound in (12). Furthermore, Lemma 3.4 implies that for a fixed value J of J_σ there are (up to a factor $(1+e^{-\Omega(m^2)})$) at most

$$\prod_{i\in[n]}2^{\lambda d_iM}=2^{\lambda M\sum_i d_i}=2^{2\lambda mM}$$

sets $\sigma\in\mathcal{I}_S$ with $J_\sigma=J$. The upper bound then follows.

To prove (13) we consider the number of possible choices for $J \cap L_{i,a}$, $J \cap R_{i,a}$, $J \cap L_{j,a}$, and $J \cap R_{j,a}$ for every equation $a : x_i + x_j = z_a$ ($z_a \in \{0,1\}$). For given S, let us define

$$X_{i,a} = \begin{cases} L_{i,a} & \text{if } i \in S; \\ R_{i,a} & \text{if } i \notin S. \end{cases}$$

Then there are two cases, determined by the status of a.

(1) Equation a is satisfied by the assignment derived from S. Then there are 2^M choices for each of $J \cap X_{i,a}$, $J \cap X_{j,a}$, giving 4^M in all.
(2) Equation a is not satisfied by the assignment derived from S. Then the subgraph of G induced by $X_{i,a} \cup X_{j,a}$ is a matching of size M and hence contains 3^M independent sets.

Multiplying the estimates from the two cases over all $a \in \mathcal{A}$ proves (13) and the lemma. □

We now proceed to the proof of Theorem 3.1. Let $Z_I = Z_I(G)$ denote the logarithm of the number of independent sets of $G(\mathcal{A})$. Let $Z_C = Z_C(\mathcal{A})$ denote the maximum number of consistent equations in $\mathcal{A}$.

Let Y_I be some estimate of Z_I satisfying $|Y_I/Z_I - 1| \le \varepsilon = 10^{-4}$. Using Y_I, we define

$$Y_C = \left(\frac{Y_I}{M} - m \ln 3\right) \frac{1.001}{\ln(4/3)}.$$

A simple calculation will then show that $1 \le Y_C/Z_C \le 12/11 - \varepsilon$, so that Y_C determines Z_C with sufficient accuracy to beat the approximability bound for E2LIN2.

From Lemma 3.6 we see that

$$Y_I \ge (1-\varepsilon)M(Z_C \ln(4/3) + m \ln 3).$$

Hence, since $Z_C \ge m/2$,

$$\frac{Y_C}{1.001} \ge (1-\varepsilon)Z_C - \frac{\varepsilon m \ln 3}{\ln(4/3)} \ge Z_C\left(1 - \frac{\varepsilon \ln 12}{\ln(4/3)}\right) \ge 0.9991 Z_C,$$

which implies that $Y_C \ge Z_C$. On the other hand, Lemma 3.6 also implies that

$$Y_I \le (1+\varepsilon)\big[M\big(Z_C \ln(4/3) + m \ln 3 + 2m\lambda \ln 2\big) + n \ln 2\big],$$

where $\lambda \le 0.009$. Hence

$$\begin{aligned} \frac{Y_C}{1.001} &\le (1+\varepsilon)Z_C + \frac{\varepsilon m \ln 3}{\ln(4/3)} + \frac{(1+\varepsilon)2m\lambda \ln 2}{\ln(4/3)} + \frac{(1+\varepsilon)\ln 2}{n \ln(4/3)} \\ &\le Z_C\left(1 + \varepsilon + \frac{2\varepsilon \ln 3}{\ln(4/3)} + \frac{4(\ln 2)(1+\varepsilon)\lambda}{\ln(4/3)} + O\left(\frac{1}{n^2}\right)\right) \\ &\le Z_C\left(1.0877 + O\left(\frac{1}{n^2}\right)\right), \end{aligned}$$

which implies that Y_C/Z_C is bounded away from $12/11$ for n large enough. Summarizing, the existence of a polynomial-time algorithm, meeting the specification in Theorem 3.1, for estimating the number of independent sets in a 25-regular graph

would entail the existence of a randomized (two-sided error) algorithm for approximating the solution to an E2LIN2 instance with relative error better than 12/11. (The algorithm is randomized because the reduction is too.) Because the latter problem is NP-hard, we could deduce that NP $\subseteq$ BPP. But this inclusion in turn implies that RP = NP (see Papadimitriou [11, Problem 11.5.18]). Thus we have established Theorem 3.1.

Appendix.

Proof of Claim 2.2. We start by computing partial derivatives of φ up to order two:

$$\frac{\partial\varphi}{\partial\alpha} = -\ln\alpha - (\Delta-1)\ln(1-\alpha) + \Delta\ln(1-\alpha-\beta), \tag{14}$$

$$\frac{\partial\varphi}{\partial\beta} = -\ln\beta - (\Delta-1)\ln(1-\beta) + \Delta\ln(1-\alpha-\beta), \tag{15}$$

$$\frac{\partial^2\varphi}{\partial\alpha^2} = -\frac{1}{\alpha} + \frac{\Delta-1}{1-\alpha} - \frac{\Delta}{1-\alpha-\beta}, \tag{16}$$

$$\frac{\partial^2\varphi}{\partial\beta^2} = -\frac{1}{\beta} + \frac{\Delta-1}{1-\beta} - \frac{\Delta}{1-\alpha-\beta}, \tag{17}$$

$$\frac{\partial^2\varphi}{\partial\alpha\,\partial\beta} = -\frac{\Delta}{1-\alpha-\beta}. \tag{18}$$

Parts (i)–(iv) of Claim 2.2 may then be verified as follows:

(i) From (16), it can easily be checked that $\partial^2\varphi/\partial\alpha^2 < 0$ on the interior of $\mathcal{T}$, and hence φ can have no interior local minima. On $\alpha = 0$, φ has a maximum at $\beta = \frac{1}{2}$ using (15), but then from (14) we find $\partial\varphi/\partial\alpha = +\infty$ at $\alpha = 0$, $\beta = \frac{1}{2}$. Similarly $\beta = 0$. On $\alpha + \beta = 1$, both $\partial\varphi/\partial\alpha, \partial\varphi/\partial\beta = -\infty$, so φ can have no maximum.

(ii) Since both $\partial^2\varphi/\partial\alpha^2, \partial^2\varphi/\partial\beta^2 < 0$, φ has a maximum if and only if the Hessian of φ has a positive determinant. The condition for this is $\alpha + \beta + \Delta(\Delta-2)\alpha\beta \le 1$, as may be checked from (16)–(18).

(iii) From (14) and (15), the conditions for a stationary point of φ may be written

$$\beta = f(\alpha), \quad \alpha = f(\beta),$$

where

$$f(x) = 1 - x - x^{1/\Delta}(1-x)^{1-1/\Delta} = (1-x)\left[1 - \left(\frac{x}{1-x}\right)^{1/\Delta}\right] \quad (0 \le x \le 1).$$

Thus, at any stationary point,

$$\alpha = f(f(\alpha)). \tag{19}$$

Clearly $f(x) \le 0$ for $x \ge \frac{1}{2}$, so $\alpha < \frac{1}{2}$ at any stationary point. Similarly $\beta < \frac{1}{2}$. To study the roots of (19), the change of variable $y = (\alpha/(1-\alpha))^{1/\Delta}$ proves to be convenient. With a little calculation we may express α, $f(\alpha)$, and $f(f(\alpha))$ in terms of y:

$$\alpha = \frac{y^\Delta}{1+y^\Delta}, \tag{20}$$

$$f(\alpha) = (1-\alpha)(1-y) = \frac{1-y}{1+y^\Delta},$$

and

$$\begin{aligned} f(f(\alpha)) &= (1 - f(\alpha)) - \big(f(\alpha)(1 - f(\alpha))^{\Delta-1}\big)^{1/\Delta} \\ &= \left(\alpha + \frac{y}{1+y^\Delta}\right) - \frac{\big((1-y)(y+y^\Delta)^{\Delta-1}\big)^{1/\Delta}}{1+y^\Delta} \\ &= \alpha + \frac{y}{1+y^\Delta}\left[1 - \left(\frac{(1-y)(1+y^{\Delta-1})^{\Delta-1}}{y}\right)^{1/\Delta}\right], \end{aligned}$$

and hence (19) is equivalent to

$$(1 + y^{\Delta-1})^{\Delta-1} = \frac{y}{1-y} \qquad (0 \le y < 1). \tag{21}$$

Note that the implicit mapping from α to y is a bijection, so we may legitimately study the solution set of (19) through that of (21). Note also that (21) has a root y' satisfying $y + y^\Delta = 1$, and this exists for any $\Delta > 0$. The reader may check that $y + y^\Delta = 1$ is equivalent to $\alpha = f(\alpha)$, and thus y' satisfies $\alpha = \beta$. To analyze (21) in general, let

$$g(y) = (\Delta - 1)\ln(1 + y^{\Delta-1}) + \ln(1 - y) - \ln y,$$

so $g(y) = 0$ has the same roots as (21). Then one may check that $g'(y) = 0$ if and only if

$$h(y) \stackrel{\text{def}}{=} \Delta(\Delta - 2)y^{\Delta-1} - (\Delta - 1)^2 y^\Delta - 1 = 0.$$

But $h(0) = -1$, $h(1) = -2$, and h has a single maximum on $[0, 1]$ at $y'' = (\Delta - 2)/(\Delta - 1)$. Now $h(y'') = (\Delta - 2)^\Delta/(\Delta - 1)^{\Delta-1} - 1 > 0$ if and only if $\Delta \ge 6$, and $h(y'') < 0$ otherwise. Therefore h has two roots in $[0, 1]$ if $\Delta \ge 6$; otherwise, it has no roots. Thus g has a single root in $[0, 1]$ if $\Delta \le 5$; otherwise, it has at most three roots. In the latter case, however, $g(0) = +\infty$, $g(1) = -\infty$, $g(y') = 0$, and a simple calculation shows

$$g'(y') = \frac{(\Delta - 1)^2(1 - y')^2 - 1}{y'(1 - y')} > 0$$

if and only if $\Delta \ge 6$, and $g'(y') < 0$ otherwise. These facts imply that g has exactly three roots if $\Delta \ge 6$.

Now the reader may check that the point (α', α') corresponding to y' (i.e., given by solving $y' = (\alpha'/(1 - \alpha'))^{1/\Delta}$) satisfies

$$\alpha + \beta + \Delta(\Delta - 2)\alpha\beta \le 1,$$

i.e.,

$$\left(\frac{1-\alpha}{\alpha}\right)\left(\frac{1-\beta}{\beta}\right) \ge (\Delta - 1)^2$$

if and only if $y' \ge y''$. This holds if and only if $\Delta \le 5$. Thus this point is a maximum for $\Delta \le 5$; otherwise, it is a saddle-point.

Thus φ has one stationary point in $\mathcal{T}$ (on $\alpha = \beta$) if $\Delta \le 5$, and this is a maximum.

(iv) By the above, if $\Delta \geq 6$, φ has no boundary maximum on $\mathcal{T}' = \{(\alpha, \beta) \in \mathcal{T} : \alpha \leq \beta\}$ and therefore by continuity has a maximum in the interior of $\mathcal{T}'$. By symmetry there is also a maximum in $\mathcal{T} \setminus \mathcal{T}'$. Thus, when $\Delta \geq 6$, φ has two symmetrical maxima and a single saddle-point on the line $\alpha = \beta$. Numerical values for the two maximum points can be obtained by solving (21) for y. Since we are assured that (21) has exactly three roots, we may locate these roots to arbitrary precision by repeated function evaluations. Once y is known to adequate precision, α can be recovered from (20). □

Proof of Claim 2.3. Let $\Omega = \{1, \ldots, N\}$ be an enumeration of the state space. When x is an N-vector and P an $N \times N$ matrix, we will use x_A to mean the vector $(x_i : i \in A)$ and P_{AB} to mean the matrix $(P_{ij} : i \in A, j \in B)$. First note that

$$\begin{aligned} \mathrm{d_{TV}}(p_{t+1}, p_t) = \mathrm{d_{TV}}(p_t P, p_{t-1} P) &= \tfrac{1}{2} \max_{||z||_\infty \leq 1} (p_t - p_{t-1}) P z \\ &\leq \tfrac{1}{2} \max_{||w||_\infty \leq 1} (p_t - p_{t-1}) w = \mathrm{d_{TV}}(p_t, p_{t-1}), \end{aligned}$$

since $||Pz||_\infty \leq ||z||_\infty$. Hence, by induction, $\mathrm{d_{TV}}(p_{t+1}, p_t) \leq \mathrm{d_{TV}}(p_1, p_0)$ and hence, using the triangle inequality, $\mathrm{d_{TV}}(p_t, p_0) \leq t\, \mathrm{d_{TV}}(p_1, p_0)$. Now, for $\emptyset \subset S \subset \Omega$, define

$$\Phi(S) = \sum_{i \in S} \sum_{j \in \overline{S}} \pi_i P_{ij} / \pi(S).$$

Thus $\Phi = \min\{\Phi(S) : S \subset \Omega \text{ and } 0 < \pi(S) \leq \frac{1}{2}\}$ is the "conductance" of $\mathcal{M}$. (Conductance is normally considered in the context of time-reversible Markov chains. However, both the definition and the line of argument employed here apply to non–time-reversible chains.) Now

$$\sum_{\substack{i \in A \\ j \in \overline{A}}} \pi_i P_{ij} \leq \sum_{\substack{i \in A \\ j \in \overline{A} \cap M}} \pi_i P_{ij} + \sum_{\substack{i \in A \cap M \\ j \in \overline{A}}} \pi_i P_{ij} \leq \pi(\overline{A} \cap M) + \pi(A \cap M) = \pi(M).$$

So by setting $(p_0)_A = \pi_A / \pi(A)$, $(p_0)_{\overline{A}} = 0$, we have that

$$\mathrm{d_{TV}}(p_1, p_0) = \tfrac{1}{2} ||\pi_A - \pi_A P||_1 / ||\pi_A||_1 = ||\pi_A P_{A\overline{A}}||_1 / ||\pi_A||_1 = \Phi(A) \leq \pi(M)/\pi(A).$$

But $\mathrm{d_{TV}}(\pi, p_0) \geq \frac{1}{2}$, because $\pi(A) \leq \frac{1}{2}$, and hence

$$\mathrm{d_{TV}}(\pi, p_t) \geq \mathrm{d_{TV}}(\pi, p_0) - \mathrm{d_{TV}}(p_t, p_0) \geq \tfrac{1}{2} - t\Phi(A).$$

Thus we cannot achieve $\mathrm{d_{TV}}(\pi, p_t) \leq e^{-1}$ until

$$t \geq (\tfrac{1}{2} - e^{-1})/\Phi \geq \pi(A)/8\pi(M).$$

By an averaging argument there must exist some initial state $x_0 \in A$ for which $\tau(x_0) \geq \pi(A)/8\pi(M)$. □

Proof of Claim 3.5. Differentiating (8), we have

$$\begin{aligned} \frac{\partial \psi}{\partial \alpha} &= a(-\ln \alpha + \ln(1-\alpha) - \beta b \delta), \\ \frac{\partial \psi}{\partial \beta} &= b(-\ln \beta + \ln(1-\beta) - \alpha a \delta), \end{aligned} \tag{22}$$

and

$$\frac{\partial^2\psi}{\partial\alpha^2} = \frac{-a}{\alpha(1-\alpha)}, \qquad \frac{\partial^2\psi}{\partial\beta^2} = \frac{-b}{\beta(1-\beta)}, \qquad \frac{\partial^2\psi}{\partial\alpha\,\partial\beta} = -ab\delta. \tag{23}$$

The following three facts about ψ are easily verified:

$$\psi(\alpha,\beta) \geq \psi(1-\alpha,\beta) \qquad \text{if } \alpha \leq \tfrac{1}{2}, \tag{24}$$

$$\psi(\alpha,\beta) \geq \psi(\alpha,1-\beta) \qquad \text{if } \beta \leq \tfrac{1}{2}, \tag{25}$$

$$\psi(\alpha,\beta) \geq \psi(\beta,\alpha) \qquad \text{if } \beta \leq \alpha \leq 1-\beta. \tag{26}$$

We wish to determine the regions where $\psi \geq a\ln 2$. These are connected neighborhoods of the local maxima of ψ. From (22) we see that ψ has no boundary maxima for α,β in the unit square $\mathcal{U}$. Thus, from (23), ψ has only local maxima or saddle-points in $\mathcal{U}$, and a stationary point is a local maximum if and only if

$$\alpha(1-\alpha)\beta(1-\beta) \leq 1/(ab\delta^2). \tag{27}$$

Thus, at any local maximum, either $\beta(1-\beta) \leq 1/(b\delta)$ or $\alpha(1-\alpha) \leq 1/(a\delta)$. If the former holds, this and $b\delta \geq 11.5$ (which holds for η sufficiently small) imply that $\beta < 0.1$, and hence $\beta < 1.2/b\delta$. An identical argument holds for α. Let us denote the rectangle $[\ell_\alpha, u_\alpha] \times [\ell_\beta, u_\beta]$ by $[\ell_\alpha, u_\alpha \,|\, \ell_\beta, u_\beta]$. Thus any local maximum of ψ must lie in the region $[0,1\,|\,0,1.2/b\delta] \cup [0,1.2/a\delta\,|\,0,1]$ and hence in the enclosing region $[0,1\,|\,0,1.2/b\delta] \cup [0,1.2/b\delta\,|\,0,1]$. (Recall that $a \geq b$.) In the square $[0,1.2/b\delta\,|\,0,1.2/b\delta]$, we have $\alpha,\beta \leq 1.2/b\delta < 0.11$ and hence

$$\psi(\alpha,\beta) < 2a(-0.11\ln(0.11) - 0.89\ln(0.89)) < a\ln 2.$$

Then, from (24) and (25), we also have $\psi(\alpha,\beta) < a\ln 2$ in $[1-1.2/b\delta, 1\,|\,0,1.2/b\delta]$ and $[0,1.2/b\delta\,|\,1-1.2/b\delta,1]$. Now, if $\beta \leq 1.2/b\delta$, let $\rho = 1-2\alpha$ and consider the upper bound

$$\psi(\alpha,\beta) \leq \Psi(\rho,\beta) \stackrel{\text{def}}{=} a(\ln 2 - \tfrac{1}{2}\rho^2) + b\beta(1-\ln\beta) - \tfrac{1}{2}(1-\rho)\beta ab\delta. \tag{28}$$

For fixed β, it is easily shown that Ψ is maximized if $\rho = \frac{1}{2}b\delta\beta \leq 0.6$. If $b\delta\beta = 1.2$, then $\rho = 0.6$ and

$$\max_\rho \Psi(\rho,\beta) \leq a(\ln 2 - 0.18) + 0.11a(1-\ln(0.11)) - 0.24a < a\ln 2.$$

Thus $\psi < a\ln 2$ everywhere on the boundary of $[1.2/b\delta, 1-1.2/b\delta\,|\,0,1.2/b\delta]$ (but not including the shared boundary with $\mathcal{U}$). Hence, by (26), $\psi < a\ln 2$ everywhere on the boundary of $[0,1.2/b\delta\,|\,1.2/b\delta, 1-1.2/b\delta]$. Moreover, $\psi(\alpha,\beta) \geq \psi(\beta,\alpha)$ for all points (α,β) in $[1.2/b\delta, 1-1.2/b\delta\,|\,0,1.2/b\delta]$. It follows that it is sufficient to determine β^* such that $\psi(\alpha,\beta) < a\ln 2$ everywhere in $[1.2/b\delta, 1-1.2/b\delta\,|\,\beta^*, 1.2/b\delta]$. To this end, again consider

$$\Psi_0(\beta) = \max_\rho \Psi(\rho,\beta) = a\ln 2 + b\beta(1-\ln\beta) - \tfrac{1}{2}ab\beta\delta + \tfrac{1}{8}ab^2\beta^2\delta^2.$$

Now $\Psi_0 < a\ln 2$ if

$$b\beta\delta^2 - 4\delta + 8(1-\ln\beta)/a < 0.$$

This inequality is satisfied, provided

$$2\left(1-\sqrt{1-2b\beta(1-\ln\beta)/a}\right) < \ b\beta\delta < \ 2\left(1+\sqrt{1-2b\beta(1-\ln\beta)/a}\right).$$

The right-hand inequality is clearly irrelevant since we are assuming that $\beta \le 1.2/b\delta$. Thus we need consider only the left-hand inequality; i.e., for fixed $\gamma = b\beta < 1.2/\delta$, we require that

$$\gamma\delta > 2\max_{a,b}\left(1-\sqrt{1-2\gamma(1-\ln\gamma+\ln b)/a}\right),$$

where the maximum is over $\frac{1}{2}-\eta \le b \le a \le 1$. Considering b first, the maximum occurs when $b=a$. So we have

$$\gamma\delta > \max_{\frac{1}{2}-\eta\le a\le 1} 2\left(1-\sqrt{1-2\gamma(1-\ln\gamma+\ln a)/a}\right). \tag{29}$$

But, because $a \ge \gamma$, the maximum now occurs when $a = \frac{1}{2}-\eta$. Thus it is enough to require that

$$\gamma\delta > 2\left(1-\sqrt{1-4\gamma(1-\ln\gamma-\ln 2)}\right),$$

because this will imply (29), provided that η is sufficiently small. To achieve $\gamma = 0.004$, it is sufficient that $\delta \ge 23.9$. □

REFERENCES

[1] E. A. Bender, *The asymptotic number of non-negative integer matrices with given row and column sums*, Discrete Math., 10 (1974), pp. 217–223.

[2] P. Berman and M. Karpinski, *On Some Tighter Inapproximability Results, Further Improvements*, Electronic Colloquium on Computational Complexity, Report TR98-065, 1998.

[3] G. Brightwell and P. Winkler, *Graph homomorphisms and phase transitions*, J. Combin. Theory Ser. B, 77 (1999), pp. 221–262.

[4] M. Dyer, A. Frieze, and M. Jerrum, *On counting independent sets in sparse graphs*, in Proceedings of the 40th IEEE Symposium on Foundations of Computer Science (FOCS'99), IEEE Computer Society Press, Los Alamitos, CA, 1999, pp. 210–217.

[5] M. Dyer and C. Greenhill, *On Markov chains for independent sets*, J. Algorithms, 35 (2000), pp. 17–49.

[6] M. Jerrum, L. Valiant, and V. Vazirani, *Random generation of combinatorial structures from a uniform distribution*, Theoret. Comput. Sci., 43 (1986), pp. 169–188.

[7] M. Jerrum, *Large cliques elude the Metropolis process*, Random Structures Algorithms, 3 (1992), pp. 347–359.

[8] M. Jerrum and A. Sinclair, *The Markov chain Monte Carlo method: An approach to approximate counting and integration*, in Approximation Algorithms for NP-Hard Problems, D. Hochbaum, ed., PWS Publishing, Boston, 1996, pp. 482–520.

[9] M. Luby and E. Vigoda, *Approximately counting up to four*, in Proceedings of the 29th Annual ACM Symposium on Theory of Computing, ACM Press, New York, 1997, pp. 682–687.

[10] J. Håstad, *Some optimal inapproximability results*, in Proceedings of the 29th Annual ACM Symposium on Theory of Computing, ACM Press, New York, 1997, pp. 1–10.

[11] C. H. Papadimitriou, *Computational Complexity*, Addison-Wesley, Reading, MA, 1994.