London Taught Course Centre

2023/24 examination

Graph Theory

Answers

1 This is a standard theorem (Google will lead you to it directly); it's more or less what Kruskal actually proved (the version in the lecture notes, week 2, of unlabelled trees, is equivalent to the $|S| = 1$ case).

The proof is an easy modification of the lecture notes proof; one strengthens the statement to ordered labelled topological minor on rooted labelled trees (as in the lecture notes) and follows the proof there, adding 'labelled' wherever it is necessary. The only point where one needs to do more is that it does not suffice to find one (i, j) such that the rooted forest $T_i - r_i$ is an ordered labelled topological minor of $T_j - r_j$, since we cannot simply map r_i to r_j; the label of r_i need not precede the label of r_j. But by Propositions 6 and 8 of the lecture notes, at this point in the proof it follows that there is actually some infinite chain $i_1 < i_2 < \dots$ such that $T_{i_a} - r_{i_a}$ is an ordered labelled topological minor of $T_{i_b} - r_{i_b}$ whenever $a < b$. Since the labels are well-quasi-ordered, the sequence of labels of $r_{i_1}, r_{i_2}, \dots$ cannot contain an infinite strictly descending sequence or infinite antichain. Proposition 6 now says there exist $a < b$ such that r_{i_a} has label equal to or preceding that of r_{i_b}, and we thus have the desired fact that T_{i_a} is an ordered labelled topological minor of T_{i_b}.

2 (a) This is (a case of) the Andrásfai-Erdős-Sós theorem, and just citing it is fine. Here is a proof.

Let first I_1 and I_2 be disjoint independent sets in $V(H)$ of size greater than $2m/5$. Such sets exist since we can pick an edge xy and then $N(x)$ and $N(y)$ are examples of such sets. Either I_1 and I_2 form a bipartition of H and we are done, or there is some z in neither I_1 nor I_2. If z has a neighbour w in I_1, then z has at most $m/5$ neighbours outside I_1, otherwise w is adjacent neither to the vertices of I_1 nor to the vertices $N(z) \setminus I_1$, and this leaves less than $2m/5$ possible neighbours, a contradiction.

It follows that if z has a neighbour in I_1, then it has $m/5$ neighbours in I_1. If now also z has a neighbour y in I_2, then y cannot be adjacent to any of the vertices of I_2, nor to any vertex of $N(z) \cap I_1$, and these two sets are disjoint of sizes at least $2m/5$ and $m/5$ respectively. Again this is a contradiction as it leaves only $2m/5$ vertices to be neighbours for y.

It follows that we can add z to one of I_1 and I_2 to get a larger pair of disjoint independent sets in $V(H)$; repeating this we obtain the desired bipartition of H.

(b) This is a classic NP-completeness result, mentioned as such in the lectures. If you figured it out on your own, well done—but you should have realised that this must be easy to find online, and indeed there are several different routes that show up on the first page of a Google search; or you could just cite it.

(c) There are several ways to do this. Probably the easiest is to take a vertex-minimal graph F which does not contain K_3 and which cannot be coloured with C colours; we proved such graphs exist in the course. Suppose $t = v(F)$. Now for each sufficiently large n we construct an n-vertex graph G by 'blowing up' F; that is, by replacing vertices of F with independent sets and edges with complete bipartite graphs. We choose the sizes of these independent sets to be $\lfloor n/t \rfloor$ and $\lceil n/t \rceil$ in order to obtain n vertices in total. Now G does not contain a copy of K_3, and since each vertex of F is in an edge (otherwise F would not be minimal) each vertex of G has degree at least

$n/(2t)$. Provided that $\log n > 2t$, which is true for sufficiently large n, the graph G is an example as desired.

3 (a) Summing over v, we have $\sum_{v\in V(H)} \sum_{\substack{T\in K_3(H)\\ v\in T}} w(T) \le v(H)$; since each K_3-copy appears exactly three times in the sum on the left, we have $3\sum_{T\in K_3(H)} w(T) \le v(H)$ and the desired inequality follows.

(b) Given $d > 0$, we can assume $d \le 2$ (since otherwise $R(G)$ will be empty and the question is trivial) and we choose $\varepsilon = 10^{-6}d$. Let K be the constant returned by the Regularity Lemma, and assume $n > 10^8 K$. Let $R(G)$ be as in the question.

Given a fractional triangle factor w of $R(G)$ with weight at least $\frac{1}{3}$, we begin the following algorithm. Start with y the everywhere zero fractional triangle factor of $R(G)$ and S the empty set. Choose $T \in K_3\big(R(G)\big)$ such that $y(T) \le 0.9999w(T)$, and pick a triangle in G with one vertex in each part indexed by T which is vertex-disjoint from the triangles in S. Add this triangle to S, and increase $y(T)$ by $\frac{1}{|V_1|}$. Repeat until $y(T) \ge 0.9999w(T)$ for every $T \in K_3\big(R(G)\big)$.

This algorithm can only fail if at some point it is not possible to pick a triangle in G as stated. So we aim to show this cannot occur. Our first task is to argue that S cannot cover too many vertices in any given V_j with $1 \le j \le k$. By definition, for any given $T \in K_3\big(R(G)\big)$ that contains j, the number of triangles in S that we picked crossing the parts indexed by T is at most $0.9999w(T)|V_1| + 1 \le 0.99999w(T)|V_j|$. Summing over all choices of T containing j, the number of triangles in S that meet V_j is at most $0.99999\sum_{j\in T} w(T)|V_j| \le 0.99999|V_j|$, so that at least $10^{-5}|V_j|$ vertices of V_j are not contained in triangles of S.

Now given any ijj', let V_i' be the vertices of V_i not covered by S, and define similarly V_j' and $V_{j'}'$. As in lectures, the number of vertices of V_i' with fewer than $(d-\varepsilon)|V_j'|$ neighbours in V_j' is at most $\varepsilon|V_i|$, and similarly with fewer than $(d-\varepsilon)|V_{j'}'|$ neighbours in $V_{j'}'$. Since $|V_i'| > 2\varepsilon|V_i|$, there is a vertex u of V_i' with at least $(d-\varepsilon)|V_j'|$ neighbours in V_j' and at least $(d-\varepsilon)|V_{j'}'|$ neighbours in $V_{j'}'$. Since $(d-\varepsilon)10^{-5} > \varepsilon$, the density of $\big(N_G(u)\cap V_j', N_G(u)\cap V_{j'}'\big)$ is at least $d-\varepsilon$, in particular there is an edge vw between these two sets and this gives a triangle uvw crossing the parts indexed by T which is disjoint from all members of S. This shows the above algorithm does not fail.

Since at the termination of the algorithm the number of triangles in S which meet V_j is at least $0.9999|V_j|$ (proved much as the upper bound above: we have $\sum_{j\in T} w(T) = 1$ from (a))), the number of vertices covered in total by S is at least $0.9999(n - |V_0|) \ge 0.9999(1-\varepsilon)n \ge 0.999n$. Thus $|S| \ge 0.333n$.