

The computational complexity of counting list H-colourings, and related problems

Mark Jerrum

Queen Mary, University of London

Simons Institute, 30th March 2016

Joint work with Andreas Galanis and Leslie Ann Goldberg (Oxford)

H-colouring

Fix an undirected graph H , which may have loops but not parallel edges.

Given a graph G , an *H-colouring of G* is a homomorphism from G to H , that is, a mapping $\sigma : V(G) \rightarrow V(H)$ such that $\{u, v\} \in E(G)$ implies $\{\sigma(u), \sigma(v)\} \in E(H)$, for all $u, v \in V(G)$.

Consider the problem:


Name #H-Col.

Instance A graph G .


Output The number of H-colourings of G .

We can view the vertices of H as a set of allowed colours $Q = V(H)$.


Examples




2-colourings



Independent sets



3-colourings



Widom-Rowlinson

H-colouring (continued)

Say that a graph is *irreflexive* if no vertices have loops, and *reflexive* if all vertices have loops. Say that a graph is *trivial* if it is an irreflexive complete bipartite graph or a reflexive complete graph.

Theorem (Dyer and Greenhill (2000))

If every connected component of H is trivial then $\#H\text{-Col} \in \text{FP}$; otherwise $\#H\text{-Col}$ is $\#P$ -complete.

We have only incomplete information on the complexity of *approximating* $\#H\text{-Col}$. But we do know that if H is connected and not trivial then $\#H\text{-Col}$ is $\#BIS$ -hard. (Leslie's talk on Monday.)

List H-colouring

Now consider the graph G together with a collection of sets $\mathbf{S} = \{S_v \subseteq Q : v \in V(G)\}$ specifying allowed colours at each of the vertices.

A *list H-colouring* of (G, \mathbf{S}) is an H-colouring σ of G satisfying $\sigma(v) \in S_v$, for all $v \in V$.

Name #List-H-Col.

Instance A graph G and a collection of colour sets $\mathbf{S} = \{S_v \subseteq Q : v \in V(G)\}$, where $Q = V(H)$.

Output The number of list H-colourings of (G, \mathbf{S}) .

It is of no importance whether we allow or disallow loops in G since a loop at vertex $v \in V(G)$ can be encoded within the set S_v .

List H-colouring (continued)

We are interested in mapping the computational complexity of $\#List-H-Col$ as a function of H . NB: the graph H is part of the problem specification, not part of the instance.

The positive direction of Dyer and Greenhill's result holds in the presence of lists, so their dichotomy applies equally to $\#List-H-Col$.

We therefore concentrate on the complexity of *approximating* $\#List-H-Col$.

The complexity of approximate counting

A *fully-polynomial randomised approximation scheme* or FPRAS is a randomised algorithm that produces approximate solutions within specified relative error with high probability in polynomial time.

An AP-reduction from problem Π to problem Π' is a randomised Turing reduction that yields close approximations to Π when provided with close approximations to Π' . It meshes with FPRAS in the sense that the existence of an FPRAS for Π' implies the existence of an FPRAS for Π .

The problem of counting satisfying assignments of a Boolean formula is denoted by $\#\text{SAT}$. Every counting problem in $\#\text{P}$ is AP-reducible to $\#\text{SAT}$. The hardest counting problems in $\#\text{P}$ are those that are *$\#\text{SAT}$ -equivalent*, i.e., AP-interreducible with $\#\text{SAT}$.

How hard can approximate counting be?

Even the hardest problems, such as $\#SAT$ can be solved, in the FPRAS sense, by a polynomial-time algorithm equipped with an NP oracle. This follows from the bisection technique of Valiant and Vazirani. So approximate counting is not *that* hard.

Assuming $NP \neq RP$, no $\#SAT$ -equivalent problem admits an FPRAS.

There are many problems that are not known to admit an FPRAS, but at the same time don't appear to be $\#SAT$ -equivalent. One of them is $\#BIS$, the problem of counting independent sets in a bipartite graph.

Problems equivalent to $\#BIS$

Quite a large number of problems are known to be AP-interreducible with $\#BIS$: the ferromagnetic Ising model with inconsistent fields, downsets in a partial order, the partition function of the Widom-Rowlinson model, stable matchings, etc. We call these problems $\#BIS$ *equivalent*.

The result Leslie talked about on Monday can be stated: if H is non-trivial, then $\#H\text{-Col}$ is $\#BIS$ -hard.

It transpires that we now have enough categories — FPRAS, $\#BIS$ -equivalent and $\#SAT$ -equivalent — for our discussion.

Observation about list colouring

Suppose H' is an induced subgraph of H . The computational complexity of $\#List-H-Col$ is not less than that of $\#List-H'-Col$. (We can use the lists to carve out H' from H ; this gives an AP-reduction from $\#List-H'-Col$ to $\#List-H-Col$.)

Another way of looking at this: the set of graphs H satisfying “ $\#List-H-Col$ is $\#BIS$ -easy” is a hereditary graph class (a set of graphs closed under taking induced subgraphs).

This observation was exploited by Feder, Hell and Huang to classify the complexity of the decision version “does graph G have an H -colouring” in terms of hereditary graph classes (1998–2003).


Two hereditary graph classes

Say that a 0,1-matrix $A = (A_{i,j} : 1 \leq i \leq n, 1 \leq j \leq m)$ has *staircase form* if the 1s in each row are contiguous and the following condition is satisfied: letting $\alpha_i = \min\{j : A_{i,j} = 1\}$ and $\beta_i = \max\{j : A_{i,j} = 1\}$, we require that the sequences (α_i) and (β_i) are non-decreasing.

A graph is a bipartite permutation graph if the rows and columns of its biadjacency matrix can be (independently) permuted so that the resulting biadjacency matrix has staircase form.

A graph is a proper interval graph if the rows and columns of its adjacency matrix can be (simultaneously) permuted so that the resulting adjacency matrix has staircase form.

Example of a proper interval graph



... and its staircase presentation:

$$\begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix}.$$


Main result

Theorem


Suppose that H is a connected undirected graph (possibly with loops).

- (i) If H is an irreflexive complete bipartite graph or a reflexive complete graph then $\#List-H-Col$ is in FP .*
- (ii) Otherwise, if H is an irreflexive bipartite permutation graph or a reflexive proper interval graph then $\#List-H-Col$ is $\#BIS$ -equivalent.*
- (iii) Otherwise, $\#List-H-Col$ is $\#SAT$ -equivalent.*


Examples (reprise)




2-colourings




Independent sets



3-colourings



Widom-Rowlinson



Whatever

Sketch proof

We'll concentrate on the #SAT-hardness part.


Fix a graph H that is covered by part (iii) of the theorem. We need to identify an induced subgraph of H that we know to be hard.

If H is neither reflexive nor irreflexive, then it must contain K'_2 , the complete graph on two vertices with a single loop, as an induced subgraph. #List- K'_2 -Col is counting independent sets, which is #SAT-equivalent.


So now suppose H is irreflexive, but not a bipartite permutation graph.

Excluded subgraph characterisation


From the excluded subgraph characterisation of bipartite permutation graphs, we know that H must contain one of X_3 , X_2 , T_2 or a cycle C_ℓ of length $\ell \neq 4$ as an induced subgraph.



X_3



X_2



T_2

Just need to show that $\#List-X_3\text{-Col}$, $\#List-X_2\text{-Col}$, $\#List-T_2\text{-Col}$ and $\#List-C_{\ell \neq 4}\text{-Col}$, are $\#SAT$ -equivalent.


The case of X_3

We consider the case X_3 in detail, and the others more swiftly, as they all follow the same general pattern. The gadget in this case is

$$\{1, 2\} - \{4, 7\} - \{3, 6\} - \{4, 5\} - \{2, 1\},$$

i.e., a path of length four with $S_0 = \{1, 2\}, \dots, S_4 = \{2, 1\}$.

Note that, reading along the path, the first element of each pair traces out a path in X_3 , and the second element also.




The case of X_3

We consider the case X_3 in detail, and the others more swiftly, as they all follow the same general pattern. The gadget in this case is

$$\{1, 2\} - \{4, 7\} - \{3, 6\} - \{4, 5\} - \{2, 1\},$$

i.e., a path of length four with $S_0 = \{1, 2\}, \dots, S_4 = \{2, 1\}$.

Note that, reading along the path, the first element of each pair traces out a path in X_3 , and the second element also.




The case of X_3

We consider the case X_3 in detail, and the others more swiftly, as they all follow the same general pattern. The gadget in this case is

$$\{1, 2\} - \{4, 7\} - \{3, 6\} - \{4, 5\} - \{2, 1\},$$

i.e., a path of length four with $S_0 = \{1, 2\}, \dots, S_4 = \{2, 1\}$.

Note that, reading along the path, the first element of each pair traces out a path in X_3 , and the second element also.




The case of X_3

We consider the case X_3 in detail, and the others more swiftly, as they all follow the same general pattern. The gadget in this case is

$$\{1, 2\} - \{4, 7\} - \{3, 6\} - \{4, 5\} - \{2, 1\},$$

i.e., a path of length four with $S_0 = \{1, 2\}, \dots, S_4 = \{2, 1\}$.

Note that, reading along the path, the first element of each pair traces out a path in X_3 , and the second element also.




The case of X_3

We consider the case X_3 in detail, and the others more swiftly, as they all follow the same general pattern. The gadget in this case is

$$\{1, 2\} - \{4, 7\} - \{3, 6\} - \{4, 5\} - \{2, 1\},$$

i.e., a path of length four with $S_0 = \{1, 2\}, \dots, S_4 = \{2, 1\}$.

Note that, reading along the path, the first element of each pair traces out a path in X_3 , and the second element also.



The case of X_3 (continued)

Recall the gadget:

$$\{1, 2\} - \{4, 7\} - \{3, 6\} - \{4, 5\} - \{2, 1\}.$$

The interactions between spins at the endpoints of a single edge can be described by a 2×2 matrix. For example, for the first edge the matrix is

$$\begin{matrix} & 4 & 7 \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \end{matrix}$$

The possible 2×2 matrices that can occur are:

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

The case of χ_3 (continued)

The interactions between spins at the endpoints of the gadget are given by a product of these 2×2 matrices, specifically,

$$D' = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 2 & 3 \\ 3 & 5 \end{pmatrix}.$$

Note that $\det D' = 1 > 0$. However there is a “twist” in the labellings of the rows and columns

$$D' = \begin{matrix} & & 2 & 1 \\ & 1 & & \\ 2 & & \begin{pmatrix} 2 & 3 \\ 3 & 5 \end{pmatrix} & \end{matrix}$$

so the gadget is actually antiferromagnetic.

The case of X_3 (continued)

The fact that $\det D' = 1$, and the twisting of the row and column labels is an automatic consequence of the way the gadget is constructed.

“Untwisting” the matrix D' yields

$$D = \begin{matrix} & \begin{matrix} 1 & 2 \end{matrix} \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} 3 & 2 \\ 5 & 3 \end{pmatrix} \end{matrix}.$$

There is an involution of X_3 that transposes vertices 1 and 2. This yields another gadget with the roles of 1 and 2 reversed. Now place the current gadget in parallel with its twin.

The case of χ_3 (concluded)

The symmetrised gadget has interaction matrix

$$D = \frac{1}{2} \begin{pmatrix} 1 & 2 \\ 9 & 10 \\ 10 & 9 \end{pmatrix}$$

which can be recognised as an antiferromagnetic Ising model.

The partition function of an antiferromagnetic Ising model is #SAT-equivalent.

#List- χ_2 -Col, #List- T_2 -Col and #List- $C_{\ell \neq 4}$ -Col can be handled analogously.

Completing the proof

Reflexive proper interval graphs follow the same pattern.

The $\#BIS$ -hardness part of the proof is also done by extracting $\#BIS$ -hard induced subgraphs in H .

The $\#BIS$ -easiness part exploits the matrix characterisation of bipartite permutation graphs and proper interval graphs. The technique appears (slightly disguised) in a paper by Chen, Dyer, Goldberg, Jerrum, Lu, McQuillan and Richerby.

Comparison with the weighted case

Look at binary spin systems with non-negative weights. This involves

- replacing the $0, 1$ (bi)adjacency matrix of H by a weighted (bi)adjacency matrix with general non-negative real weights;
- replacing the lists by unary functions from Q to $\mathbb{R}_{\geq 0}$.

The natural replacement for the staircase condition (in the bipartite case) is: there are permutations of the rows and columns of A such that every 2×2 submatrix of A has non-negative determinant.

With these changes, the trichotomy holds in the weighted situation too (Goldberg and Jerrum, 2015).

Comparison with the weighted case (continued)

However the unweighted result is not a special case of the weighted one as weighted binary interactions have to be matched by general (non-negative real) unary weights.

Example. Let

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

This is the interaction matrix of a modified independent set model. Li, Lu and Yin show (following Weitz) that there is an FPTAS for this model. (Lists are not a hinderance.)

However, $\det A < 0$ so, with unary weights, computing the partition function becomes #SAT-equivalent.

Bounded degree

An important advantage of the unweighted result is that the trichotomy continues to hold if the instance has bounded degree, in fact degree 6. This is tight as, for example, there is a FPTAS for independent sets in a graph of maximum degree 6 (Weitz, 2006).

Comparison with general counting CSPs

Chen, Dyer, Goldberg, Jerrum, Lu, McQuillan and Richerby proved a classification theorem for (non-negative real) weighted $\#$ CSPs, with unary functions given free.

In the bijunctive case (all allowed functions have arity at most 2), it provides a trichotomy, which substantially generalises the one just cited for a single binary function.

Although this trichotomy is decidable, it does not have the explicit form of the one described here.