The Complexity of Approximate Counting
(How hard is it to evaluate, or approximately evaluate, generating functions?)

Leslie Ann Goldberg, Oxford

AofA 2024

35th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms

Bath, 17-21 June, 2024

Warm up: The independence polynomial

Graph $G = (V, E)$

A set $I \subseteq V$ is an independent set of G if it contains no edges of G.

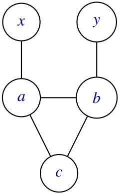

- $\emptyset$ is an independent set of G.
- Any size-1 subset of V is an independent set of G.
- G has 6 larger independent sets.

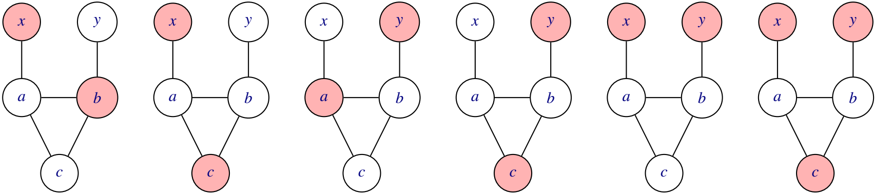

The independence polynomial

$$Z_G(\lambda) = \sum_{I \in \mathcal{I}_G} \lambda^{|I|}.$$

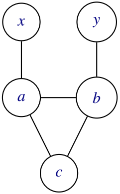

- $\emptyset$ is an independent set of G.
- Any size-1 subset of V is an independent set of G.
- G has 6 larger independent sets.

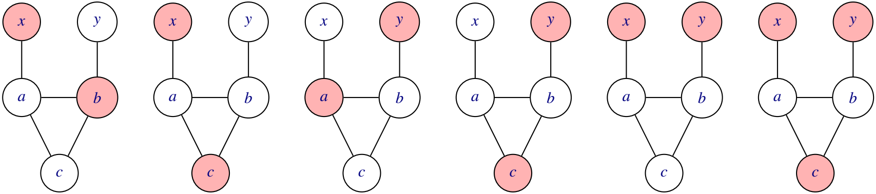

$$Z_G(\lambda) = 1 + 5\lambda + 5\lambda^2 + \lambda^3$$

Gibbs/Boltzmann distribution

Gibbs measure: measure on independent sets where the probability of I is $\propto \lambda^{|I|}$.

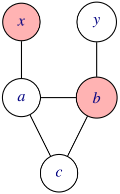

The probability of this independent set is $\lambda^2/Z_G(\lambda)$, where $Z_G(\lambda) = 1 + 5\lambda + 5\lambda^2 + \lambda^3$.

The normalising factor $Z_G(\lambda)$ (the independence polynomial of G) is also called the partition function of G in statistical physics.

How hard is it to compute the independence polynomial of G, given a fixed “activity” λ?

$$Z_G(\lambda) = \sum_{I \in \mathcal{I}_G} \lambda^{|I|}$$

Start with $\lambda = 1$...

$Z_G(1)$ is the number of independent sets of G. This is known to be #P-complete (Valiant, 1979), even when the graph is restricted to have degree at at most 3 (Greenhill, 2000).

But most people are more interested in approximating $Z_G(\lambda)$.

Approximating the Partition Function

Fix λ.

$$Z_G(\lambda)(1-\varepsilon) \leqslant \hat{Z} \leqslant Z_G(\lambda)(1+\varepsilon)$$

Running time at most $\mathrm{poly}(n, 1/\varepsilon)$.

FPRAS or FPTAS.

Robust.

When is there an FPTAS/FPRAS?

$\mathcal{G}_\Delta$: Graphs with max degree at most Δ.

$$\lambda_c = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^{\Delta}} \sim \frac{e}{\Delta}.$$

• Weitz 2006: If $\lambda < \lambda_c$ then $Z_G(\lambda)$ can be efficiently approximated on graphs $G \in \mathcal{G}_\Delta$. (There is an FPTAS)

• Sly 2010; Galanis, Štefankovič, Vigoda 2012; Sly, Sun 2012: If $\lambda > \lambda_c$ then, for some $\kappa > 1$, $Z_G(\lambda)$ cannot be efficiently approximated within a factor of κ^n on Δ-regular graphs (unless NP=RP).

What is the magic value λ_c?

Let $Z^{\text{in}}_{G,v}(\lambda) = \sum_{I \in \mathcal{I}_G;\, v \in I} \lambda^{|I|}$. In the Gibbs measure, the probability that v is occupied is

$$p_v(G) = \frac{Z^{\text{in}}_{G,v}(\lambda)}{Z_G(\lambda)}$$

Given a Δ-regular tree T of height h with root r,

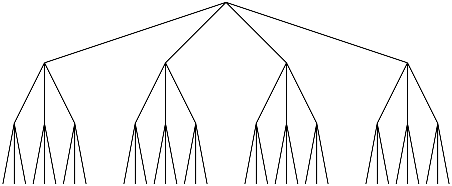

The occupation ratio of the tree: $p[h] = p_r(T)$.

$p[h]$ converges to a limit as $h \to \infty$ iff $\lambda \leqslant \lambda_c$ (Kelly 1985)

The complexity of approximating $Z_G(\lambda)$ for $G \in \mathcal{G}_\Delta$ depends on whether $p[h]$ converges.

Another Example: Partition function of the Ising model

- "Spins" $\{0, 1\}$

- parameter β (real number, associated with the temperature of the model)

Graph $G = (V, E)$ 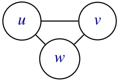

A "configuration" σ assigns a spin to every vertex 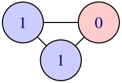

Spins interact along the edges: Same spins contribute a factor of β so $w(\sigma) = \beta$.

The partition function: $Z_G = \sum_{\sigma \in [q]^V} w(\sigma)$

Example: Partition function of the Ising model

$$Z_G = \sum_{\sigma \in [q]^V} w(\sigma)$$

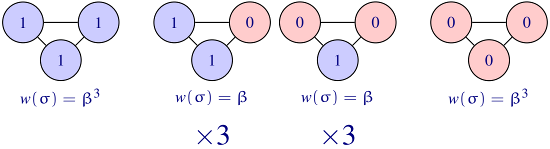

$Z_G = 2\beta^3 + 6\beta$.

Again can consider difficulty of approximately computing Z_G, or sampling from the Gibbs distribution.

Partition Functions, more generally

- Fix "Spins" $[q] = \{0, \ldots, q-1\}$ e.g., $q = 2$ for the Ising model

- Fix symmetric matrix $A \in \mathbb{R}^{q \times q}$

e.g., for the Ising model $A = \begin{pmatrix} \beta & 1 \\ 1 & \beta \end{pmatrix}$.

- Given a graph $G = (V, E)$

- configuration $\sigma \in [q]^V$ weight

$$w(\sigma) = \prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)}$$

- Partition function Z_G associates G with the real number

$$Z_G = \sum_{\sigma \in [q]^V} w(\sigma)$$

Examples: Ising model, Potts model... Independent sets

$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$

Example
Partition functions count graph homomorphisms

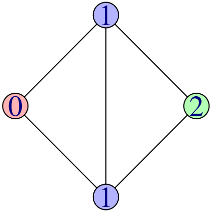

Graph $G = (V, E)$

Graph H

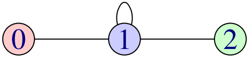

Homomorphism from G to H: A map $\sigma \in V(H)^V$ that maps every edge of G to an edge of H

Example
Partition functions count graph homomorphisms

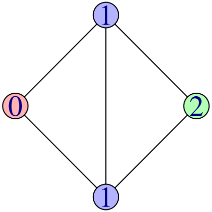

Graph $G = (V, E)$

Configuration
$\sigma \in \{0, 1, 2\}^V$ with weight $w(\sigma) = 1$ is a homomorphism from G to H.

Graph H

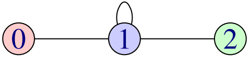

Homomorphism from G to H: A map $\sigma \in V(H)^V$ that maps every edge of G to an edge of H

A: Adjacency matrix of H

$$A = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

Z_G: number of homomorphisms from G to H.

Why computing partition fns is called “counting”

Complexity of computing partition functions

Dichotomy Theorem: For every symmetric matrix A, one of the following holds:
(1) the corresponding partition function Z_G can be computed in polynomial time (as a function of n, the number of vertices of G), or
(2) Z_G is #P-hard to compute.

We can tell which, given A.

- Dyer and Greenhill 2000: 0-1-matrices
- Bulatov and Grohe 2005: non-negative real algebraic matrices.
- G, Grohe, Jerrum, Thurley 2010: real algebraic matrices.
- Cai, Chen, Lu 2013: complex algebraic matrices.

the nonnegative real case

$$A = \begin{array}{c|ccccc} & 0 & 1 & 2 & 3 & 4 \\ \hline 0 & 0 & 3.5 & 0 & 0 & 0 \\ 1 & 3.5 & 1 & 1 & 0 & 0 \\ 2 & 0 & 1 & 0 & 0 & 0 \\ 3 & 0 & 0 & 0 & 0 & 1 \\ 4 & 0 & 0 & 0 & 1 & 0 \end{array}$$

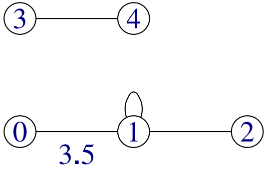

corresponding weighted graph

Each non-bipartite connected component of this graph corresponds to one block and each bipartite connected component corresponds to two blocks.

Computing Z_G is in polynomial time if the rank of every block of A is 1 and #P-hard otherwise.

When A can have negative numbers

Roughly, computing Z_G is tractable if each of the blocks of A can be written as a tensor product of a positive matrix of rank 1 and a tractable Hadamard matrix.

$$\begin{pmatrix} 2 & 3 \\ 4 & 6 \end{pmatrix} \otimes \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = \begin{pmatrix} 2 & 3 & 2 & 3 \\ 4 & 6 & 4 & 6 \\ 2 & 3 & -2 & -3 \\ 4 & 6 & -4 & -6 \end{pmatrix}$$

Hadamard matrices

A Hadamard matrix is a square matrix H with entries from $\{-1, 1\}$ such that $H \cdot H^T$ is a diagonal matrix.

$$\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}.$$

(the order is 1, 2, or a multiple of 4 - open whether there is one for every multiple of 4)

- A symmetric Hadamard matrix H is tractable if it has a "quadratic representation" and hard otherwise.

A "quadratic representation" for order 2^k (roughly)

$$A = \begin{array}{c|cccc} & 00 & 01 & 10 & 11 \\ \hline 00 & & & & \\ 01 & & -1 & & \\ 10 & & & & \\ 11 & & & & \end{array}$$

Row and column labels of A over $\mathbb{F}_2^k$ (here $k = 2$)

multivariate polynomial $h(X_1, \ldots, X_k, Y_1, \ldots, Y_k)$ over $\mathbb{F}_2$ of degree at most 2 such that

$h(0001) = 1 \Leftrightarrow H_{00,01} = -1$.

Approximating the partition function

Open even for these special cases

- A is a (symmetric) 0-1 matrix (counting graph homomorphisms) (probably won't have time)
- A is 2×2. (this talk)

$$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

Normalisation to $A_{01} = A_{10} = 1$ is wlg since easy to compute if they are 0

$$w(\sigma) = \prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)}.$$

$$Z_G = \sum_{\sigma \in \{0,1\}^V} w(\sigma).$$

Above $\beta\gamma = 1$. Reduction to Ising ($\beta = \gamma$) with consistent fields

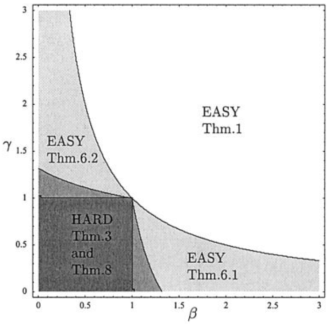

G, Jerrum, Paterson, 2003

$$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

$$w(\sigma) = \prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)}.$$

$$Z_G = \sum_{\sigma \in \{0,1\}^V} w(\sigma).$$

No FPRAS unless NP=RP

Easiness extends below hyperbola, hardness outside of the square

understanding $\beta\gamma < 1$

1. Inputs in $\mathcal{G}_\Delta$ (graphs of degree at most Δ)
2. Generalise the notion of Gibbs distribution to infinite graphs.

   $\Pr(\cdot = \sigma) = w(\sigma)/Z_G$.

   Infinite G: For any finite subgraph H, marginal distribution, conditioned on $\sigma(G \setminus H)$, is proportional to $w(\cdot)$.

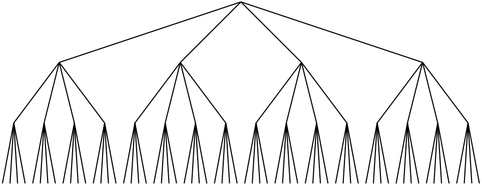

- Non-uniqueness: If the infinite Δ-regular tree has multiple Gibbs measures then no FPRAS/FPTAS for Δ-regular graphs G (NP-hard).
- Uniqueness: If $\forall d \leqslant \Delta$ the infinite d-regular tree has a unique Gibbs measure then there is an FPTAS for Z_G for $G \in \mathcal{G}_\Delta$.

Sly 2010; Galanis, Štefankovič, Vigoda 2012 Sly, Sun 2012
Weitz 2006; Sinclair, Srivastava, Thurley 2011; Li, Lu, Yin 2012

Uniqueness or not?

$f(x) = \left(\frac{\beta x+1}{x+\gamma}\right)^{\Delta-1}$

x^* is the unique positive fixed-point — the solution to $x^* = f(x^*)$.

Uniqueness: $|f'(x^*)| \leqslant 1$

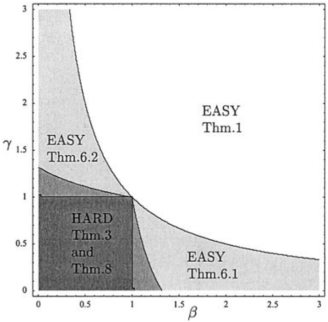

• $0 \leqslant \beta < 1$ and $0 < \gamma \leqslant 1$: non-uniqueness on the infinite Δ-regular tree for all sufficiently large Δ.

• $0 \leqslant \beta < 1$ and $\gamma > 1$: uniqueness holds on the infinite Δ-regular tree for all sufficiently large Δ.

The uniqueness threshold is not monotonic in Δ. (It is possible to be in uniqueness for Δ, but in non-uniqueness for some $d < \Delta$!)

The boundary between hard and easy is the "in uniqueness for all Δ" curve

What if β and γ can be negative?

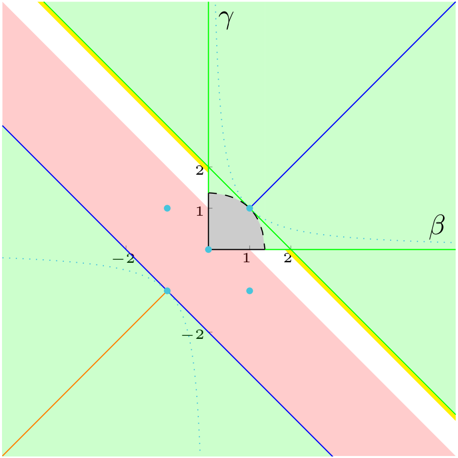

$$Z_G = \sum_{\sigma \in \{0,1\}^V} \prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)}.$$

Yumou Fei, LG, Pinyan Lu, ITCS 2024

$$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

- sky blue points: poly-time exact (already seen)
- : pos quadrant: dotted line “is” uniqueness-for-all-Δ curve. $\beta\gamma = 1$ exact poly time.
- $\beta = \gamma$ ising: Earlier work with Jerrum, hard for $\beta = \gamma \in (-1, 0)$ (even #P-hard to find sign of Z_G) and equivalent to counting PMs for $\beta = \gamma < -1$ (orange line).

What if β and γ can be negative?

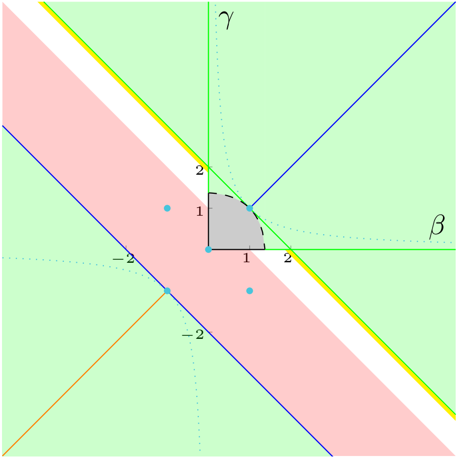

$$Z_G = \sum_{\sigma \in \{0,1\}^V} \prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)}.$$

Yumou Fei, LG, Pinyan Lu, ITCS 2024

$$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

- $\beta + \gamma = -2$, $\beta + \gamma = 2$.
- green: FPTAS on $\mathcal{G}_\Delta$ and FPRAS (all degrees)
- blue line $\beta + \gamma = -2$ has FPRAS
- red: #P-hard even to determine sign
- yellow “strip” (width $\to 0$ as move from origin) has FPTAS on $\mathcal{G}_\Delta$ so $\beta + \gamma = 2$ not a threshold
- white: open. There are hard points near $(1, 0)$. Conjecture: extend non-uniqueness curve?

How to approximate Z_G

Theorem. *Fix* $\beta \neq \gamma$ *with* $|\beta + \gamma| > 2$.

1. *There is an FPRAS for* Z_G
2. *For any positive integer* Δ *there is an FPTAS for* Z_G *for* $G \in \mathcal{G}_\Delta$.

A polynomial generalising Z_G (recall the independence polynomial)

$$A = \begin{pmatrix} A_{00} & A_{01} \\ A_{10} & A_{11} \end{pmatrix} = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

For $G = (V, E)$ and $\mathbf{x} \in \mathbb{R}^V$, let

$$Z_G(\mathbf{x}) = \sum_{\sigma \in \{0,1\}^V} \left(\prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)} \prod_{v \in V} x_v^{\sigma(v)} \right).$$

$\mathbf{x}$ is the vector of external fields.

$Z_G = Z_G(\mathbf{1})$.

Setting $x_v = x$ for all $v \in V$, we get a univariate polynomial $Z_G(x)$.

The method of Barvinok 2016

$p(z) = a_0 + a_1 x + \cdots + a_d x^d$, $\quad a_i$'s complex

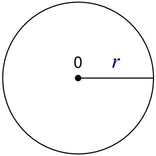

Suppose $p(x) \neq 0$ in disk of radius $r > 1$ around origin

Fix a branch of $f(x) = \ln p(x)$ for $|x| \leqslant 1$

e.g., fix the value of $\log(p(0))$ to lie in $(-\pi, \pi]$

truncate Taylor expansion of f around $x = 0$:

$f_N(x) = \sum_{j=0}^{N} \frac{x^j}{j!} f^{(j)}(0)$.

For $|x| \leqslant 1$, $f_N(x)$ is additively close to f (as a fn of d, N, r) so $|f(1) - f_N(1)| < \varepsilon$ for $N = O(\ln(d/\varepsilon))$

Patel/Regts 2017

For

$$p(x) = Z_G(x) = \sum_{\sigma \in \{0,1\}^V} \left(\prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)} \prod_{v \in V} x^{\sigma(v)} \right).$$

can compute the coefficients of $f_N(x)$ for graphs $G \in \mathcal{G}_\Delta$ by reduction to counting induced subgraphs from certain fixed graphs H into G.

So to get a FPTAS we just need to prove that there are no complex zeroes in a disk of radius > 1 in the complex plane

Theorem. *Fix $\beta \neq \gamma$ with $|\beta + \gamma| > 2$.*

1. *There is an FPRAS for Z_G*
2. *For any positive integer Δ there is an FPTAS for Z_G for $G \in \mathcal{G}_\Delta$.*

$$A = \begin{pmatrix} A_{00} & A_{01} \\ A_{10} & A_{11} \end{pmatrix} = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

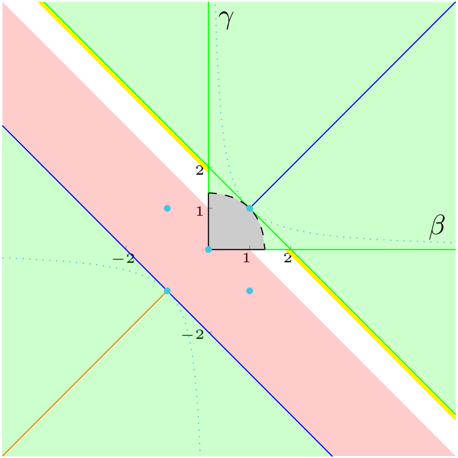

FPTAS for $G \in \mathcal{G}_\Delta$ when $|\beta + \gamma| > 2$ from zero-freeness in a disk of radius $r > 1$

Contraction method of Asano 1970 (used to give simple proof of Lee-Yang circle theorem and extended by Ruelle 1971) Suffices to show that $\gamma x_u x_v + x_u + x_v + \beta$ has no zeroes when $|x_u| < r$ and $|x_v| < r$

$Z_G(\mathbf{x}) = \sum_{\sigma \in \{0,1\}^V} \left(\prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)} \prod_{v \in V} x_v^{\sigma(v)} \right)$

$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}$.

A typical case: $\beta > \gamma$, $\beta + \gamma > 2$. Showing $\gamma x_u x_v + x_u + x_v + \beta$ has no zeroes when $|x_u| < r$ and $|x_v| < r$

The unique sol'n to $\gamma\, x\, g(x) + x + g(x) + \beta = 0$ is $g(x) = -(x+\beta)/(\gamma x + 1)$.

But g maps the open disk of radius r around the origin to the outside of this circle...

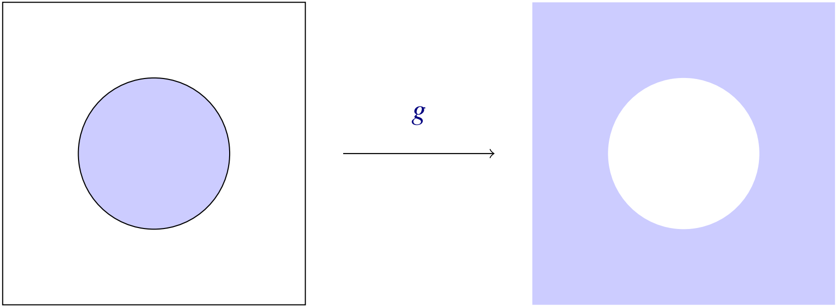

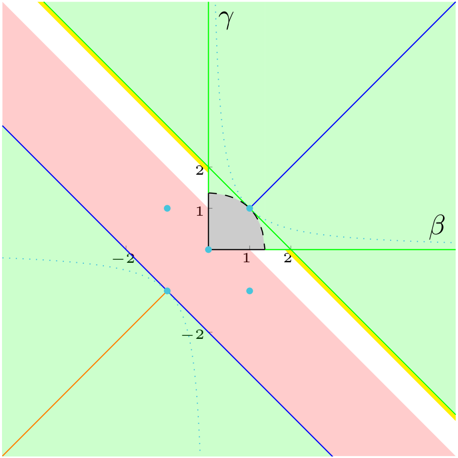

FPRAS for all G when
$\beta \neq \gamma$, $|\beta + \gamma| \geqslant 2$

$Z_G = \sum_{\sigma \in \{0,1\}^V} \prod_{\{u,v\} \in E} A_{\sigma(u),\sigma(v)}$
$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}$.