CMSC 39600-1 Online Algorithms Autumn 2004 Teacher: Adam Kalai

Homework due 10/27. Please write the name of any collaborator you worked with.

These problems are based on questions asked in class.

1. We have a slot machine M with a known distribution over the sequence of its rewards. There is a fixed $\alpha \in (0, 1)$, and we would like to play so as to maximize our expected discounted reward,

$$\text{Expected discounted reward} = E[\sum_{t=0}^{\infty}(1-\alpha)^t r^t],$$

where r^t is the reward we receive at time t. Recall the definition of the Gittins index. Imagine considering two machines, M and C_x where C_x is a machine that always returns reward x. The Gittins index of a machine M is

$$G(M) = sup\{x | \text{we would rather play } M \text{ than } C_x\}.$$

That is, x is our indifference point between playing M and C_x.

(a) What is $G(C_x)$ in terms of x.

(b) Let $\frac{1}{2}C_x + \frac{1}{2}C_y$ be the machine which is either a machine that always pays off x or a machine that always pays off y, with equal probability. (This is different than a machine which, each time, independently randomly pays off x or y.) Say $0 < x < y < 1$. Write a closed-form expression for the Gittins index of this machine.

(c) Let B_x be a machine which returns reward $2x$ with probability 1/2 and reward 0 with probability 1/2, each period, independently. (This is different than $\frac{1}{2}C_{2x} + \frac{1}{2}C_0$.) Consider a machine $\frac{1}{2}B_x + \frac{1}{2}B_y$ which is either B_x or B_y with equal probability, for $0 < x < y < 1$. For example, this assigns probability 1/16 to the sequence $2x, 0, 2x$ and 1/16 to the sequence $2y, 0, 2y$, but 0 probability to the sequence $2x, 0, 2y$. Write a closed-form expression for the Gittins index of this machine. How does this compare to the answer of (b)?

2. Online gradient descent. Recall the gradient descent algorithm and guarantees of Zinkevich. Let $S \subset \mathbb{R}^n$ be a convex (compact) n-dimensional set of diameter D. Let $f^1, f^2, \ldots, f^T$ be a sequence of differentiable convex functions such that $\|\nabla f^t(x)\| \leq G$ for some G and all $x \in S$. Then the OGD algorithm is to take

$$x^{t+1} = \Pi_S \left(x^t - \frac{D}{G\sqrt{T}} \nabla f_t(x^t) \right).$$

Here $\Pi_S(x) = \arg\min_{x' \in S} \|x - x'\|$ finds the closest point in S to x.[1] In class, we saw that the total cost of such an online algorithm is near the best single point in hindsight:

$$\sum_{t=1}^{T} f^t(x^t) \leq \min_{x \in S} \left(\sum_{t=1}^{T} f^t(x) \right) + DG\sqrt{T}.$$

Someone asked after class, "is the final point x^{T+1} a good point to use in hindsight?" Show that it isn't by giving an example of a convex set S of diameter $D = 1$ and a sequence of functions of gradient at most $G = 1$, such that

$$\sum_{t=1}^{T} f^t(x^{T+1}) \geq \min_{x \in S} \left(\sum_{t=1}^{T} f^t(x) \right) + c_1 T - c_2.$$

The above should hold for some constants $c_1 > 0, c_2$ and for all T.

3. Hannan-style algorithms. Consider the following online min-cut problem. There are m nodes. Each period, we choose a cut of these nodes, i.e. a partition into two sets. Then an edge between two vertices (could be any two vertices) is revealed. We pay 1 if the edge crosses our cut (i.e., the vertices are in different partitions) and pay 0 if the edge does not cross the cut (i.e., the vertices are in the same partition). Describe a frequency count algorithm for this problem (you can assume the existence of a standard offline min-cut algorithm for an arbitrary weighted graph). Describe *clearly* to what you would add perturbations. What is the dimensionality n of this problem (how many frequencies must

[1] For example, in one dimension, if $S = [0, 1]$, then $\Pi_S(x) = 0$ if $x < 0$, x if $x \in [0, 1]$, and 1 if $x > 1$. The update is $x^{t+1} = \Pi_S \left(x^t - \frac{D}{G\sqrt{T}} f_t'(x^t) \right)$.

you keep track of)? What is the maximum cost each period? Suppose we are adding perturbations $p \in [0, \sqrt{T}]^n$ at the beginning of running PFC. Give the best upper bounds (to within constant factor) you can on the parameter P, where P is the maximum cost of a cut on the pretend perturbed period 0. (The other parameters B, L=1.)