Stat 155 Game theory, Yuval Peres **Fall 2004**

Lectures 20,21,22

The game of Green Hackenbush: In the game of Green Hackenbush, we are given a finite graph, that consists of vertices and some undirected edges between some pairs of the vertices. One of the vertices is called the root, and might be thought of as the ground on which the rest of the structure is standing. We talk of 'green' Hackenbush because there is an partisan variant of the game in which edges may be coloured red or blue instead.

The aim of the players I and II is to remove the last edge from the graph. At any given turn, a player may remove some edge from the graph. This causes not only that edge to disappear, but also all those edges for which every path to the root travels through the edge the player removes.

Note firstly that, if the original graph consists of a finite number of paths, each of which ends at the root, then, in this case, Green Hackenbush is equivalent to the game of nim, where the number of piles is equal to the number of paths, and the number of chips in a pile is equal to the length of the corresponding path.

We need a lemma to handle the case where the graph is a tree:

Lemma 1 (Colon Principle) *The Srague-Grundy function of Green Hackenbush on a tree is unaffected by the following operation: for any example of two branches of the tree meeting at a vertex, we may replace these two branches by a path emanating from the vertex whose length is the nim-sum of the Sprague-Grundy functions of the two branches.*

Proof: See Ferguson, $I - 42$. The proof in outline: if the two branches consist simply of paths (or 'stalks') emanating from a given vertex, then the result is true, by noting that the two branches form a two-pile game of nim, and using the direct sum Theorem for the Sprague-Grundy functions of two games. More generally, we show that we may perform the replacement operation on any two branches meeting at a vertex, by iterating replacing pairs of stalks meeting inside a given branch, until each of the two branches itself has become a stalk.

As a simple illustration, see the figure. The two branches in this case are stalks, of length 2 and 3. The Sprague-Grundy values of these stalks equal 2 and 3, and their nim-sim is equal to 1. Hence, the replacement operation takes the form shown.

For further discussion of Hackenbush, and references about the game, see Ferguson, part I, section 6.

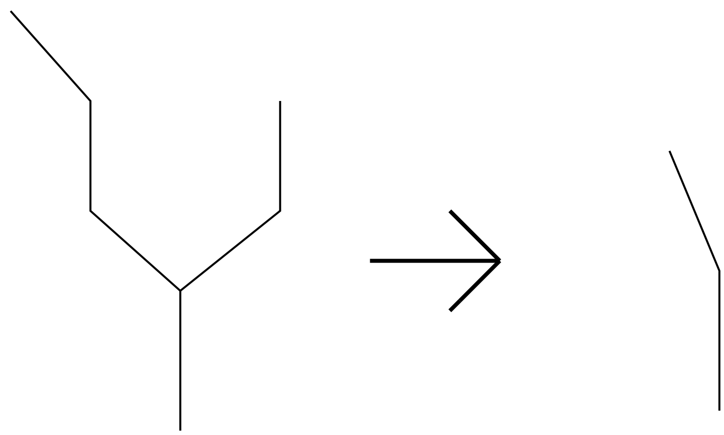

General hide-and-seek games: We now analyse a more general version of the game of hide-and-seek. A matrix of values $(b_{ij})_{n\times n}$ is given. Player II chooses a location (i,j) at which to hide. Player I chooses a row or a column of the matrix. He wins a payment of b_{ij} if the line he has chosen contains the hiding place of his opponent.

Firstly, we propose a strategy for player II, later checking that it is optimal. The player may choose a fixed permutation π of the set $\{1,\ldots,n\}$ and then hide at location (i,π_i) with a probability p_i that he chooses. Given a choice π, the optimal choice for p_i is $p_i = d_{i,\pi_i}/D_\pi$, where $d_{ij} = b_{ij}^{-1}$ and $D_\pi = \sum_{i=1}^n d_{i,\pi_i}$, because it is this choice that equalizes the expected payments. The expected payoff for the game is then $1/D_\pi$.

Thus, if Player II is going to use a strategy that consists of picking a permutation π^* and then doing as described, the right permutation to pick is one that maximizes D_π. We will in fact show that doing this is an optimal strategy, not just in the restricted class of those involving permutations in this way, but over all possible strategies.

To find an optimal strategy for Player I, we need an analogue of Konig's lemma. In this context, a *covering* of the matrix $D = (d_{ij})_{n\times n}$ will be a pair of vectors $u = (u_1,\ldots,u_n)$ and $w = (w_1,\ldots,w_n)$ such that $u_i + w_j \geq d_{ij}$ for each pair (i,j). (We assume that u and v have non-negative components). The analogue of the Konig lemma:

Lemma 2 *Consider a minimal covering (u^*,v^*). (This means one for which $\sum_{i=1}^n (u_i+w_i)$ is minimal). Then:*

$$\sum_{i=1}^n (u_i + w_i) = \max_\pi D_\pi. \tag{1}$$

Proof: Note firstly that a minimal covering exists, because the map

$$(u, v) \to \sum_{i=1}^{n} (u_i + v_i),$$

and attains its infimum, if at all, on the closed and bounded set $\{(u, w) : u_i, w_i \leq nM\}$, where $M = \max_{k,l} D_{k,l}$.

Note also that we may assume that $\min_i u_i^* > 0$.

That the left-hand-side of (1) is at least the right-hand-side is straightforward. Indeed, for any π, we have that $u_i^* + w_{\pi_i}^* \geq d_{i,\pi_i}$. Summing over i, we obtain this inequality.

Showing the other inequality is harder, and requires Hall's marriage lemma, or something similar. We need a definition of 'knowing' to use the Hall lemma. We say that row i knows column j is

$$u_i^* + w_j^* = d_{ij}.$$

Let's check Hall's condition. Suppose that k rows $i_1, \ldots, i_k$ know between them only $l < k$ columns $j_1, \ldots, j_l$. Define $\tilde{u}$ from u^* by reducing these rows by a small amount ϵ. Leave the other rows unchanged. The condition that ϵ must satisfy is in fact that

$$\epsilon < \min_i u_i^*$$

and also

$$\epsilon < \min \{u_i + w_j - d_{ij} : (i, j) \, \text{such that} \; u_i + w_j > d_{ij}\}.$$

Similarly, define $\tilde{w}$ from w^* by adding ϵ to the l columns known by the k rows. Leave the other columns unchanged. That is, for the columns that are changing,

$$\tilde{w}_{j_i} = w_{j_i}^* + \epsilon \text{ for } i \in \{1, \ldots, l\}.$$

We claim that $(\tilde{u}, \tilde{w})$ is a covering of the matrix. At places where the equality $d_{ij} = u_i^* + w_j^*$ holds, we have that $d_{ij} = \tilde{u}_i + \tilde{w}_j$, by construction. In places where $d_{ij} < u_i^* + w_j^*$, then

$$\tilde{u}_i + \tilde{w}_j \geq u_i^* - \epsilon + w_j^* > d_{ij},$$

the latter inequality by the assumption on the value of ϵ.

The covering $(\tilde{u}, \tilde{w})$ has a strictly smaller sum of components than does (u^*, w^*), contradicting the fact that this latter covering was chosen to be minimal.

We have checked that Halls' condition holds. Hall's lemma provides a mathcing of columns and rows. This is a permutation π^* such that, for each i, we have that

$$u_i^* + w_{pi_i}^* = d_{i,\pi(i)},$$

from which it follows that

$$\sum_{i=1}^{n} u_i^* + \sum_{i=1}^{n} w_i^* = D_{\pi^*}.$$

We have found a permutation π^* that gives the other inequality required to prove the lemma. □

We have therefore found a pair of optimal strategies for the players. To summarise, player I chooses row i with probability $u_i^*/D_{\pi^*}^*$, and column j with probability w_j^*/D_π^*. Against this strategy, if player II chooses (i, j), then the payoff will be

$$\frac{u_i^* + v_j^*}{D_{\pi^*}} b_{ij} \geq \frac{d_{ij} b_{ij}}{D_{\pi^*}} = D_{\pi^*}^{-1}.$$

We deduce that the permutation strategy for player II described earlier is indeed optimal.