Lecture 20: November 5

Instructor: Alistair Sinclair *Scribes: Bhaskar Roberts & Kada Situ*

20.1 Strong spatial mixing

Recall from the last lecture that our goal is to approximate the partition function of the hard-core model, given by

$$Z_G(\lambda) := \sum_{\substack{I \subseteq V \\ ind.set}} \lambda^{|I|}, \tag{20.1}$$

where G is any graph with a maximum degree Δ. Last time we stated a theorem that claims that, if λ is less than the critical threshold value

$$\lambda < \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^{\Delta}} =: \lambda_c(\Delta), \tag{20.2}$$

then there exists a FPTAS for $Z_G(\lambda)$, and we described Weitz's algorithm to construct such an approximation scheme. In this lecture we will complete the proof of the theorem by giving an analysis of Weitz's algorithm.

Recall that, given a graph G of maximum degree Δ, Weitz's algorithm constructs the *self-avoiding walk tree* $T_{SAW}(G, v)$ rooted at a vertex v. This tree enumerates all self-avoiding walks starting from v, ending each walk at a leaf whenever a cycle is closed. The tree also assigns a *boundary condition* (occupied or unoccupied) to each such leaf depending on the direction in which the cycle was traversed. The key fact we proved last time is that the occupation probability of v under the Gibbs distribution on G is exactly the same as the occupation probability of the root in the SAW tree. This means that we can compute this probability in the tree rather than in the graph, which we can do using a simple recursive procedure.

However, the SAW tree will in general have exponential size (compared to the size of G itself), so we can't work with the entire tree. Instead, we want to *truncate* the tree at depth $O(\log n)$, so that its size will be $\Delta^{O\log n}$, which is polynomial in n, so we can do the computation on it. This will only work if the error introduced by the truncation is small. To prove that this is the case, we can appeal to the correlation decay property (because we are assuming $\lambda < \lambda_c(\Delta)$), which should ensure that the error at depth $O(\log n)$ is inverse polynomial.

There is, however, one crucial issue we have to sort out before we can appeal to correlation decay: we saw that $\lambda < \lambda_c(\Delta)$ implies exponential decay of correlations *in the Δ-regular tree without boundary conditions.* However, in the SAW tree there are boundary conditions, and some of them may be very close to the root (due to short cycles in G), so they will still be present after we truncate the tree.

To address this issue, we need a stronger notion of decay of corrections that holds even in the presence of boundary conditions close to the root.

Definition 20.1. *Let $S \subset V$ be any subset of vertices of the infinite tree, and τ be any fixed configuration on S. Let $p_r^\tau = \Pr[$root occupied given $\tau]$. Then* strong spatial mixing (SSM) *holds if there exists a constant*

$c > 0$ s.t.

$$|p_r^\tau - p_r^{\tau'}| \le \exp(-c \cdot dist(r, R)).$$

for any two configurations τ, τ' on S that agree on $S \setminus R$.

This definition is almost identical to that of Weak Spatial Mixing in the previous lecture, except that SSM requires the correlation to decay exponentially only with the distance to the *disagreeing portion* of the boundary conditions, even if τ, τ' have (the same) fixed spins close to the root. It's clear from the above definition that SSM implies WSM, since $\text{dist}(r, S) \le \text{dist}(r, R)$. But the converse is not true in general. Even though it may seem that the presence of additional (fixed) boundary conditions can only make the root less sensitive to distant boundary conditions, this is not the case, e.g., for the ferromagnetic Ising model, where it is possible to construct examples in which WSM holds but SSM does not. (Essentially the examples use boundary conditions to "shift" the occupation probability of the root into a range of values where it is *more* sensitive to the spins at the distant leaves; see, e.g., [SST14].)

The following theorem states that, in the case of the hard-core model, WSM in fact does imply SSM.

Theorem 20.2. *For the hard-core model on T_Δ, WSM $\implies$ SSM, i.e., SSM holds for all $\lambda < \lambda_c(\Delta)$.*

Before we prove this key theorem, let's complete the description and analysis of the algorithm assuming that SSM holds. Here is the algorithm:

1. Given G and a vertex v of G, construct the first ℓ levels of $T_{SAW}(G, v)$, where $\ell = \frac{1}{c} ln(\frac{5n}{\varepsilon})$, with c being the constant in the definition of SSM.
2. Assign unoccupancy probabilities to the leaves as follows. Note that there are two kinds of leaves: those that are created by truncation, and those that were leaves in the original tree (above the point where we do the truncation).
   - The leaf nodes that aren't created by truncation either already have a boundary condition, or they are free (in which case they were leaves in G). If a leaf is occupied we assign it value 0, and if it's unoccupied we assign it 1. (Recall that we're computing the probabilities that vertices are *un*occupied.) If a leaf is free we assign it the value $\frac{1}{1+\lambda}$, which is the probability that a vertex with no occupied neighbors is unoccupied. (Note that in fact the leaves with boundary conditions can be removed as follows: an unoccupied leaf can just be removed; an occupied leaf can be removed, together with its parent. This leaves just free leaves, each of which gets the above boundary condition. But of course the shape of the tree is highly non-uniform, reflecting the effect of the boundary conditions.)
   - The leaves that were created artificially by truncation in step 1 are assigned an arbitrary value in the range $[\frac{1}{1+\lambda}, 1]$, which is the range of legal unoccupation probabilities.
3. Use the tree recurrence on the SAW tree $T = T_{SAW}(G, v)$ to compute $\Pr_T[\sigma(v) = 0]$, the unoccupation probability of the root v.
4. Fix $\sigma(v) = 0$, and repeat the above for another vertex v' until we exhaust all vertices in G. (Note that the fixed spin $\sigma(v) = 0$ is just another boundary condition, so the algorithm works as before.)

In this way, we end up with an estimate of the probability that all vertices are unoccupied, which is just the probability of the empty set in the Gibbs measure $\pi[\emptyset]$. (The empty set is always an independent set.) But since $\pi[\emptyset] = 1/Z_G(\lambda)$, taking the reciprocal of our estimate immediately gives us an estimate for $Z_G(\lambda)$ with the same (multiplicative) accuracy. Furthermore, in order to get a $1\pm\varepsilon$ estimate of $Z_G(\lambda)$, we need to get an $1\pm\frac{\varepsilon}{n}$ estimate of each marginal $Pr[\sigma(v) = 0]$. If we can show that we can obtain each such estimate in polynomial time, then we will have an FPTAS. Note that this algorithm is deterministic.

Claim 20.3. *The running time of the above procedure is polynomial in n and ε^{-1}.*

Proof. The size of the the tree is

$$O(\Delta^l) = O\left(\Delta^{\frac{1}{c}ln(\frac{5n}{\varepsilon})}\right) = O\left((\frac{5n}{\varepsilon})^{\frac{1}{c}ln\Delta}\right)$$

Clearly, this is polynomial in n and $\frac{1}{\varepsilon}$ □

Claim 20.4. *The accuracy for each marginal is* $1\pm\frac{\varepsilon}{n}$.

Proof. Redefine p_v^τ to denote the *un*occupation probability of the root v with boundary condition τ. For any two sets of boundary conditions τ, τ' that differ only on the *truncated* leaves, and are legal on those leaves, by SSM we have

$$|p_v^\tau - p_v^{\tau'}| \leq \exp(-c \cdot l) = \exp\left(-c \cdot \frac{1}{c} \cdot \ln(\frac{5n}{\varepsilon})\right) = \frac{\varepsilon}{5n}.$$

But since the minimum unoccupation probability for any vertex is $\frac{1}{1+\lambda}$, we have

$$p_v^\tau, p_v^{\tau'} \geq \frac{1}{1+\lambda} \geq \frac{1}{5},$$

using the fact that $\lambda < \lambda_c(\Delta) \leq \lambda_c(3) = 4$. Hence

$$\left|\frac{p_v^\tau}{p_v^{\tau'}} - 1\right| \leq \frac{\varepsilon}{n}.$$

But now we're done because if *any* pair of legal boundary conditions give estimates that are within $\leq 1 + \frac{\varepsilon}{n}$ of each other, then the estimate obtained with any arbitrary legal boundary condition (which is what the algorithm outputs) must be within $\leq 1\pm\frac{\varepsilon}{n}$ of the true value. □

20.2 Proof of Theorem 20.2

In this section, we will prove Theorem 20.2, namely that on the Δ-regular tree, for $\lambda < \lambda_c(\Delta)$, WSM $\implies$ SSM. This is the main technical meat in the analysis of Weitz's algorithm. Weitz's original proof is ingenious but tailored to the hard-core model. Instead we use technology that has since proved useful in proving the SSM property for a wider range of models. This particular version is due to Restrepo *et al.* [RST+13] and is typical of these arguments.

Let $d = \Delta - 1$, so that we are working on the d-ary tree. Consider an arbitrary node v in the tree, along with its children $\{v_1, \dots, v_d\}$. Let p_v denote the probability that v is unoccupied. (We suppress the boundary condition τ for ease of notation.) As in the last lecture, we can write the following recurrence for p_v:

$$p_v = \frac{1}{1 + \lambda \prod_{i=1}^{d} p_{v_i}} =: f(p_{v_1}, \dots, p_{v_d}). \tag{20.3}$$

We want to show that the distance between any two valid computations on the tree decreases at a uniform rate under this recurrence, i.e., for some $\gamma > 0$:

$$|p_v - p'_v| \leq (1-\gamma) \max_i |p_{v_i} - p'_{v_i}|. \tag{20.4}$$

This will imply that

$$|p_r - p'_r| \leq (1-\gamma)^l \max_{\text{Leaves } L_i} |p_{L_i} - p'_{L_i}|,$$

which is exactly what we need to show for SSM.

Unfortunately, however, the stepwise decay property (20.4) doesn't hold in general for a uniform γ. But we can get around this by considering instead the decay of distance under some function of p_v.

Definition 20.5. *A* message *(or* potential*) is a continuously differentiable function $\phi : [\frac{1}{1+\lambda}, 1] \to \mathbb{R}$ with positive derivative. Hence ϕ is increasing and invertible on its range, and ϕ^{-1} is also continuously differentiable with positive derivative.*

The choice of an appropriate message is problem-dependent, and not very well understood. For this application to the hard-core model, the following message works:

$$\phi(x) := \frac{1}{s} \log \frac{x}{s-x}, \quad \text{where } s := \frac{d+1}{d}.$$

Note: A plausible derivation of this message is as a modification of the simpler message $\phi(x) = \log \frac{x}{1-x}$ that arises naturally in the analysis of the ferromagnetic Ising model with zero field.

Abbreviating p_{v_i} to p_i and writing $m_i = \phi(p_i)$, we can turn the recurrence (20.3) into a recurrence on messages, as follows:

$$m = \phi\Big(f[\phi^{-1}(m_1)], \dots, f[\phi^{-1}(m_d)]\Big) =: F(m_1, \dots, m_d).$$

We will write $\mathbf{m}$ to denote the vector of values $(m_1, \dots, m_d)$.

Claim 20.6. $\exists \gamma > 0$ *s.t.* $\forall \mathbf{m}, \mathbf{m}' \in (\phi[\frac{1}{1+\lambda}, 1])^d$,

$$|F(\mathbf{m}) - F(\mathbf{m}')| \le (1-\gamma)\|\mathbf{m} - \mathbf{m}'\|_\infty$$

First we will show that proving Claim 20.6 is sufficient to prove Theorem 20.2. Then we will go back and prove the Claim.

Lemma 20.7. *Claim 20.6 $\implies$ Theorem 20.2*

Proof. Assuming Claim 20.6 is true, we have for any l,

$$|m_r - m_r'| \le (1-\gamma)^l \max_{\text{leaves } L_i \text{ at depth } l} |m_{L_i} - m_{L_i}'|. \tag{20.5}$$

We just have to translate this to a similar statement with m replaced by p. Since ϕ is monotone,

$$\begin{aligned} \max |m_{L_i} - m_{L_i}'| &\le \phi(1) - \phi\Big(\frac{1}{1+\lambda}\Big) \\ &= \frac{1}{s}\Big[\log \frac{1}{s-1} - \log \frac{(1+\lambda)^{-1}}{s - (1+\lambda)^{-1}}\Big] \\ &= \frac{1}{s} \log \Big(\frac{s + s\lambda - 1}{s-1}\Big) \\ &\le c \log d, \end{aligned} \tag{20.6}$$

for a constant c, since $\lambda \le 4$ and $s = \frac{d+1}{d}$. Next, by the mean value theorem applied to ϕ:

$$|m_v - m_v'| = |\phi(p_v) - \phi(p_v')| \ge |p_v - p_v'| \inf_{p \in [\frac{1}{1+\lambda}, 1]} \phi'(p). \tag{20.7}$$

But it is simple to compute

$$\phi'(x) = \frac{1}{x(s-x)} \ge \frac{4}{s^2} \ge \frac{16}{9},$$

using the fact that $s = \frac{d+1}{d} \le \frac{3}{2}$. Hence (20.7) implies that $|p_v - p'_v| \le |m_v - m'_v|$, and therefore by (20.5) and (20.6) we get

$$|p_r - p'_r| \le |m_r - m'_r| \le (1-\gamma)^l \cdot c \log d \le \exp(-c'l),$$

as required. □

All that remains to prove Theorem 20.2 is to prove Claim 20.6.

Proof of Claim 20.6. By the multivariate mean value theorem:

$$|F(\mathbf{m}) - F(\mathbf{m}')| \le \sup_{\mathbf{m}} \|\nabla F(m_1, \ldots, m_d)\|_1 \cdot \|\mathbf{m} - \mathbf{m}'\|_\infty$$

where $\nabla F = \left(\frac{\partial F}{\partial m_i}\right)_i$. So it is sufficient to prove that $\sup_{\mathbf{m}} \|\nabla F(\mathbf{m})\|_1 \le 1 - \gamma$, where

$$\|\nabla F(\mathbf{m})\|_1 = \sum_{i=1}^{d} \left|\frac{\partial F}{\partial m_i}\right|.$$

Recall that $F = \phi \circ f \circ \phi^{-1}$, where

$$\phi(x) = \frac{1}{s} \log \frac{x}{s-x}, \quad \phi^{-1}(y) = \frac{y e^{sy}}{e^{sy}+1}, \quad f(x_1, \ldots, x_d) = \frac{1}{1 + \lambda \prod_i x_i}.$$

Now by direct calculation using the chain rule, we see **[exercise!]** that:

$$\frac{\partial F}{\partial m_i} = \frac{1-p}{s-p}(s - p_i)$$

where $p_i = \phi^{-1}(m_i)$ and $p := 1/(1 + \lambda \prod_i p_i)$.

Thus

$$\begin{aligned}
\|\nabla F(\mathbf{m})\|_1 &= \sum_{i=1}^{d} \frac{1-p}{s-p}(s - p_i) \\
&\le \frac{1-p}{s-p} \cdot d \cdot \left[s - \Big(\prod_{i=1}^{d} p_i\Big)^{1/d}\right] \\
&= \frac{1-p}{s-p} \cdot d \cdot \left[s - \Big(\frac{1-p}{\lambda p}\Big)^{1/d}\right]. \qquad (20.8)
\end{aligned}$$

The second line follows from the AM-GM inequality applied to the p_i

We want an upper bound for the expression in (20.8), which is provided by the following fact.

Proposition 20.8. *The function*

$$h(x) = \frac{(1-x)\left[1 + \frac{1}{d} - \left(\frac{1-x}{\lambda x}\right)^{1/d}\right]}{1 + \frac{1}{d} - x}$$

satisfies $\max_{x \in [0,1]} h(x) \le \frac{w}{1+w}$ *where w is the unique solution to* $w(1+w)^d = \lambda$.

Note that (20.8) is precisely $d \cdot h(p)$.

Proof. Write

$$h(x) = \frac{(1-x)(1+\frac{1}{d}-\Phi(x))}{1+\frac{1}{d}-x},$$

where $\Phi(x) := (\frac{1-x}{\lambda x})^{1/d}$. Now note that $\Phi'(x) = \frac{-\Phi(x)}{dx(1-x)}$, so Φ is monotonically decreasing on $[0,1]$, going from $+\infty$ to 0 in this interval. Thus it has a unique fixed point, which one can check **[exercise!]** is $x^* = \frac{1}{1+w}$ with w defined as above. Moreover, $\Phi(x) > x$ iff $x < x^*$. But now we can compute the derivative of h as

$$h'(x) = \frac{(1+\frac{1}{d})(\Phi(x)-x)}{dx(1+\frac{1}{d}-x)^2} \qquad \begin{cases} > 0 & \text{for } x < x^*; \\ < 0 & \text{for } x > x^*. \end{cases}$$

Hence $h(x)$ is maximized at x^*, and $h(x^*) = \frac{(1-x^*)(1+\frac{1}{d}-\Phi(x^*))}{1+\frac{1}{d}-x^*} = 1 - x^* = \frac{w}{1+w}$, as required. □

Now, if we set $\lambda = \lambda_c(d) = \frac{d^d}{(d-1)^{d+1}}$, then it's easy to check **[exercise!]** that $w = \frac{1}{d-1}$, and hence $h(x) \le \frac{w}{1+w} = \frac{1}{d}$, so $d \cdot h(x) \le 1$.

But one can also readily check **[exercise!]** that $w = w(\lambda)$ is monotonically strictly increasing with λ, and of course $\frac{w}{1+w}$ is increasing with w. Hence for any $\lambda < \lambda_c(d)$, there exists $\gamma > 0$ such that

$$\sup_{\mathbf{m}} \|\nabla F(\mathbf{m})\|_1 \le d \sup_x h(x) \le 1 - \gamma.$$

This completes the proof of Claim 20.6, and hence also of Theorem 20.2. □

20.3 Concluding remarks

1. An essentially equivalent version of the SAW tree construction (but not the correlation decay algorithm) appeared earlier in the context of matchings in the work of Godsil [God81].

2. Weitz's algorithm is deterministic, yielding an FPTAS for the partition function. It's interesting to note, however, that the SSM property that underlies the algorithm indirectly implies that MCMC also works for this problem, yielding an FPRAS when $\lambda < \lambda_c(\Delta)$, on all graphs of maximum degree Δ provided the graph is *amenable*, in the sense that the number of vertices in a ball of radius r around a vertex grows only polynomially with r. (This applies, e.g., to lattices $\mathbb{Z}^d$, but not to tree-like graphs.) This follows from a generic implication that, on amenable graphs, SSM implies $O(n \log n)$ mixing time of the Glauber dynamics [DSVW04]. Note that, for the hard-core model, the SSM property we have proved for trees applies also to general graphs of maximum degree Δ via the SAW tree construction.

3. For some special families of graphs, one may be able to do better than for arbitrary graphs of maximum degree Δ by exploiting additional structure of the SAW tree, notably the fact that its "average degree" (suitably defined) may be significantly less than Δ. An important example is the hard-core model on the Cartesian lattice $\mathbb{Z}^d$. Here one can show, e.g., that the Weitz algorithm works on (arbitrarily large square regions of) $\mathbb{Z}^2$ for $\lambda < 2.08$, which is much larger than $\lambda_c(4) = 1.69$, with analogous (though decreasing) improvements for larger dimensions d [SSŠY17]. It is notable that these lower bounds on the uniqueness threshold[1], which are currently the best known, have come from algorithmic investigations in Computer Science rather than from Physics arguments.

[1]Strictly speaking, we should insert the caveat that a single uniqueness threshold is not known to exist for the hard core model on $\mathbb{Z}^d$; this is because the property of uniqueness is not known to be monotone in λ (though it is conjectured to be so)—so it's conceivable that there are multiple thresholds.

4. An analogous algorithmic result holds also for the *antiferromagnetic Ising model* [SST14] (and indeed for any antiferromagnetic spin system with two spin values; "antiferromagnetic" here means roughly that the potential function gives higher weight to neighboring spins that disagree than agree—this is true for the hard-core model due to the hard constraint forbidding adjacent +1 spins). Recall that the antiferromagnetic Ising model on graph $G = (V, E)$ has partition function

$$Z_G(\lambda, \mu) = \sum_{S \subseteq V} \lambda^{|E(S,\bar{S})|} \mu^{|S|},$$

where $\lambda = \exp(2\beta) \geq 1$ (nearest-neighbor interaction) and $\mu = \exp(\beta h)$ (external field). The uniqueness diagram for the Ising model on the Δ-regular tree is sketched in Figure 20.1. (This picture can be verified via the same kind of recurrence analysis we used for the hard-core model.) Note in particular

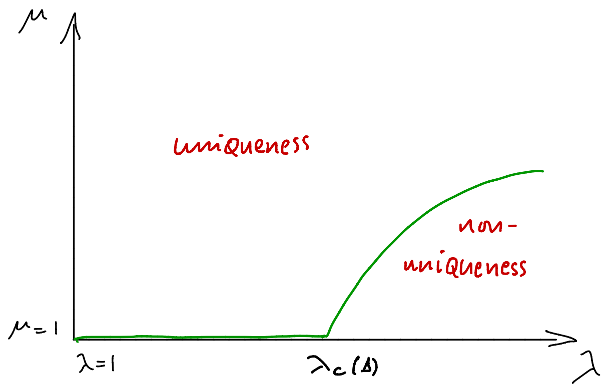

Figure 20.1: Phase diagram for the anti-ferromagnetic Ising model on the Δ-regular tree.

that, for $1 \leq \lambda \leq \lambda_c(\Delta)$, where $\lambda_c(\Delta) = \frac{\Delta}{\Delta-2}$ is the threshold for the model with zero field ($\mu = 1$), we have a unique Gibbs measure for all external fields; and for $\lambda > \lambda_c(\Delta)$ there exists a critical field $\mu_c(\lambda, \Delta) > 0$ such that uniqueness occurs for $|\log \mu| \geq \log \mu_c(\lambda, \Delta)$. I.e., as the interaction parameter λ gets stronger, we need a larger field to break the separation of the system into + and − phases. Within the uniqueness region, by definition we get WSM on the tree, and the same arguments as those above yield an FPTAS. Specifically:

- Weitz's SAW tree construction still works in exactly the same way, since it depends only on the fact that there are two spin values.
- An analogous stepwise decay to that in the prove of Theorem 20.2 can be used to prove that WSM implies SSM, though the details are different: namely, the recurrence is different and one needs a different message function.

And as in the case of the hard-core model, there are complementary negative results showing that no FPRAS can exist outside the uniqueness region unless NP=RP [SS12, GGŠ$^+$14]. (We'll be seeing a proof of this for the hard-core model in the next lecture.)

5. The situation for *ferromagnetic* 2-spin systems is much murkier than the clean picture we have seen for the antiferromagnetic case. First, as we saw in an earlier lecture, there is an FPRAS (based on MCMC) for *any* graph, at *all* values of the parameters [JS93]. Hence we can't expect any negative complexity theoretic results in this case. On the other hand, it is still an open problem to find a *deterministic* approximation algorithm (based on correlation decay or any other technique) throughout this range. There is still a uniqueness threshold on the Δ-regular tree here (exactly analogous to the antiferromagnetic case—indeed the same phase diagram as in Figure 20.1 holds here as the tree is

bipartite), so we might expect Weitz's algorithm to work at least in this regime. But even this is not known, since while the SAW tree construction works (as it does for any 2-spin system), as noted earlier the implication that WSM implies SSM no longer holds in the ferromagnetic case. On the other hand, if we don't restrict to graphs of bounded degree, then correlation decay algorithms are known essentially up to the uniqueness threshold for all ferromagnetic 2-spin systems [LLY13, GL18]. And this positive result is complemented by a #BIS-hardness result for these models on the other side of this threshold (excluding the special case of the ferromagnetic Ising model, which as we've seen is tractable everywhere) [LLZ14].

6. It is an open problem to extend Weitz's SAW tree framework to spin systems with more than two spins (such as colorings or the Potts model). Some attempts in this direction (which fall quite far short of what can be done with other methods, such as MCMC) can be found in [GK12, GKM15].

References

[DSVW04] M. Dyer, A. Sinclair, E. Vigoda, and D. Weitz. Mixing in time and space for lattice spin systems. *Random Structures & Algorithms*, 24:461–479, 2004.

[GGŠ+14] A. Galanis, Q. Ge, D. Štefankovič, E. Vigoda, and Y. Yin. Improved inapproximability results for counting independent sets in the hard-core model. *Random Structures & Algorithms*, 45:78–110, 2014.

[GK12] D. Gamarnik and D. Katz. Correlation decay and deterministic fptas for counting colorings of a graph. *Journal of Discrete Algorithms*, pages 29–47, 2012.

[GKM15] D. Gamarnik, D. Katz, and S. Misra. Strong spatial mixing of list coloring of graphs. *Random Structures & Algorithms*, pages 599–613, 2015.

[GL18] H. Guo and P. Lu. Uniqueness, spatial mixing and approximation for ferromagnetic 2-spin systems. *ACM Transactions on Computing Theory*, 10:17, 2018.

[God81] C.D. Godsil. Matchings and walks in graphs. *Journal of Graph Theory*, 5:285–297, 1981.

[JS93] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22:1087–1116, 1993.

[LLY13] L. Li, P. Lu, and Y. Yin. Correlation decay up to uniqueness in spin systems. *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 67–84, 2013.

[LLZ14] J. Liu, P. Lu, and C. Zhang. The complexity of ferromagnetic two-spin systems with external fields. *Proceedings of RANDOM*, pages 843–856, 2014.

[RST+13] R. Restrepo, J. Shin, P. Tetali, E. Vigoda, and L. Yang. Improved mixing condition on the grid for counting and sampling independent sets. *Probability Theory and Related Fields*, 156:75–99, 2013.

[SS12] A. Sly and N. Sun. The computational hardness of counting in two-spin models on d-regular graphs. *Prodeedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 361–369, 2012.

[SSŠY17] A. Sinclair, P. Srivastava, D. Štefankovič, and Y. Yin. Spatial mixing and the connective constant: Optimal bounds. *Probability Theory and Related Fields*, 168:153–197, 2017.

[SST14] A. Sinclair, P. Srivastava, and M. Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *Journal of Statistical Physics*, 155:666–686, 2014.