MATH 642:582–Fall, 2003

Assignment 4–Due November 4 (Version: November 2, 2003)

1. (a) Let $\mathcal{H}$ be a hypergraph on X. A $\mathcal{H}\times X$ matrix M is *compatible* with $\mathcal{H}$ if $M(E,x) = 0$ whenever $x \notin E$. Prove that $\mathcal{H}$ has a system of distinct representatives (SDR) if and only if there is a nonnegative matrix compatible with $\mathcal{H}$ whose row sums are each at least 1, and whose column sums are each at most 1.

   (b) Let $\mathcal{H}$ be a hypergraph such that for every edge E and $x \in E$, $|E| \geq deg(x)$. Prove that $\mathcal{H}$ has an SDR.

2. For a hypergraph $\mathcal{H}$ recall that $\nu(\mathcal{H})$ is the size of the largest matching. Let $\mu(\mathcal{H})$ be the average over all (not necessarily maximal) matchings M of the size (number of edges) of M. Let $r(\mathcal{H})$ denote the size of the largest edge of $\mathcal{H}$. Prove that $\nu(\mathcal{H}) \leq \mu(\mathcal{H})\max\{2, r(\mathcal{H})\}$. (Hint: use induction on the number of edges.)

3. Recall that for a hypergraph $\mathcal{H}$, the vertex cover number $\tau(\mathcal{H})$ is the size of the smallest set of vertices that has nonempty intersection with every edge of $\mathcal{H}$. A hypergraph $\mathcal{H}$ is τ-critical if for any edge E of $\mathcal{H}$, $\tau(\mathcal{H}-\{E\}) < \tau(\mathcal{H})$. A *fractional vertex cover* of $\mathcal{H}$ is a function c on the vertices with the property that for each edge E, $\sum_{x\in E} c(x) \geq 1$. The *weight* of c is $\sum_{x\in V(\mathcal{H})} c(x)$. The *fractional cover number* $\tau^*(\mathcal{H})$ of $\mathcal{H}$ is the minimum weight of a fractional vertex cover.

   Let $\mathcal{H}$ be a τ-critical hypergraph. Prove that either $\mathcal{H}$ is a matching or $\tau^*(\mathcal{H}) < \tau(\mathcal{H})$.

4. The purpose of this problem is to prove a generalization of the Hall theorem (which, in particular, gives a different proof of Hall's theorem). A hypergraph $\mathcal{G}$ is a *shrinking* of a hypergraph $\mathcal{H}$ if there is a 1-1 function $f : \mathcal{H} \longrightarrow \mathcal{G}$ such that $f(E) \subseteq E$ for all $E \in \mathcal{H}$. (In this definition, we use the common abuse of notation that $\mathcal{H}$ and $\mathcal{G}$ are each identified with their set of edges. Intuitively, $\mathcal{G}$ can be obtained by individually "shrinking" (removing vertices from) some or all of the edges of $\mathcal{H}$.)

   Define the *surplus* of a hypergraph $\mathcal{H}$ to be $|\cup\mathcal{H}| - |\mathcal{H}|$ (where $\cup\mathcal{H}$ is the union of the edges of $\mathcal{H}$.) The surplus may, of course, be positive or negative. The *global surplus* of $\mathcal{H}$ is the minimum of the surplus of $\mathcal{F}$ over all nonempty $\mathcal{F} \subseteq \mathcal{H}$.

   Theorem. Let $\mathcal{H}$ be a hypergraph with global surplus $s \geq 0$. Then $\mathcal{H}$ has a shrinking $\mathcal{G}$ that is $s+1$-uniform (all edges have size $s+1$) and global surplus s.

   (Not to hand in: (i) Why is this a generalization of Hall's theorem? (ii) In the case that $s = 1$ the conclusion is equivalent to: $\mathcal{H}$ can be shrunk to a graph (all edges of size 2) that has no cycles,i.e., is a forest.)

   (a) If $\mathcal{F} \subseteq \mathcal{H}$ is nonempty and the surplus of $\mathcal{F}$ is equal to the global surplus of $\mathcal{H}$ we say that $\mathcal{F}$ is *surplus critical* for $\mathcal{H}$. Prove that if $\mathcal{F}$ and $\mathcal{G}$ are surplus critical for $\mathcal{H}$ then $\mathcal{F}\cap\mathcal{G}$ is either empty or is also surplus critical for $\mathcal{H}$.

   (b) An edge E of a hypergraph $\mathcal{H}$ is *unshrinkable* if replacing E by a proper subset of E reduces the global surplus of $\mathcal{H}$. Prove that if E is unshrinkable then there is a surplus critical $\mathcal{F} \subseteq \mathcal{H}$ that contains E, and has the property that any other member of $\mathcal{F}$ is disjoint from E.

(c) Prove the theorem.

5. The standard algorithm for checking whether a hypergraph has an SDR (or, equivalently, checking whether a bipartite graph has a perfect matching) involves searching for augmenting paths, as discussed in section 5.4 of the book. In this problem and the next, you will prove the correctness of two completely different (and rather surprising) algorithms.

   For purposes of this problem, if $\bar{v}$ is a nonnegative nonzero real vector, *normalizing* $\bar{v}$ means to divide each entry of $\bar{v}$ by $\sum_i v_i$. Also, we say that a vector is *stochastic* if it is nonnegative and its sum is 1, and ϵ-stochastic if it is nonnegative and its sum is strictly between $1+\epsilon$ and $1/(1+\epsilon)$.

   Let $\mathcal{H}$ with n vertices and n edges, and let M be the incidence matrix. Let M_0 be the matrix obtained by normalizing each column of M. Consider the following algorithm.

   For i from 1 to $10n^3 \log_2 n$ do: Let N_i be the matrix obtained by normalizing each row of M_{i-1}. Let M_i be the matrix obtained by normalizing each column of N_i.

   The purpose of this problem is to prove the following remarkable result: **Theorem.** Let $M_f(\mathcal{H})$ be the final M_i produced by the algorithm. Then $\mathcal{H}$ has an SDR if and only if M_f is $1/n$-stochastic.

   (a) Prove that if there is a matrix compatible with $\mathcal{H}$ whose columns are stochastic and whose rows are $1/n$-stochastic, then $\mathcal{H}$ has an SDR.

   (b) For each $i \geq 0$, prove that $per(M_{i+1}) \geq per(M_i)$. Furthermore prove that if $\epsilon \in [0,1]$ and M_i is not ϵ-stochastic then $per(M_{i+1}) \geq (1+\frac{\epsilon^2}{8})per(M_i)$. (Here $per(M)$ denotes the permanent of M).

   (c) Prove the Theorem.

6. If X is a set and w is a real valued function on X, we extend the domain of w to subsets of X by defining $w(Y) = \sum_{y \in Y} w(y)$. The quantity $w(Y)$ is called the *weight* of the subset Y with respect to w.

   Let $\mathcal{H}$ be a hypergraph on X. If w is a weight function on X, let $w_{\min}(\mathcal{H})$ be the minimum weight of any edge of $\mathcal{H}$. (This is defined to be infinite for the hypergraph having no edges). We say that w *is uniquely minimized over* $\mathcal{H}$ if there is exactly one edge of weight $w_{\min}(\mathcal{H})$.

   Theorem Let $\mathcal{H}$ be a hypergraph on X with $|X| = n$. Let Ω be the set of all functions from X to $[2n]$. If w is chosen uniformly at random from Ω, then with probability at least 1/2, w is uniquely minimized over $\mathcal{H}$.

   The first two parts of this problem are devoted to the proof of this theorem. The remaining three parts (which can be done independently of the first two parts) apply it to develop yet another neat algorithm for testing whether a bipartite graph has a perfect matching.

   (a) For a vertex x of $\mathcal{H}$ the star $\mathcal{H}$ over x,$\mathcal{H}[x]$ is the partial hypergraph with edges set $\{E \in \mathcal{H} : x \in E\}$. The link of $\mathcal{H}$ over x, $\mathcal{H}\langle x \rangle$, is the hypergraph on $X - x$

with edge set $\{E - x : E \in \mathcal{H}[x]\}$, and $\mathcal{H}_{X-x}$ is the hypergraph on $X - x$ whose edge set is $\{E \in \mathcal{H} : x \notin E\}$.

Prove that for a hypergraph $\mathcal{H}$ and a vertex weighting w the following are equivalent:

i. $\mathcal{H}$ uniquely minimizes w

ii. For every vertex x,
$w_{\min}(\mathcal{H}_{X-x}) - w_{\min}(\mathcal{H}\langle x\rangle) \neq w(x)$.

(b) Prove the theorem.

(c) Let $G = \{X, Y; E\}$ be a bipartite graph with $|X| = |Y| = k$. Let $\mathcal{M}(G)$ be the set of all real valued matrices M indexed by $X \times Y$ that satisfy $M(x, y) = 0$ if $(x, y) \notin E$. Prove that G has a perfect matching if and only if there is a matrix $M \in \mathcal{M}$ such that $det(M) \neq 0$.

(d) Let G be as in the previous part. Construct a matrix $M \in \mathcal{M}(G)$ randomly as follows: for each $(x, y) \in E$, select $j(x, y) \in [2k^2]$ at random and set $M(x, y) = 2^{j(x,y)}$. Prove that if G has a perfect matching then the probability that $det(M) = 0$ is less than $1/2$.

(e) Use the results of the previous problems to design a *probabilistic test* for whether a bipartite graph has a matching. Such a test should have the following property: given a bipartite graph G and a number $\epsilon > 0$ answers "Yes" or "No" such that: (1) if G has a perfect matching then the test yields "Yes" with probability at least $1 - \epsilon$ and (2) if G has no perfect matching then the test always says "No".