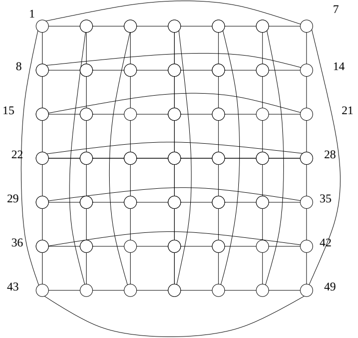
1
7
8
14
15
21
22
28
29
35
36
42
43
49

Problem 6.4

Gibbs sampling and mean field: Consider the Ising model with binary variables $X_s \in \{0, 1\}$, and a factorization of the form:

$$p(x; \eta) \quad \propto \quad \Big\{ \sum_{s \in V} \eta_s x_s + \sum_{(s,t) \in E} \eta_{st} x_s x_t \Big\}.$$

To make the problem symmetric, assume a 2-D grid with toroidal (donut-like) boundary conditions, as illustrated in Figure 1(c).

(a) Derive the Gibbs sampling updates for this model. Implement the algorithm for $\eta_{st} = 0.5$ for all edges, and $\eta_s = (-1)^s$ for all $s \in \{1, \ldots, 49\}$ (using the node ordering in Figure 1(c)). Run a burn-in period of 1000 iterations (where one iteration amounts to updating each node once). For each of 5000 subsequent iterations, collect a sample vector, and use the 5000 samples to form Monte Carlo estimates $\widehat{\mu}_s$ of the moments $\mathbb{E}[X_s]$ at each node. Output a 7×7 matrix of the estimated moments. Repeating this same experiment a few times will provide an idea of the variability in your estimate.

(b) Derive the naive mean field updates (based on a fully factorized approximation), and implement them for the same model. Compute the average ℓ_1 distance $\frac{1}{49} \sum_{i=1}^{49} |\tau_s - \widehat{\mu}_s|$ between the mean field estimated moments τ_s, and the Gibbs estimates $\widehat{\mu}_s$.

Problem 6.5

Bonus: Sum-product and loopy graphs:

In many real-world applications (e.g., error-correcting codes, data compression, sensor networks, computer vision, bioinformatics etc.), the sum-product algorithm is applied to graphs with cycles. In sharp contrast to trees, the sum-product updates are no longer guaranteed to converge, or to compute the correct marginal distributions. Indeed, the results can be very surprising! As an illustration of this phenomenon, consider a model on a toroidal grid (Figure 1(c)), given by $p(z; \theta) \propto \exp\big\{\theta \sum_{(s,t) \in E} z_s z_t\big\}$ where $Z_s \in \{-1, +1\}$ and $\theta \in \mathbb{R}$. Since this model is symmetric, the correct marginal distributions are $p(z_s = 1; \theta) = 0.5$ for all nodes, and for all coupling strengths $\theta \in \mathbb{R}$. However, if the loopy sum-product algorithm is run on this problem, it exhibits the strange behavior shown in Figure 2. Show that for all $\theta < \theta_{crit} = \frac{1}{2} \log 2 \approx 0.3466$, the loopy form of sum-product will compute the correct symmetric marginal distributions $p(z_s = 1; \theta) = 0.5$. Explain why it breaks down for $\theta > \theta_{crit}$.

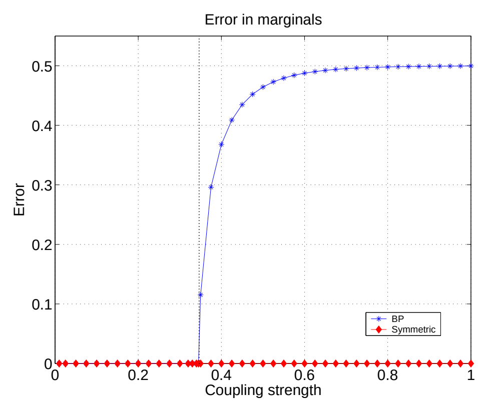

Figure 2. Break-down of the sum-product algorithm on the homogeneous Ising model. For all $\theta < \theta_{crit} \approx 0.3466$, the sum-product algorithm computes the correct symmetric marginals. Beyond this point, it outputs increasingly inaccurate answers, as shown by the large average ℓ_1 distance between the SP estimate and symmetric truth.