Assignment 1

CSCI 4113/6101
Instructor: Norbert Zeh
Solutions

Question 1: Maximum Matching as an LP

1a

Part (b) of this question suggests that a graph for which (1) and (2) have optimal solutions with different objective function values must be non-bipartite, that is, it must have a cycle of odd length. The simplest such graph is itself a cycle of length 3:

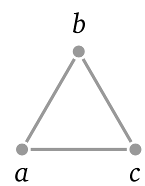

Since every edge shares an endpoint with each of the other two edges, a maximum matching of this graph contains only one edge, say $\{a, b\}$. The constraints in (1) also imply this: If we set $x_{a,b} = 1$, then the constraints

$$x_{a,b} + x_{a,c} \le 1,$$
$$x_{a,b} + x_{b,c} \le 1$$

imply that $x_{a,c} = x_{b,c} = 0$. So, an optimal solution of (1) has objective function value 1 (not a surprise because (1) is an ILP formulation of the maximum matching problem.)

An optimal solution of (2) is

$$x_{a,b} = x_{a,c} = x_{b,c} = \frac{1}{2},$$

which has objective function value $\frac{3}{2}$.

This is indeed a feasible solution because

$$x_{a,b} + x_{a,c} = 1,$$
$$x_{a,b} + x_{b,c} = 1,$$
$$x_{a,c} + x_{b,c} = 1.$$

It is an optimal solution because for any feasible solution, we have

$$\begin{aligned} x_{a,b} + x_{a,c} &\le 1, \\ x_{a,b} + x_{b,c} &\le 1, \\ x_{a,c} + x_{b,c} &\le 1, \end{aligned}$$

and, therefore,

$$2x_{a,b} + 2x_{a,c} + 2x_{b,c} \le 3,$$

that is,

$$x_{a,b} + x_{a,c} + x_{b,c} \le \frac{3}{2}.$$

1B

Consider an optimal solution $\hat{x}$ of (2), and assume that $\hat{x}$ minimizes the number of edges $e \in E$ such that $0 < \hat{x}_e < 1$. We prove that $\hat{x}$ must be an integral solution. Assume the contrary. Then let $H = (V, E')$ be the subgraph of G such that $E' = \{e \in E \mid 0 < \hat{x}_e < 1\}$.

Case 1: H contains a cycle

Let $C = \langle v_0, \dots, v_k \rangle$ be a cycle in H. Then k is even because every cycle in a bipartite graph has even length. Let

$$\delta = \min(\min\{1 - \hat{x}_{v_{2i},v_{2i+1}} \mid 0 \le i < k/2\}, \min\{\hat{x}_{v_{2i+1},v_{2i+2}} \mid 0 \le i < k/2\}).$$

In words, we partition the edges in C into the even-numbered edges and odd-numbered edges. δ is the smaller of the maximum amount by which we can increase the values associated with even-numbered edges and the maximum amount by which we can increase the values associated with odd-numbered edges while keeping all values between 0 and 1.

Now let $\tilde{x}$ be defined as

$$\tilde{x}_e = \begin{cases} \hat{x}_e & \text{if } e \notin \mathrm{C} \\ \hat{x}_e + \delta & \text{if } e = \{x_{2i}, x_{2i+1}\}, \text{ for some } i \in [k/2-1]_0 \\ \hat{x}_e - \delta & \text{if } e = \{x_{2i+1}, x_{2i+2}\}, \text{ for some } i \in [k/2-1]_0, \end{cases}$$

that is, we increase the values associated with even-numbered edges in C by δ, decrease the values associated with odd-numbered edges in C by δ, and leave the values associated with all other edges unchanged.

For every edge $e \notin C$, we have $\tilde{x}_e = \hat{x}_e$. Since $0 \le \hat{x}_e \le 1$ (because $\hat{x}$ is a feasible solution of (2)), this implies that $0 \le \tilde{x}_e \le 1$.

If e is an even-numbered edge of C, then $0 \le \hat{x}_e < \hat{x}_e + \delta = \tilde{x}_e \le 1$, because $\delta \le 1 - \hat{x}_e$.

If e is an odd-numbered edge of C, then $0 \le \hat{x}_e - \delta = \tilde{x}_e < \hat{x}_e \le 1$, because $\delta \le \hat{x}_e$.

Thus, $0 \le \tilde{x}_e \le 1$, for all $e \in E$.

For every vertex $v \in V$, if $v \notin C$, then $\tilde{x}_{v,w} = \hat{x}_{v,w}$, for every edge $\{v, w\}$ incident to v. Thus,

$$\sum_{w \in V} \tilde{x}_{v,w} = \sum_{w \in V} \hat{x}_{v,w} \leq 1.$$

If $v = v_i$, for some $i \in [k-1]_0$, then assume w.l.o.g. that i is even. The case when i is odd is analogous. Then $\tilde{x}_{v_{i-1},v_i} = \hat{x}_{v_{i-1},v_i} - \delta$, $\tilde{x}_{v_i,v_{i+1}} = \hat{x}_{v_i,v_{i+1}} + \delta$, and $\tilde{x}_{v_i,w} = \hat{x}_{v_i,w}$, for every other edge $\{v_i, w\}$ incident to v_i. Thus,

$$\sum_{w \in V} \tilde{x}_{v_i,w} = \sum_{w \in V} \hat{x}_{v_i,w} + \delta - \delta = \sum_{w \in V} \hat{x}_{v_i,w} \leq 1.$$

This shows that

$$\sum_{w \in V} \tilde{x}_{v,w} \leq 1,$$

for every vertex $v \in V$. Since we also have $0 \leq \tilde{x}_e \leq 1$, for every edge $e \in E$, this shows that $\tilde{x}$ is a feasible solution of (2).

It is an *optimal* solution of (2) because C has even length. In particular, this implies that we have $\tilde{x}_e = \hat{x}_e + \delta$ for half the edges in C, and $\tilde{x}_e = \hat{x}_e - \delta$ for the other half of the edges in C. For all edges $e \notin C$, we have $\tilde{x}_e = \hat{x}_e$. Therefore,

$$\sum_{e \in E} \tilde{x}_e = \sum_{e \in E} \hat{x}_e,$$

which implies that $\tilde{x}$ is optimal because $\hat{x}$ is optimal and $\tilde{x}$ is feasible.

Now observe that because the edges in C are the only edges with $\tilde{x}_e \neq \hat{x}_e$ and all edges in C satisfy $0 < \hat{x}_e < 1$, we have

$$|\{e \in E \mid 0 < \tilde{x}_e < 1\}| \leq |\{e \in E \mid 0 < \hat{x}_e < 1\}|. \tag{S1}$$

By the choice of δ, we have $\delta = 1 - \hat{x}_e$ for an even-numbered edge e in C or $\delta = \hat{x}_e$ for an odd-numbered edge in C. For this edge e, we have $\tilde{x}_e \in \{0, 1\}$ and $0 < \hat{x}_e < 1$, the latter because $e \in C \subseteq H$. Thus, (S1) is strict. Since $\tilde{x}$ is an optimal solution of (2), this contradicts the choice of $\hat{x}$. Therefore, H cannot contain a cycle.

CASE 2: *H* DOES NOT CONTAIN A CYCLE

Since H cannot contain a cycle, it must be a forest. If H has no edges, then $\hat{x}$ is an integral solution. So assume that H is not empty. Then there exist two leaves x_0 and x_k of H with a path $P = \langle x_0, \ldots, x_k \rangle$ between them. We use P to define a solution $\tilde{x}$ similar to the previous case. In particular, we define

$$\delta = \min(\min\{1 - \hat{x}_{v_{2i},v_{2i+1}} \mid 0 \leq i < k/2\}, \min\{\hat{x}_{v_{2i-1},v_{2i}} \mid 1 < i \leq k/2\}).$$

and then set $\tilde{x}_e = \hat{x}_e + \delta$ for all even-numbered edges in P, $\tilde{x}_e = \hat{x}_e - \delta$ for all odd-numbered edges in P, and $\tilde{x}_e = \hat{x}_e$ for all other edges of H.

By the same argument as in the previous case,

$$|\{e \in E \mid 0 < \tilde{x}_e < 1\}| < |\{e \in E \mid 0 < \hat{x}_e < 1\}|. \tag{S2}$$

Also by the same argument as in the previous case, all edges in H satisfy $0 \le \tilde{x}_e \le 1$, and we have

$$\sum_{w \in V} \tilde{x}_{v,w} = \sum_{w \in V} \hat{x}_{v,w} \le 1,$$

for all $v \in V \setminus \{v_0, v_k\}$. For v_0, note that $\{v_0, v_1\}$ is the only edge in G incident to v_0 with $0 < \hat{x}_{v_0,v_1} < 1$. Thus, every other edge $\{v_0, w\}$ incident to x_0 satisfies $\hat{x}_{v_0,w} \in \{0,1\}$. Since $\hat{x}_{v_0,v_1} > 0$ and

$$\sum_{w \in V} \hat{x}_{v_0,w} \le 1,$$

this implies that $\hat{x}_{v_0,w} = 0$, for each such edge. Since we have $\tilde{x}_{v_0,w} = \hat{x}_{v_0,w}$ for each such edge, this shows that

$$\sum_{w \in V} \tilde{x}_{v_0,w} = \tilde{x}_{v_0,v_1} \le 1.$$

An analogous argument shows that

$$\sum_{w \in V} \tilde{x}_{v_k,w} = \tilde{x}_{v_{k-1},v_k} \le 1.$$

Thus, $\tilde{x}$ is once again a feasible solution of (2) that satisfies (S2).

Finally, observe that P contains at least as many even-numbered edges as odd-numbered edges. Thus,

$$\sum_{e \in E} \tilde{x}_e \ge \sum_{e \in E} \hat{x}_e,$$

that is, $\tilde{x}$ is an optimal solution, once again contradicting the choice of $\hat{x}$.

1c

Part (b) of the question shows that a solution $\hat{x}$ of (1) is an optimal solution of (1) if and only if it is an optimal solution of (2). Thus, since $\sum_{e \in E} \hat{x}_e = |M = \{e \in E \mid \hat{x}_e = 1\}|$, for any such solution $\hat{x}$, we only need to prove that every solution of (1) defines a matching M of G and that every matching M of G is defined by an integral solution of (2).

To this end, observe first that every subset $M \subseteq E$, matching or not, can be defined as $M = \{e \in E \mid \hat{x}_e = 1\}$, for

$$\hat{x}_e = \begin{cases} 1 & \text{if } e \in M \\ 0 & \text{otherwise.} \end{cases}$$

This shows that there exists a bijections between subsets of E and assignments $\hat{x}$ of values in $\{0,1\}$ to the edges in E. Thus, it suffices to prove that $M = \{e \in E \mid \hat{x}_e = 1\}$ is a matching if and only if $\hat{x}$ is a feasible solution of (1), that is, an integral feasible solution of (2).

We have $0 \le \hat{x}_e \le 1$, for all $e \in E$. Thus, feasibility of $\hat{x}$ depends only on the constraints

$$\sum_{w \in V} x_{v,w} \le 1 \quad \forall v \in V.$$

If every vertex $v \in V$ satisfies its corresponding constraint, then M contains at most one edge incident to v, so M is a matching. If some vertex $v \in V$ violates its corresponding constraint, then M contains at least two edges incident to v, so M is not a matching. Thus, M is a matching if and only if $\hat{x}$ is a feasible solution of (1), and we have a one-to-one correspondence between optimal solutions of (1) and maximum matchings of G.

1D

Here is an algorithm based immediately on the proof of (b): We start by computing an arbitrary optimal solution $\tilde{x}$ of (2). Now, as long as $\tilde{x}$ is not integral, we repeat the following steps: We construct the graph $H = (V, E')$ defined in (b). We try to find a cycle in H, which can be done in linear time using DFS, as you hopefully learned in CSCI 3110. If there exists such a cycle, we update $\tilde{x}$ as discussed in (b). The new solution has at least one fractional variable less. If H does not contain a cycle, then it must contain a degree-1 vertex (leaf). Running DFS from this leaf finds a path P from this leaf to another leaf, in linear time. Once again, we update $\tilde{x}$ as discussed in (b).

The correctness of this algorithm follows from (b): Each time we update $\tilde{x}$, this maintains feasibility of $\tilde{x}$, does not decrease the objective function value of $\tilde{x}$, and reduces the number of fractional variables in $\tilde{x}$. Thus, $\tilde{x}$ is a feasible solution of (2) at all times and, after at most m iterations, all variables in $\tilde{x}$ are integral, that is, $\tilde{x}$ is the desired solution $\hat{x}$ at this point.

Since each iteration of the algorithm takes linear time, as just discussed (finding C or P is the "hard" part; computing δ and updating the values associated with the edges in C or P takes two iterations over the edges in C or P and thus also takes linear time), and we just observed that we obtain an integral solution after at most m iterations, the conversion from a fractional optimal solution of (2) to an integral optimal solution of (2) takes at most $O((n+m)m)$ time.

Question 2: Vertex Colouring as an ILP

2A

As observed in the statement of the question, if the colouring uses k colours, we can assume that these k colours are called 0 through $k-1$. We can also assume that $k \le n$ because every graph can be coloured with n colours. Thus, we can represent an arbitrary colouring and the goal to minimize the number of colours used by the following constraints and objective function:

$$\begin{aligned} &\text{Minimize } k \\ \text{s.t. } & k - c_v \ge 1 && \forall v \in V \\ & c_v \in [n-1]_0 && \forall v \in V. \end{aligned}$$

We need to add constraints that ensure that two adjecent vertices have different colours. Here is how we do this. For every edge $\{u, v\}$, we introduce a variable $s_{u,v} \in \{0, 1\}$, and we impose the two constraints that

$$\begin{aligned} c_u - c_v + n s_{u,v} &\ge 1, \\ c_v - c_u - n s_{u,v} &\ge 1 - n. \end{aligned}$$

Note that for a fixed graph G, n is a constant, so these are indeed linear constraints. Why do they do the right thing? If $s_{u,v} = 0$, then the first constraint says that $c_u \geq c_v + 1$, so $c_u \neq c_v$. The second constraint is trivially satisfied because $c_u, c_v \in [n-1]_0$, so $1 - n \leq c_v - c_u \leq n - 1$. If $s_{u,v} = 1$, then the first constraint is trivially satisfied because, as just observed $c_u - c_v \geq 1 - n$, so $c_u - c_v + n \geq 1$. The second constraint then ensures that $c_v - c_u - n \geq 1 - n$, that is, $c_v - c_u \geq 1$, so $c_u \neq c_v$ again. By setting $s_{u,v} = 0$ or $s_{u,v} = 1$, we choose whether we want $c_u > c_v$ or $c_v > c_u$, but one of these two inequalities is being enforced.

The final ILP, then, is

$$
\begin{aligned}
&\text{Minimize } k \\
\text{s.t.} \quad & k - c_v \geq 1 && \forall v \in V \\
& c_u - c_v + n s_{u,v} \geq 1 && \forall \{u,v\} \in E \\
& c_v - c_u - n s_{u,v} \geq 1 - n && \forall \{u,v\} \in E \\
& c_v \in [n-1]_0 && \forall v \in V \\
& s_{u,v} \in \{0,1\} && \forall e \in \{u,v\}.
\end{aligned}
\tag{S3}
$$

2B

Consider an assignment $\hat{c}$ of colours to the vertices of G, with $\hat{c}_v \in [n-1]_0$, for all $v \in V$.

If $\hat{c}$ is a valid colouring, we can complete it to a solution $(\hat{c}, \hat{s}, \hat{k})$ of (S3) by setting

$$\hat{s}_{u,v} = \begin{Bmatrix} 0 & \text{if } \hat{c}_u > \hat{c}_v \\ 1 & \text{otherwise} \end{Bmatrix} \quad \forall \{u,v\} \in E$$

and

$$\hat{k} = 1 + \max\{\hat{c}_v \mid v \in V\}.$$

This solution $(\hat{c}, \hat{s}, \hat{k})$ is a feasible solution of (S3). Indeed, we clearly have $\hat{c}_v \in [n-1]_0$, for all $v \in V$, and $\hat{s}_{u,v} \in \{0,1\}$, for all $\{u,v\} \in E$. By the definition of $\hat{k}$, we have $\hat{k} - \hat{c}_v \geq 1$, for all $v \in V$. For every edge $\{u,v\}$, we have $\hat{c}_u > \hat{c}_v$ or $\hat{c}_v > \hat{c}_u$ because $\hat{c}$ is a valid colouring. If $\hat{c}_u > \hat{c}_v$, then $\hat{s}_{u,v} = 0$, so $\hat{c}_u - \hat{c}_v + n\hat{s}_{u,v} \geq 1$ and $\hat{c}_v - \hat{c}_u - n\hat{s}_{u,v} \geq 1 - n$ (because $\hat{c}_u, \hat{c}_v \in [n-1]_0$). If $\hat{c}_v > \hat{c}_u$, then $\hat{c}_u - \hat{c}_v \geq 1 - n$ and $\hat{s}_{u,v} = 1$, so $\hat{c}_u - \hat{c}_v + n\hat{s}_{u,v} \geq 1$, and we also have $\hat{c}_v - \hat{c}_u \geq 1$, so $\hat{c}_v - \hat{c}_u - n\hat{s}_{u,v} \geq 1 - n$.

Conversely, assume that $(\hat{c}, \hat{s}, \hat{k})$ is a feasible solution of (S3). Then $\hat{c}$ is a valid colouring. To prove this, all we need to prove is that $\hat{c}_u \neq \hat{c}_v$, for every edge $\{u,v\} \in V$. If $\hat{s}_{u,v} = 0$, then the constraint $c_u - c_v + n s_{u,v} \geq 1$ enforces that $\hat{c}_u > \hat{c}_v$. If $\hat{s}_{u,v} = 1$, then the constraint $c_v - c_u - n s_{u,v} \geq 1 - n$ enforces that $\hat{c}_v > \hat{c}_u$. In both cases, $\hat{c}_u \neq \hat{c}_v$, so $\hat{c}$ is a valid colouring of G.

2C

Let $\hat{c}$ be an optimal colouring of G (one that uses the minimum number of colours), and let $(\tilde{c}, \tilde{s}, \tilde{k})$ be an optimal solution of (S3). As already observed, we can assume that $\hat{c}_v \in [\ell - 1]_0$ if $\hat{c}$ uses ℓ colours. As shown in (b), there exists a feasible solution $(\hat{c}, \hat{s}, \hat{k})$ of (S3) corresponding to $\hat{c}$, and $\tilde{c}$ is a valid colouring of G. Moreover, given that $\hat{c}_v \in [\ell - 1]_0$, we can choose $\hat{k} = \ell$.

Since all vertices $v \in G$ satisfy $\tilde{c}_v \in [\tilde{k} - 1]_0$, due to the first constraint of (S3), $\tilde{c}$ is a colouring of G with at most $\tilde{k}$ colours. Since $\hat{c}$ is an optimal colouring, this shows that $\hat{k} = \ell \leq \tilde{k}$.

Conversely, $\left(\hat{c}, \hat{s}, \hat{k}\right)$ has objective function value $\hat{k}$ in (S3). Since $\left(\tilde{c}, \tilde{s}, \tilde{k}\right)$ is an optimal solution of (S3), this implies that $\hat{k} \geq \tilde{k}$.

Together, these two inequalities show that $\ell = \hat{k} = \tilde{k}$, that is, both $\hat{c}$ and $\tilde{c}$ are optimal colourings of G, and both $\left(\hat{c}, \hat{s}, \hat{k}\right)$ and $\left(\tilde{c}, \tilde{s}, \tilde{k}\right)$ are optimal solutions of (S3). This proves that an optimal colouring of G gives an optimal solution of (S3), and vice versa.