Approximating the Tutte Polynomial (and the Potts partition function)

Leslie Ann Goldberg, University of Liverpool

(based on joint work with Mark Jerrum)

Counting, Inference and Optimization on Graphs, Princeton

November 2-5, 2011

The Tutte polynomial of a graph $G = (V, E)$

$$T(G; x, y) = \sum_{A \subseteq E} (x-1)^{\kappa(V,A)-\kappa(V,E)} (y-1)^{|A|-(|V|-\kappa(V,A))}$$

$\kappa(V, A) =$ number of connected components of the graph (V, A)

$$T(G;x,y) = \sum_{A\subseteq E} (x-1)^{\kappa(V,A)-\kappa(V,E)}(y-1)^{|A|-(|V|-\kappa(V,A))}$$

If G is connected, $T(G;1,1)$ counts spanning trees.

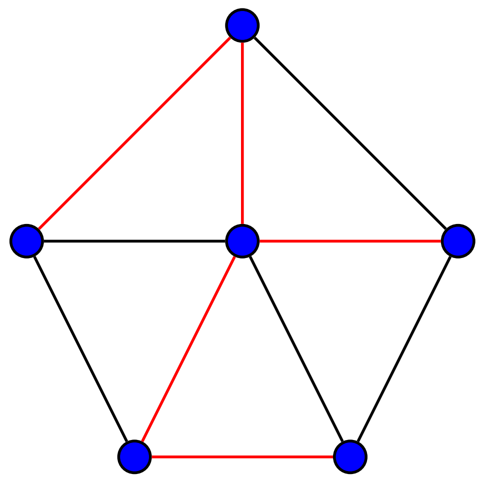

$$T(G;x,y)=\sum_{A\subseteq E}(x-1)^{\kappa(V,A)-\kappa(V,E)}(y-1)^{|A|-(|V|-\kappa(V,A))}$$

If G is connected, $T(G;2,1)$ counts forests.

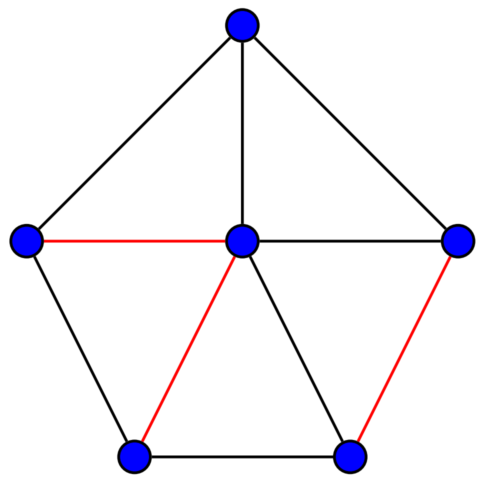

Combinatorial interpretation of the Tutte polynomial

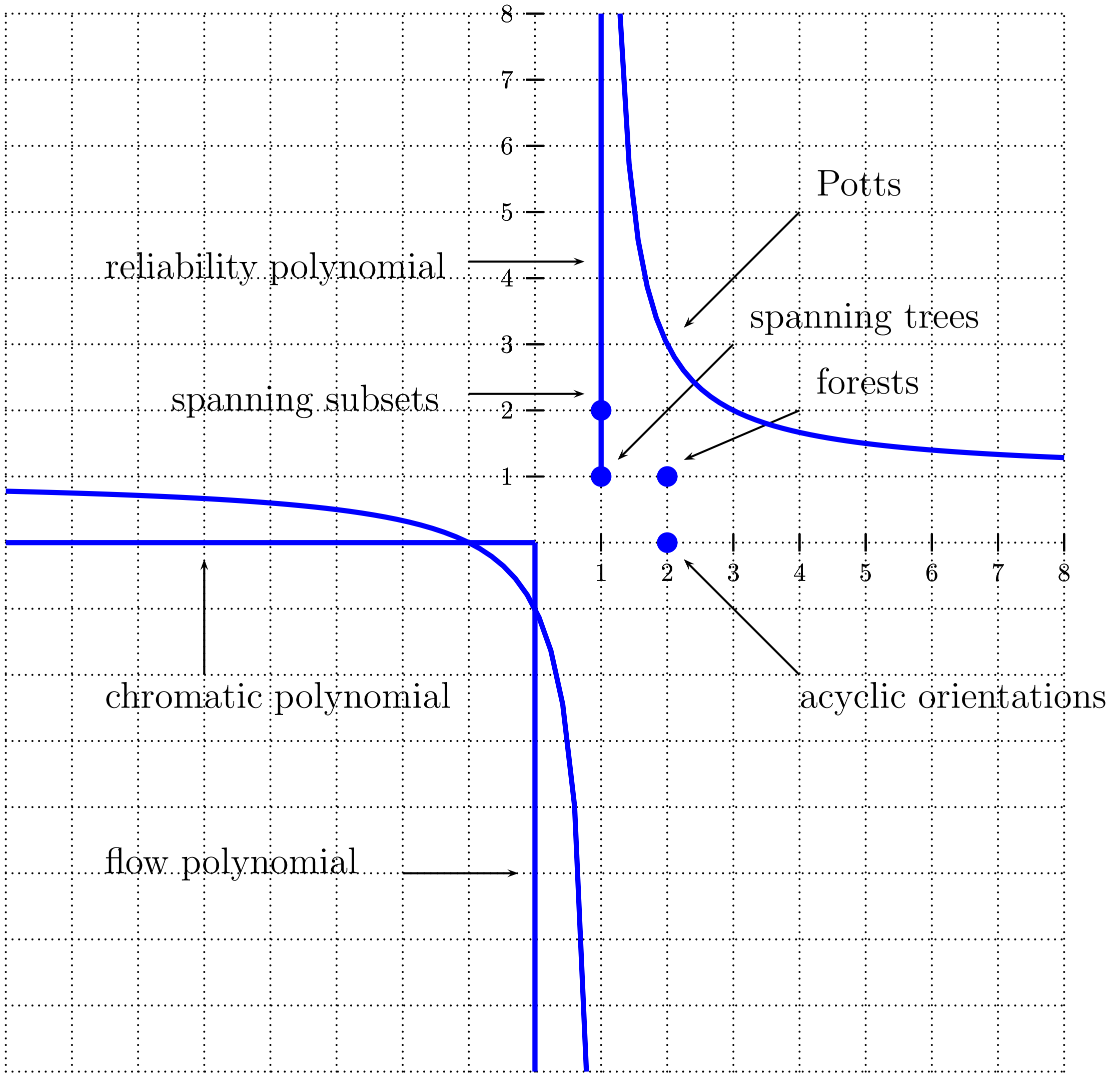

Partition function of the q-state Potts model at $(x-1)(y-1) = q$

Complexity of evaluating the Tutte polynomial

For fixed rationals x and y, Jaeger, Vertigan and Welsh (1990) studied the complexity of the following problem.

Name. $\textsc{Tutte}(x, y)$.
Input. A graph $G = (V, E)$.
Output. $T(G; x, y)$.

They showed that for all (x, y), $\textsc{Tutte}(x, y)$ is either #P-hard or computable in polynomial time.

$$T(G; x, y) = \sum_{A \subseteq E} (x-1)^{\kappa(V,A)-\kappa(V,E)} (y-1)^{|A|-|V|+\kappa(V,A)}$$

Complexity of evaluating the Tutte polynomial

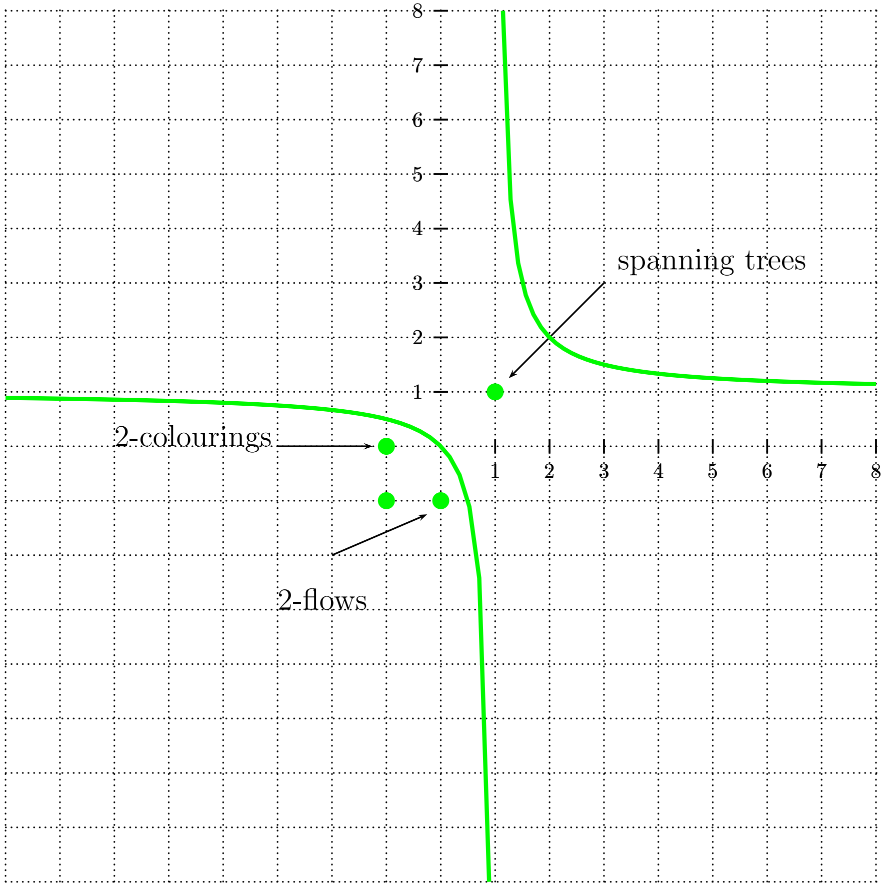

$(x-1)(y-1) = 1.$

Approximate evaluation

Reminder from Mark's talk:

Definition. An FPRAS is a randomised algorithm that produces a result that is correct to within relative error $1 \pm \varepsilon$ with high probability. It must run in time $\mathrm{poly}(n, \varepsilon^{-1})$, where n is the input size.

Approximate evaluation

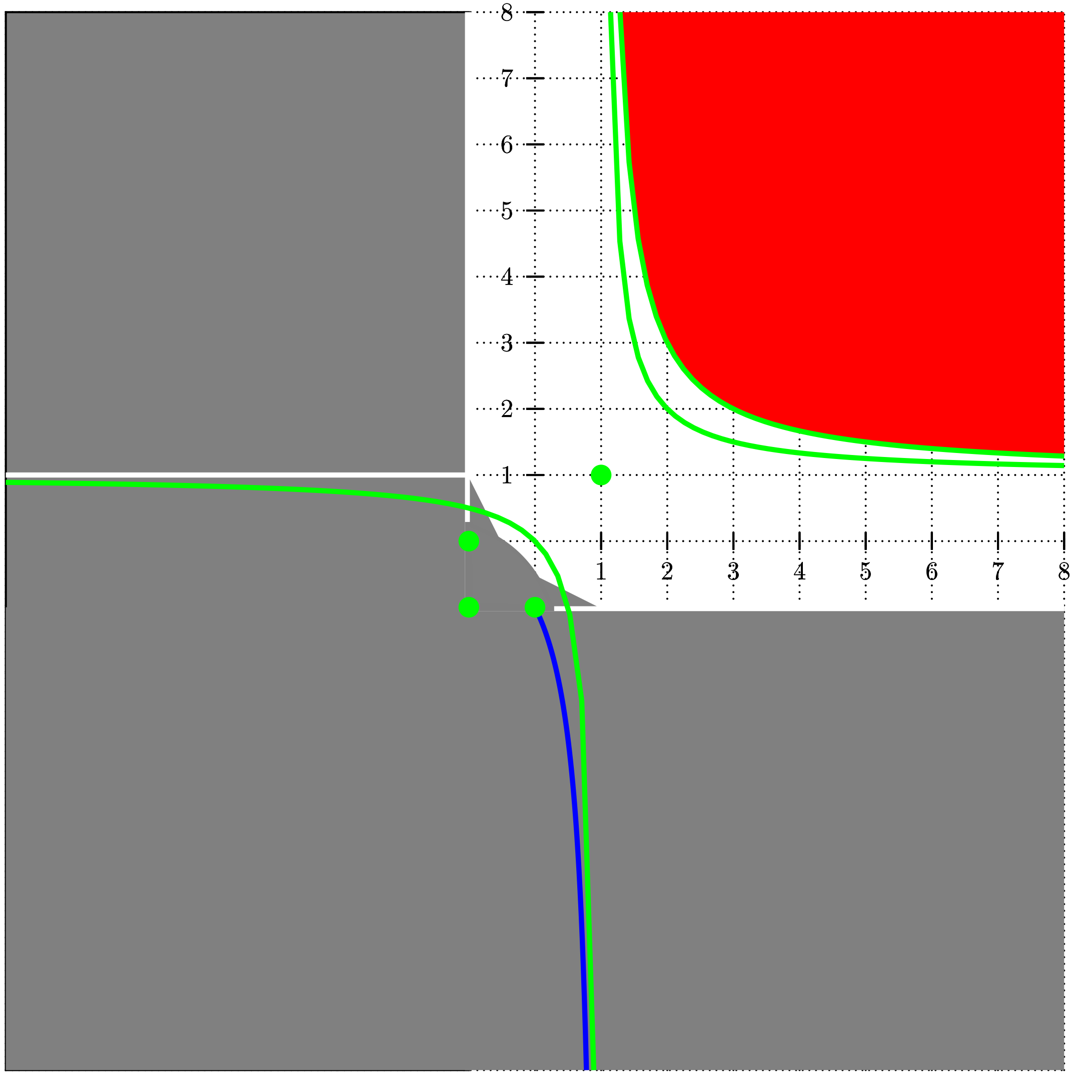

FPRASable points (green)

- Points where exact evaluation is possible in polynomial time
- Points on the upper branch of the hyperbola $(x-1)(y-1)=2$. This is due to Jerrum and Sinclair's FPRAS (1993) for the partition function of the Ising model in the ferromagnetic case. (Mark already talked about this) The Ising model is the 2-state Potts model.

The random cluster formulation of the Tutte polynomial

The multivariate Tutte polynomial of G is

$$Z_{\mathrm{Tutte}}(G; q, \gamma) = \sum_{F \subseteq E} q^{\kappa(V,F)} \prod_{e \in F} \gamma_e,$$

where q and $\gamma = \{\gamma_e\}_{e \in E}$ are commuting indeterminates.

If $(x-1)(y-1) = q$ and $\gamma_e = y - 1$ for all e then

$$T(G; x, y) = q^{-\kappa(V,E)} \gamma^{-|V|+\kappa(V,E)} Z_{\mathrm{Tutte}}(G; q, \gamma).$$

The partition function of the q-state Potts model

If q is positive integer then (Fortuin and Kastelyn) $Z_{\text{Tutte}}(G; q, \gamma)$ is equal to

$$Z_{\text{Potts}}(G; q, \gamma) = \sum_{\sigma : V \to [q]} \prod_{e \in E} \big(1 + \gamma_e \delta_e(\sigma)\big),$$

where $[q] = \{1, \ldots, q\}$ is a set of q spins or colours, and $\delta_e(\sigma)$ is 1 if e is monochromatic in σ and 0 otherwise.

The Ising model is the case $q = 2$. Jerrum and Sinclair's FPRAS is for $\gamma_e = \gamma > 0$ (the ferromagnetic case).

Non-FPRASable points

At every grey point, we (2008) showed that there is no FPRAS unless RP=NP.

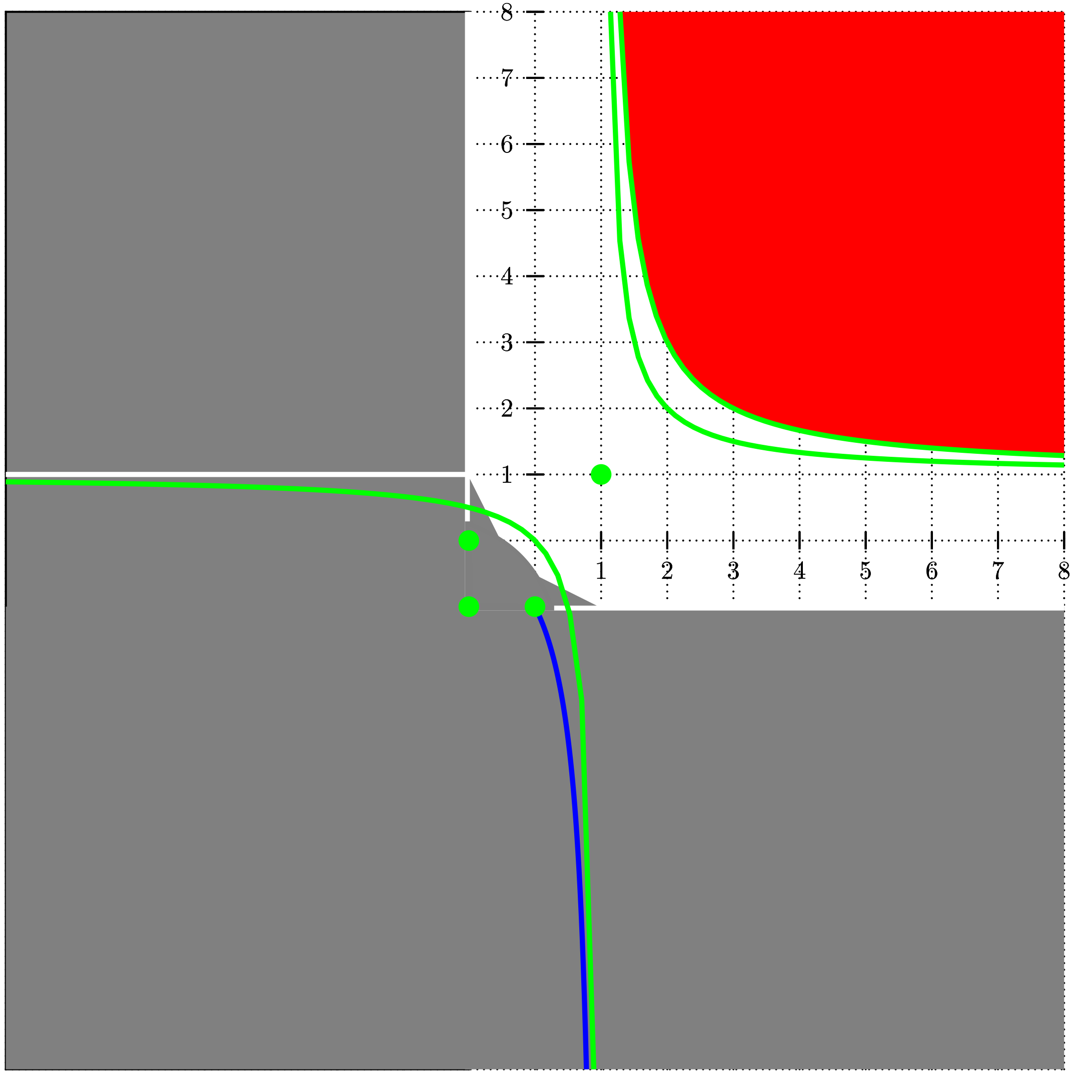

- $x < -1$ except $q = 0, 1$
- $y < -1$ except $q = 1, 2$
- In the vicinity of the origin in the triangle $y < -1 - 2x$
- In the vicinity of the origin in the triangle $x < -1 - 2y$
- In the vicinity of the origin and $q > 1.5$

This is straightforward for some points, for example, $T(G; -2, 0)$ is the number of proper 3-colourings of G (so the decision problem is NP-hard). On the other hand, $T(G; 0, -5)$ is the number of nowhere-zero 6-flows, and Seymour has shown that there is a nowhere-zero 6-flow iff G has no cut edge, so the decision problem is in P.

Implementing edge weights using series and parallel compositions (see Brylawski, JVW, Sokal)

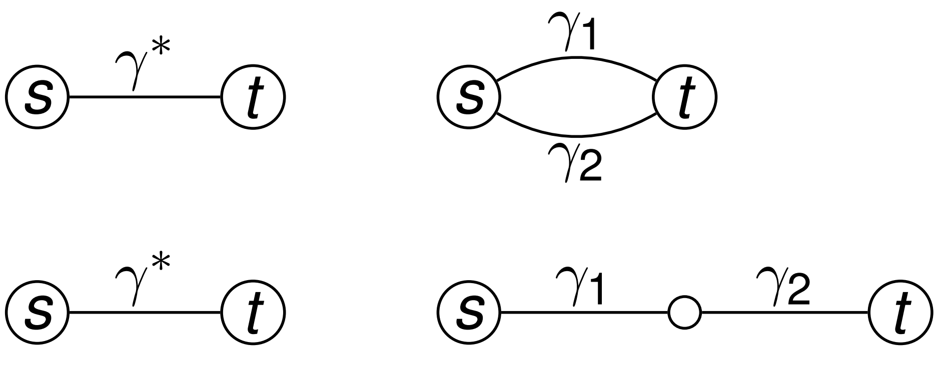

$$1 + \gamma^* = (1 + \gamma_1)(1 + \gamma_2)$$

$$1 + \frac{q}{\gamma^*} = \left(1 + \frac{q}{\gamma_1}\right)\left(1 + \frac{q}{\gamma_2}\right)$$

Our key tool (for $q \notin \{0, 1, 2\}$): If copies of γ can be used to implement some $\gamma^* \notin [-2, 0]$ ($y \notin [-1, 1]$) and also some $\gamma^* \in (-2, 0)$ then there is no FPRAS for evaluating $Z_{\text{Tutte}}(G; q, \gamma)$ (where γ is the constant function mapping every edge to weight γ).

What about the other points?

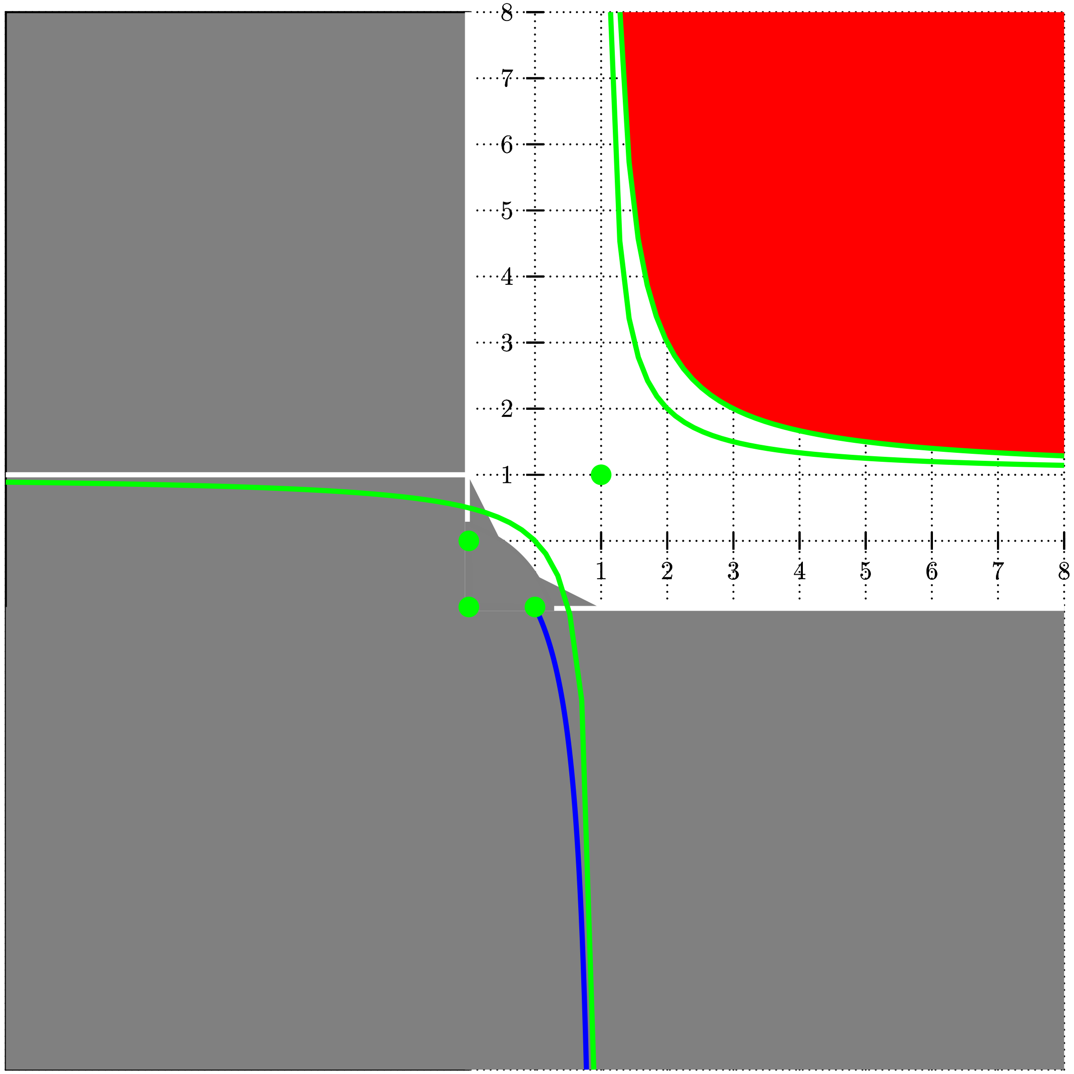

- On the blue hyperbola segment ($q = 2$), approximate evaluation is equivalent in difficulty to approximately counting perfect matchings in a graph.
- Alon, Frieze and Welsh (1995) gave an FPRAS for the region $x \geq 1$, $y \geq 1$ when G is "dense" (minimum degree $\Omega(n)$).
- Recently (2010), we gave a negative result for the red region, subject to the hardness of the complexity class $\#\mathrm{RH}\Pi_1$, which Mark told you about.

Reminders from Mark's talk

The problems in $\#\mathrm{RH}\Pi_1$ can be expressed in terms of counting the number of models of a logical formula from a certain syntactically restricted class.

Complete for $\#\mathrm{RH}\Pi_1$ wrt Approximation-preserving (AP)-reductions:

Problem

Name: #BIS.
Instance: A bipartite graph B.
Output: The number of independent sets in B.

Problem

Name: $\textsc{Tutte}(q, \gamma)$.
Instance: A graph G.
Output: $Z_{\text{Tutte}}(G; q, \gamma)$, where γ is the constant function with $\gamma_e = \gamma$ for all e.

The result in the red region: $\textsc{Tutte}(q, \gamma)$ is hard for $\#\mathrm{RH}\Pi_1$ with respect to Approximation-Preserving (AP)-reductions when $q > 2$ and $\gamma > 0$.

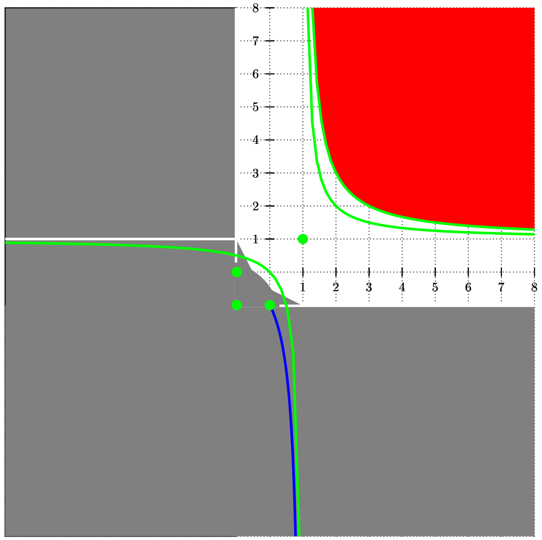

The Tutte polynomial of a hypergraph

The (multivariate) Tutte polynomial of $H = (\mathcal{V}, \mathcal{E})$ is defined as

$$Z_{\text{Tutte}}(H; q, \gamma) = \sum_{\mathcal{F} \subseteq \mathcal{E}} q^{\kappa(\mathcal{V}, \mathcal{F})} \prod_{f \in \mathcal{F}} \gamma_f ,$$

where $\kappa(\mathcal{V}, \mathcal{F})$ denotes the number of connected components in the sub-hypergraph $(\mathcal{V}, \mathcal{F})$.

From #BIS to hypergraph Tutte

Problem

Name: UNIFORMHYPERTUTTE(q, γ).
Instance: A uniform hypergraph $H = (\mathcal{V}, \mathcal{E})$.
Output: $Z_{\mathrm{Tutte}}(H; q, \gamma)$, where γ is the constant function with $\gamma_f = \gamma$ for all f.

Using standard techniques:

Lemma. *Suppose $q > 1$. Then*

$$\#\mathrm{BIS} \leq_{\mathrm{AP}} \textsc{UniformHyperTutte}(q, q-1).$$

From hypergraph Tutte to graph Tutte

Lemma. *Suppose that $q > 2$ and $\gamma, \gamma' > 0$. Then* $\textsc{UniformHyperTutte}(q, \gamma) \leq_{\mathrm{AP}} \textsc{Tutte}(q, \gamma')$.

We go via the intermediate problem of approximating $Z_{\mathrm{Tutte}}(G, q, \gamma)$ when $\gamma = \{\gamma_e\}$ contains two different edge weights (which are part of the problem instance).

The Random cluster distribution

For a graph $G = (V, E)$, associate every edge $e \in E$ with a quantity $p(e) \in [0, 1]$. For a set of edges $A \subseteq E$ define

$$\widetilde{P}(G; A, q, p) = q^{\kappa(V,A)} \prod_{e \in A} p(e) \prod_{e \in E \setminus A} (1 - p(e)).$$

Then the probability of edge-set A in the random cluster model is given by

$$P(G; A, q, p) = \widetilde{P}(G; A, q, p) / Z_{\mathrm{rc}}(G; q, p).$$

where $Z_{\mathrm{rc}}(G; q, p)$ is the appropriate normalising factor (essentially just an alternative parameterisation of the multivariate Tutte polynomial).

First-order phase transition

For the complete graph K_N and $q > 2$, the random cluster model has a first-order phase transition with edge probability $p = \lambda_c/N$ [Bollobás, Grimmett and Janson, 1996].

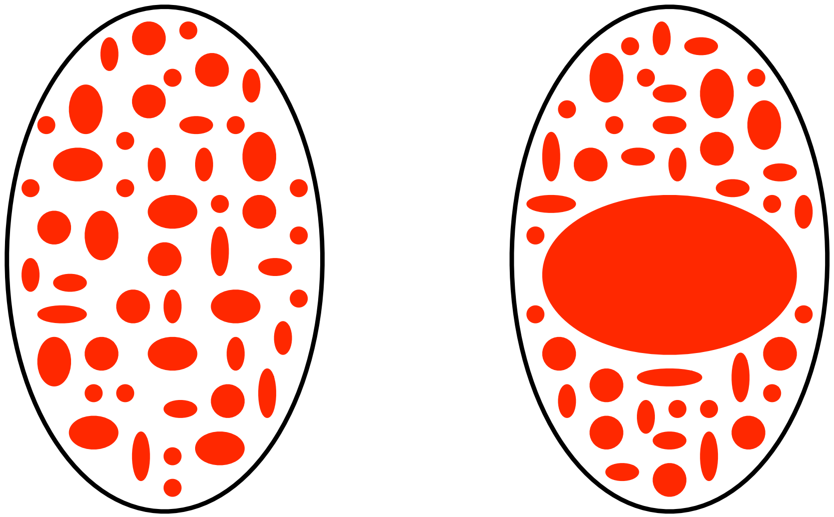

If p is carefully tuned, two phases coexist: one with all components $O(\log n)$; one with a giant component containing a constant fraction of the vertices.

Simulating a hyperedge

Gadget for simulating a hyperedge with $t = |T|$ vertices.

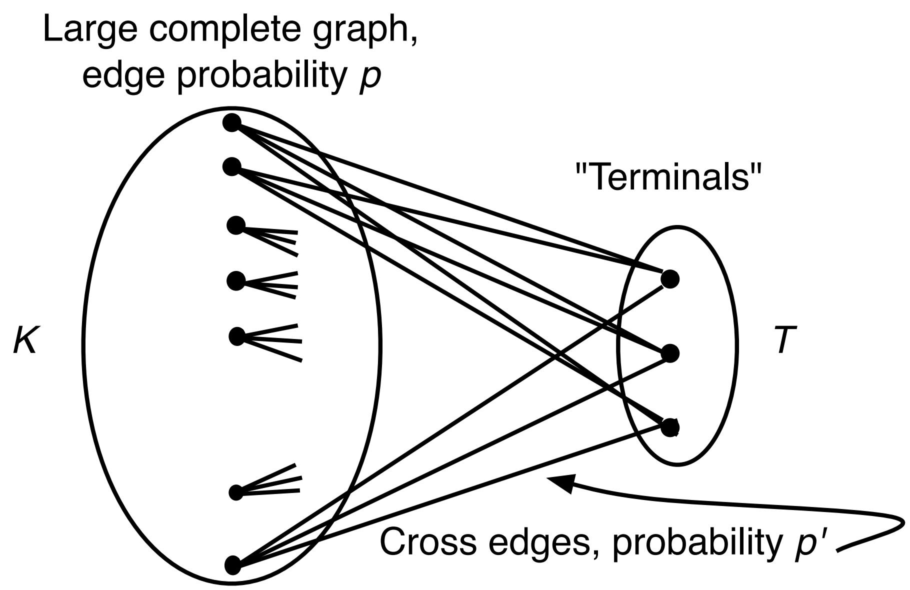

Salient features:

- $t \ll N$, where $N = |K|$,
- $p \approx \lambda_c / N$,
- $\lambda_c / N \ll p' \ll 1$.

Simulating a hyperedge (continued)

Gadget in “hyperedge excluded” state: With high probability all edges crossing from T hit a different connected component in K.

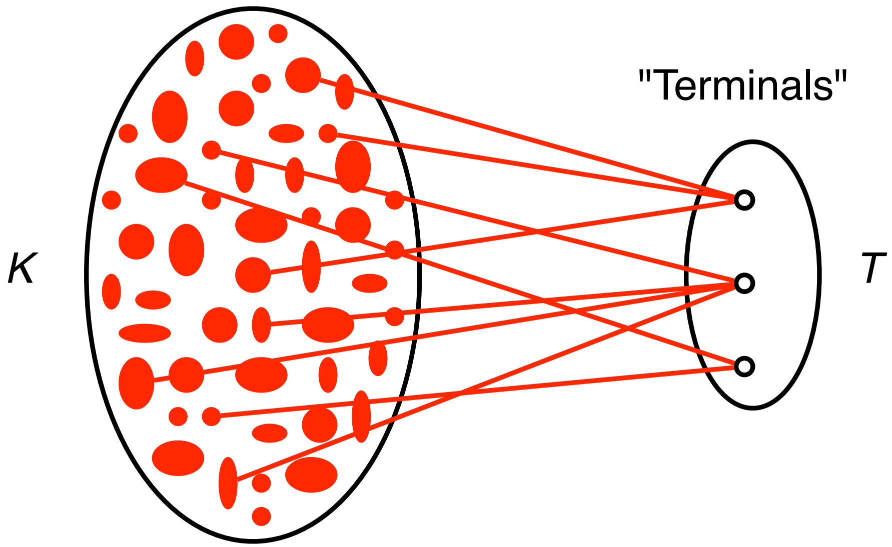

Effect: all terminals in T find themselves in different components.

Simulating a hyperedge (continued)

Gadget in “hyperedge included” state: With high probability at least one edge from each terminal hits the giant component in K.

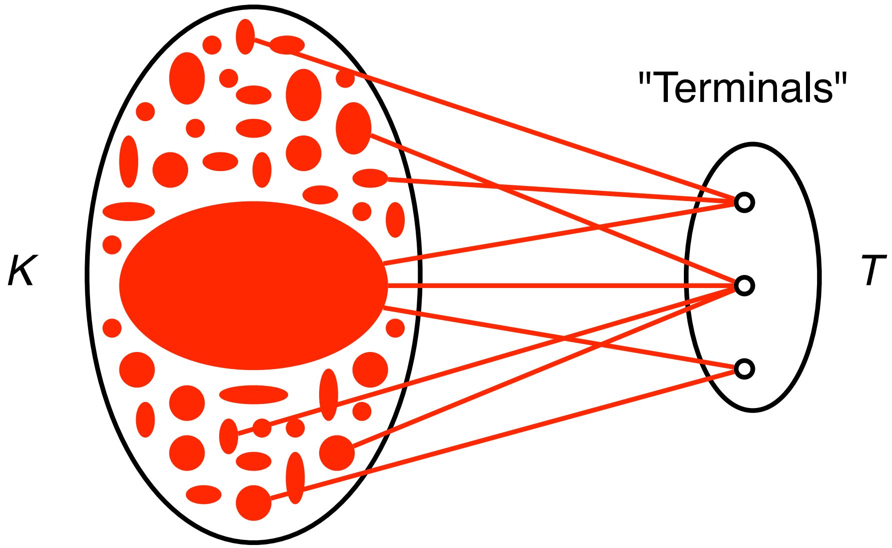

Effect: all terminals in T find themselves in the same components.

Getting #BIS hardness for the red region

So far we have reduced #BIS to a multivariate version of $Z_{\mathrm{Tutte}}(G; q, \gamma)$ in which q is fixed and $\gamma = \{\gamma_e\}$ contains just two values (not under our control).

We need to get from there to $\textsc{Tutte}(q, \gamma)$, where γ is a desired constant edge weight. This can be using implementations (series-parallel compositions) as before. Putting it all together:

Theorem. *Suppose $q > 2$ and $\gamma > 0$. Then* #BIS $\leq_{\mathrm{AP}}$ $\textsc{Tutte}(q, \gamma)$.

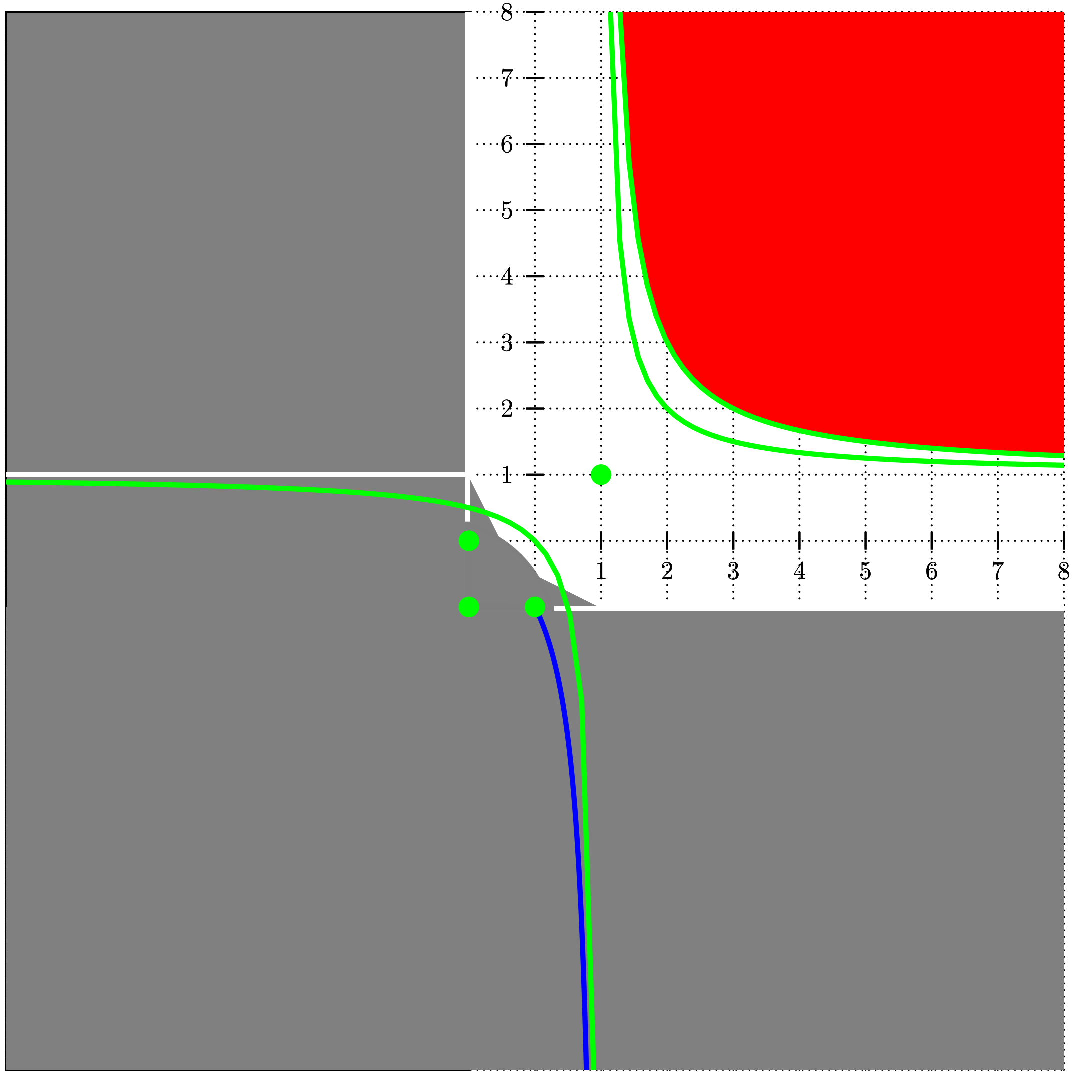
8
7
6
5
4
3
2
1
1 2 3 4 5 6 7 8

Does $\textsc{Tutte}(q, \gamma) \leq_{\mathrm{AP}}$ #BIS for $q > 2$ and $\gamma > 0$?

Bordewich 2010 showed that if any problem in #P fails to have an FPRAS, then there is an infinite approximation hierarchy within #P.