1. Computability, Complexity and Algorithms

(a) (5 points) Let $G(V, E)$ be an undirected graph. A Hamilton Path in G is a path of legth $(|V| - 1)$, while a Hamilton Cycle in G is a cycle of length $|V|$. HC is the problem of deciding if a graph has a Hamilton Cycle, while HP is the problem of deciding if a graph has a Hamilton Path. We know that HC is NP-complete. Show that HP is NP-complete.
(b) (5 points) Let G be a connected undirected graph with at least 3 vertices. Let G^3 be the graph obtained by connecting all pairs of vertices that are connected by a path in G of length at most 3. Show that for all graphs G^3 as above, HC is in P. (Hint: Show that G^3 always has a Hamilton Cycle.)

Solution: (a) Let $G(V, E)$ be a graph. For every edge $e = \{u, v\} \in E$ construct a graph $G_e(V_e, E_e)$, where $V_e = V \cup \{x, y\}$ and $E_e = E \cup \{\{u, x\}, \{v, y\}\}$. Clearly, G has a Hamilton Cycle that includes e if and only if G_e has a Hamilton Path. We may thus decide if G has a Hamilton Cycle using at most $|E|$ calls to the polynomial time algorithm that decides Hamilton Path. If any of this calls returns YES, then the answer to Hamilton Cycle is YES. If all of them return NO, then the answer to Hamilton Cycle is NO.
(b) To prove this, it suffices to consider only trees; any Hamilton cycle on the cube of a spanning tree is also a Hamilton cycle on the cube of the original graph. We now prove the following theorem, which obviously implies that any cube of a tree is Hamiltonian.
Theorem: *Let $T = (V, E)$ be a tree. For any edge $e \in E$, there is a Hamilton cycle on T^3 that contains edge e.*
Proof. By induction on $n = |V|$. The cases $n = 3$ and $n = 4$ are trivial, as the cube of the tree is always a clique.
Now assume that the claim is true for all trees with at most $n \geq 4$ vertices and let $T = (V, E)$ be a tree with $(n+1)$ vertices. Let $e = \{u, v\} \in E$ be an arbitrary edge. We show how to construct a cycle in T^3 that uses edge e. Since T is a tree, removing edge e will break T into two connected components, both of which are trees. Let T_u be the component that contains u and let T_v be the component that contains v. We may assume, without loss of generality, that T_v has at least 3 vertices. Let u' be the neighbor of u in T_u. By the inductive hypothesis, there is a Hamilton cycle H in T_u^3 that uses edges $\{u, u'\}$. There are now three cases to consider:
Case 1: T_v contains only one vertex. We get the desired Hamilton cycle by taking H and adding v between u and u'; this can be done since $d_T(u', v) = 1 + d_T(u, v) = 2$.
Case 2: T_v contains two vertices. One of these is v and we denote the other by v'. Obviously, v' is adjacent to v. We change the cycle $H = (\ldots, u', u, \ldots)$ to $H(\ldots, u', v', v, u, \ldots)$, which yields the desired result. The modification can be done, since $d_T(u', v') = 3$.
Case 3: T_v contains at least three vertices. Let v' be a neighbor of v in T_v. By the inductive hypothesis, there is a Hamilton cycle H' in T_v that contains edge $\{v, v'\}$. We have that $d_T(u', v') = 3$, so we can constrcuct a Hamilton cycle in T by removing edge $\{u, u'\}$ from H and $\{v, v'\}$ from H', and joining these cycles together by adding edges $\{u', v'\}$ and $\{u, v\}$.
Thus, in all cases, we can construct a Hamilton cycle in T^3 that contains edge $\{u, v\}$.

2. Analysis of Algorithms

Consider the following algorithm for the weighted vertex cover problem, where $w(v)$ is the weight of vertex v. Initially $t(v) := w(v)$ for all vertices. When $t(v)$ drops to 0, v is picked in the cover. $c(e)$ is the amount that we charge an edge e. In particular:

1. Initialization
$C := \emptyset$
$\forall v \in V$, $t(v) := w(v)$
$\forall e \in E$, $c(e) := 0$
2. While C is not a vertex cover do:
    Pick uncovered edge, say $\{u, v\}$
    Let $m := \min\{t(u), t(v)\}$
     $t(u) := t(u) - m$
     $t(v) := t(v) - m$
     $c(u, v) := m$
     Include in C all vertices v that have $t(v) = 0$
3. Output C

Argue that this is a factor 2 approximation algorithm for weighted vertex cover.

Solution: The (IP) formulation of weighted vertex cover, its (LP)-relaxation and the dual (DP) are:

(IP)	(LP)	(DP)
$\min \sum_{v \in V} x_v w(v)$	$\min \sum_{v \in V} x_v w(v)$	$\max_{e \in E} y_e$
$x_v + x_u \geq 1\ ,\ \forall \{u, v\} \in E$	$x_v + x_u \geq 1\ ,\ \forall \{u, v\} \in E$	$\sum_{u:\{u,v\} \in E} y_{\{u,v\}} \leq w(v)\ ,\ \forall v \in V$
$x_v \in \{0, 1\}\ ,\ \forall v \in V$	$x_v \geq 0\ ,\ \forall v \in V$	$y_e \geq 0\ ,\ \forall e \in E$

The algorithm builds integral solution to (LP), thus also solution to (IP), and solution to (DP). The complementary slackness conditions hold tight for primal variables, and with factor 2 for dual variables. In particular, we have

$$x_v > 0 \Rightarrow \sum_{u:\{u,v\} \in E} y_{\{u,v\}} = w(v)\ ,\ \forall v \in V$$

$$y_{\{u,v\}} > 0 \Rightarrow x_v + x_u \leq 2\ ,\ \forall \{u, v\} \in E$$

We thus have:

$$\begin{aligned}
\sum_{v \in V} x_v w_v &= \sum_{v \in V} \sum_{u:\{u,v\} \in E} y_{\{u,v\}} && \text{, by complementary slackness} \\
&\leq 2 \sum_{\{u,v\} \in E} y_{\{u,v\}} && \text{, by complementary slackness} \\
&\leq 2OPT(\text{DP}) = 2OPT(\text{LP}) && \text{, by duality theory} \\
&\leq 2OPT(\text{IP})
\end{aligned}$$

3. Theory of Linear Inequalities

Let $P = \{x \in \mathbb{R}^n \mid Ax \leq b\} \subseteq [0,1]^n$ with $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$ be a polytope contained in the 0/1 cube; in particular the bound inequalities $0 \leq x \leq 1$ are valid for P.

For $i \in [n]$ we consider the following procedure:

1. Generate the nonlinear system $(b - Ax)x_i \geq 0$, $(b - Ax)(1 - x_i) \geq 0$.
2. Relinearize the system by replacing $x_j x_i$ with y_j whenever $i \neq j$ and x_j whenever $i = j$. We obtain a new, higher dimensional polyhedron M_i.
3. Define $P_i := \text{proj}_x M_i$.

Finally define $P^1 := \bigcap_{i \in [n]} P_i$. This polyhedron is a strengthening of the original formulation of P.

Prove the following:

$$P^1 = \bigcap_{i \in [n]} \text{conv}\left((P \cap \{x \mid x_i = 0\}) \cup (P \cap \{x \mid x_i = 1\})\right)$$

Solution. It suffices to verify the claim separately for each P_i with $i \in [n]$. Let

$$\begin{aligned}\bar{P} &:= \text{conv}((P \cap \{x \mid x_i = 0\}) \cup (P \cap \{x \mid x_i = 1\})) \\ &= \text{conv}((\{x \mid Ax \leq b\} \cap \{x \mid x_i = 0\}) \cup (\{x \mid Ax \leq b\} \cap \{x \mid x_i = 1\})).\end{aligned}$$

We will first show that $P_i \subseteq \bar{P}$. We assume that $P \neq \emptyset$ as otherwise there is nothing to show. Let $cx \leq \delta$ be valid for $P \cap \{x \mid x_i = 0\}$ and $P \cap \{x \mid x_i = 1\}$. Then by Farkas' Lemma there exists $\tau, \lambda \geq 0$ so that

$$cx + \lambda x_i \leq \delta$$

and

$$cx + \tau(1 - x_i) \leq \delta$$

are valid for P. Hence

$$(\delta - cx + \lambda x_i)(1 - x_i) \geq 0$$

and

$$(\delta - cx + \tau(1 - x_i))x_i \geq 0$$

are valid for the nonlinear system. Adding up both inequalities we obtain

$$\delta - cx + (\tau + \lambda)(1 - x_i)x_i \geq 0$$

and using $(1 - x_i)x_i = 0$ we have $cx \leq \delta$ is valid for M_i and hence also for P_i.

It remains to show that $\bar{P} \subseteq P_i$. Suppose that $\bar{P} \neq \emptyset$; otherwise there is nothing to show. Let $\bar{x} \in P \cap \{x \mid x_i = \ell\}$ for $\ell \in \{0, 1\}$. We can lift the point by defining $y_j = \bar{x}_j \bar{x}_i$ for $i \neq j$. Then $(\bar{x}, y) \in M_i$ as $\bar{x}_i^2 = \bar{x}_i$.

4. Combinatorial Optimization

Consider the following generalization of matroids. Given a ground set E and a family $\mathcal{F}$ of subsets of E we say that $(E,\mathcal{F})$ is *k-extensible* if the following hold: (i) if $A \in \mathcal{F}$ and $B \subseteq A$, then $B \in \mathcal{F}$; (ii) consider $A, B \in \mathcal{F}$ with $A \subseteq B$; if $e \in E$ is such that $A + e \in \mathcal{F}$, then there is a set $K \subseteq B - A$ of size at most k such that $B - K + e$ belongs to $\mathcal{F}$.

Given $(E,\mathcal{F})$ k-extensible and a weight function $w : E \to \mathbb{R}$ (extended to sets as usual by $w(A) = \sum_{e\in A} w(e)$), consider the greedy algorithm for $\max_{S\in\mathcal{F}} w(S)$: (0) Start with $S = \emptyset$; (1) pick an element $e \in E - S$ with largest weight that satisfies $S + e \in \mathcal{F}$, and update $S \leftarrow S + e$ (if no such element exists, stop); (2) Repeat the previous step.

1. Let e_i be the element chosen by the greedy algorithm in step i, and let S_i be the set obtained at the end of step i (so $S_i = e_1 + \ldots + e_i$). Given any set $A \in \mathcal{F}$, let $OPT(A) = \max\{w(B) : B \supseteq A, B \in \mathcal{F}\}$ (i.e. the best extension of A in $\mathcal{F}$). Show that for all i
$$w(OPT(S_i)) \geq w(OPT(S_{i-1})) - k \cdot w(e_i).$$

2. Show that the last set S_ℓ computed by the greedy algorithm satisfies $w(S_\ell) \geq \frac{1}{k+1} w^*$, where $w^* = \max\{w(A) : A \in \mathcal{F}\}$.

3. Given a graph $G = (V,E)$ and $b \in \mathbb{R}^V_+$, a set $S \subseteq E$ is a *b-matching* if S has at most b_v edges incidents to vertex v, for all $v \in V$. Give a polytime algorithm for finding a b-matching with at least $1/3$ as many edges as the largest b-matching. (No need to analyze the running-time of the algorithm.)

Solution.

1. Notice $OPT(S_{i-1})$ is an extension of S_{i-1}. Since $S_{i-1} + e_i \in \mathcal{F}$, k-extensibility implies that there is a set $K \subseteq OPT(S_{i-1}) - S_{i-1}$ of size at most k such that $OPT(S_{i-1}) - K + e_i$ belongs to $\mathcal{F}$. Since the set $OPT(S_{i-1}) - K + e_i$ is an extension of $S_{i-1} + e_i = S_i$, we have
$$w(OPT(S_i)) \geq w(OPT(S_{i-1}) - K + e_i) = w(OPT(S_{i-1})) - w(K) + w(e_i).$$
Now notice that for every $e \in K$, we have $S_{i-1} + e \in \mathcal{F}$: this is because $OPT(S_{i-1}) \in \mathcal{F}$ contains $S_{i-1} + e$ and because of Property 1 in the definition of k-extensible systems. Thus, by definition of e_i each element of K has weight at most $w(e_i)$ and hence $w(K) \leq k \cdot w(e_i)$. Plugging this in the last displayed inequality concludes the proof.

2. Notice that the last set S_ℓ is a maximal set in $\mathcal{F}$, and hence $OPT(S_\ell) = S_\ell$; also, $w^* = w(OPT(\emptyset))$. Then applying the result from the previous question repeatedly we get
$$w(OPT(S_\ell)) \geq w(OPT(\emptyset)) - k \cdot \sum_{i=1}^{\ell} w(e_i) = w(OPT(\emptyset)) - k \cdot w(S_\ell),$$
or equivalently $w(S_\ell) \geq w^* - k \cdot w(S_\ell)$. Reorganizing gives the result.

3. Let $\mathcal{F}$ be the set of b-matchings; we show that $(E, \mathcal{F})$ is 2-extensible. Property 1 of 2-extensibility clearly holds. For Property 2, consider b-matchings S, S' with $S \subseteq S'$, and an edge (u, v) such that $S + (u, v)$ is a b-matching. We claim that there is a set $K \subseteq S' - S$ of size at most 2 such that $S' - K + (u, v)$ is a b-matching: simply let K consist of one edge of $S' - S$ incident to u (if exists) and one edge of $S' - S$ incident to v (if exists). Then $S' - K$ has at most $b_u - 1$ edges incident to u and at most $b_v - 1$ edges incident to v, hence $S' - K + (u, v)$ is a b-matching.

   Therefore, the greedy algorithm described above, applied to this system $(E, \mathcal{F})$ gives the desired approximation (and it is easy to check that it can be made to run in polynomial time).

5. Graph Theory

Let G be a simple plane graph of minimum degree at least three. Prove that G has either a vertex of degree three incident with a face of size at most five, or a face of size three incident with a vertex of degree at most five.

Solution: Suppose for a contradiction that G is a simple plane graph of minimum degree at least three containing neither of the two required configurations. We use the discharging method to obtain a contradiction. A vertex of degree d will be assigned a charge of $d-4$ and a face of size l will be assigned a charge of $l-4$. Then the sum of the charges is -8 by Euler's formula. We now redistribute the charges as follows: a vertex of degree three will send a charge of $-1/3$ to every incident face and every face of size three will send a charge of $-1/3$ to every incident vertex. This results in a new distribution of the charges with the same sum. We claim that the new charge of every vertex and every face is nonnegative. That is a contradiction, because the total sum is negative. To see that every vertex and face has nonnegative new charge let v be a vertex of G. If v has degree three, then its initial charge is -1, but it sends $-1/3$ to every incident face, and hence ends up with charge zero. If v has degree four or five, then its initial charge is nonnegative and it does not receive any charge; hence its final charge is nonnegative. Let us now assume that v has degree $d \geq 6$. Its initial charge is $d-4$ and it receives at least $-d/3$ from its incident faces for total of at least $d - 4 - d/3 \geq 0$, as desired. The analysis of face charges is analogous.

6. Probabilistic methods

Let H denote the graph on five vertices consisting of a copy of K_4 along with an additional edge attached to one of the 4 vertices. (Recall that K_4 is the complete graph on 4 vertices.) Let $G_{n,p}$ denote the usual Erdős-Rényi random graph on n vertices with the edge probability $p = p(n)$. Then

(i) Show that

$$\Pr(G_{n,p} \text{ contains a copy of } H) \to 0 \text{ if } p \ll n^{-2/3};$$

(ii) Let $G_1 = G_{n,p/2}$. Show that

$$\Pr(G_1 \text{ contains a copy of } K_4) \to 1 \text{ if } p \gg n^{-2/3}.$$

(iii) Let G be the union of two independent copies of $G_{n,p/2}$. (i.e., on the same set of vertices, but the edge set is taken as the union of edge sets.) Prove that

$$\Pr(G \text{ contains a copy of } H) \ \to 1 \text{ if } p \gg n^{-2/3}\,.$$

(iv) Conclude that

$$\Pr(G_{n,p} \text{ contains a copy of } H) \ \to 1 \text{ if } p \gg n^{-2/3}\,.$$

(In the above, we are using the fairly standard notation: $p(n) \ll f(n)$ means that $p(n)/f(n) \to 0$, as $n \to \infty$, and similarly, $p(n) \gg f(n)$ means that $p(n)/f(n) \to \infty$, as $n \to \infty$.)

Solution: (i) Recall that $n^{2/3}$ is the threshold probability for having a copy of K_4. Clearly there can not be an H, if there are not K_4's in the graph. Now, using the first moment method – the expected number of K_4's being $\binom{n}{4}p^6$ – the probability that there is a K_4 goes to 0, for $p \ll n^{-2/3}$.

(ii) Follows from a straightforward 2nd moment argument by computing $\mathrm{Var}(X)$, for $X = \#$ of copies of K_4, and observing that $\mathrm{Var}(X) = o([E(X)]^2)$.

(iii) Let $G = G_1 \cup G_2$, where G_1 and G_2 are independent copies of $G_{n,p/2}$. By Part(ii), G_1 contains a K_4 w.h.p. So it suffices to show that the probability that G_2 contains an edge with *exactly one endpoint* in this K_4 tends to 1. This probability is at least $1-(1-p/2)^{4(n-4)} \approx 1-e^{-2pn} \to 1$.

(iv) The graph G in Part(iii) can be seen as distributed as $G_{n,p'}$, where $p' = p-(p/2)^2$, since that is the probability of having an edge between a pair of vertices, independent of everything else. Since $p' < p$, Part (iv) follows from Part (iii).

7. Algebra

Let A be $n \times n$ matrix with rational entries and let $p \in \mathbb{N}$ be a prime. Show that A cannot satisfy

$$A^{n+1} - pA^n = pI,$$

where I is the identity matrix.

Solution: Suppose not. Let $m(x)$ be the minimal polynomial of A. Note that $m(x)$ has coefficients in $\mathbb{Q}$. The matrix A satisfies $A^{n+1} - pA^n - pI = 0$, and therefore $m(x)$ divides $x^{n+1} - px^n - p$. The polynomial $x^{n+1} - px^n - p$ is irreducible over $\mathbb{Q}$ by Eisenstein's criterion. Therefore $m(x) = x^{n+1} - px^n - p$. However the degree of $m(x)$ is at most n since $m(x)$ divides the characteristic polynomial of A. This is a contradiction.

7. Linear Algebra

Let A be a $n \times n$ matrix and v a vector in $\mathbb{R}^n$ such that the set $\{v,\, Av,\, A^2v, \ldots,\, A^{n-1}v\}$ is linearly independent. Show that any matrix B that commutes with A can be written as a polynomial in A.

Solution: Observe that since $\{v,\, Av,\, A^2v, \ldots,\, A^{n-1}v\}$ are n linearly independent vectors in $\mathbb{R}^n$ they form a basis.

Thus the vector Bv can be written uniquely as

$$Bv = \gamma_0 v + \gamma_1 Av + \cdots + \gamma_{n-1} A^{n-1} v$$

Set

$$C = \gamma_0 \mathrm{Id} + \gamma_1 A + \cdots + \gamma_{n-1} A^{n-1}$$

Since A commute with B we have

$$BAv = ABv = \gamma_0 v + \gamma_1 AAv + \cdots + \gamma_{n-1} A^{n-1} Av = CAv$$

Similarly we get that for every m

$$BA^m v = CA^m v$$

This implies that B and C coincide on a basis and thus $B = C$.