Lecture 6

1 Congestion minimization in the experts framework

We will view the toll congestion game from the first lecture in the experts framework. The experts algorithm will output a flow that approximately minimizes the congestion, the flow is rounded off probabilistically to obtain approximately optimal paths with high probability.

While experts provide a very general framework for algorithm design, a number of details and tricks are involved in getting efficient algorithms for a specific problems.

1.1 Recap

The formulation of the experts framework from lecture 4 shows that the loss L incurred by the experts algorithm is close to the loss L^* suffered by the best expert,

$$L \leq (1+\epsilon)L^* + \frac{\log n}{\epsilon} \tag{1}$$

If the game is played with gains instead of losses for the experts, we multiply weights of the experts by $(1+\epsilon)^{-g_i}$ and a similar analysis shows that the gain G of the experts algorithm is close to the gain G^* for the best expert,

$$G \geq (1-\epsilon)G^* - \frac{\log n}{\epsilon} \tag{2}$$

1.2 The toll congestion game

Given graph $G = (V, E)$ and pairs of vertices $(s_i, t_i), i \in [k]$ the problem is to find paths connecting each pair of vertices minimizing the maximum congestion over an edge.

The toll congestion game: The toll player has m strategies corresponding to edges $e \in E$ and the routing player has strategies corresponding to valid routings between (s_i, t_i). If the row player plays routing r and the column player chooses an edge e, the row player pays a toll of $c(e, r)$, the congestion on the edge e in routing r.

A mixed strategy for the row player is a probability distribution on the set of routings (s_i, t_i), the mixed strategies of the routing player are flows. The best response to a flow f is to play the edge with maximum congestion under f. If the routing player goes first, the optimal strategy is to play the flow minimizing the maximum congestion, so the value of the game is $\min_f \max_e c(e, f)$.

If the toll player assigns tolls w_e on the edges the expected payoff for routing r is the sum of the lengths of the paths between (s_i, t_i) in r,

$$A(r, w) = \sum_e w_e c(e, r) = \sum_{i \in [k]} w(r_i) \tag{3}$$

The best response to a toll strategy w consists of routing along shortest paths between (s_i, t_i) in the metric induced by w. In particular, given a strategy of the toll player the best response to it can be computed efficiently.

Experts for toll congestion: The zero sum game framework from lecture 5 has experts corresponding to strategies of the min player. The number of strategies for the routing player is exponential, so an efficient algorithm can not store weights for them. The solution is to assign experts to strategies of the toll player, with experts having gains (**??**) instead of losses.

The algorithm is the following:

1. The column player follows the experts algorithm, the initial weights are $w_e = 1$, the weights are updated to $w_i \to (1+\epsilon)^{g_i^t/k}$ where g_i^t is the gain for the i-th expert in round t.

2. The routing player plays the optimal response to the toll player's strategy, which is to route along the shortest paths connecting (s_i, t_i) under the metric w.

   The 'oracle' for the routing player finds shortest paths in a weighted graph.

3. The output of the algorithm is the maximum congestion c_{max} for the average vector of flows $f = \frac{1}{T}\sum_t f(t)$. The value of c_{max} is within 2ϵ of the optimal congestion C^* for $T = \frac{k \log m}{\epsilon^2}$.

The output of the procedure is a flow while the solution to the problem must be a routing. We will discuss a randomized rounding strategy for this later.

Analysis: The gain for the toll player over T rounds is at most C^*T as the routing player plays second. The gain for the best expert against the average flow $f = \sum f(i)/T$ is $c_{max}T$. The analysis of the experts algorithm (**??**) yields,

$$C^*T \geq c_{\max}(1-\epsilon)T - \frac{k \log m}{\epsilon} \tag{4}$$

Dividing by $T = \frac{k \log m}{\epsilon^2}$ in we obtain,

$$C^* \geq c_{\max}(1-\epsilon) - \epsilon \tag{5}$$

For a constant ϵ, the running time of the algorithm is $O(k^2 m \log n)$ as $O(k \log m)$ rounds are required and each round involves computing k shortest paths which can be done in time $O(km)$, say using breadth first search.

Trick to reduce running time: The running time for obtaining a constant factor approximation to the optimal congestion can be reduced to $\widetilde{O}(km)$, where $\widetilde{O}$ suppresses poly logarithmic factors. The idea is to split each round of the game into k stages, where the routing player routes one path in each stage and the toll player updates tolls accordingly. The payoffs for the modified game are 0/1 valued. [Exercise].

1.3 Rounding

The experts algorithm produces a close to optimal flow which needs to be rounded to a routing. The flow f produced by the experts algorithm is $f = \sum f_i/T$, where each f_i is a routing as it is the best response to some toll strategy. For each pair of nodes (s_i, t_i) there are a total of $l = T$ paths available and the rounding strategy is to choose one of these paths uniformly at random.

If the edge e has congestion c_e in the optimal flow f then the number of paths passing through e can be bounded by,

$$|P_e| \leq c_e T$$

The random variable $X_i = 1$ if the path through (s_i, t_i) in the rounded off strategy goes through e. The expected congestion for the edge e for the randomized strategy is,

$$E[\sum_{i \in [k]} X_i] = \sum_i \frac{|P_{e,i}|}{l} = \frac{P_e}{T} \leq c_e \tag{6}$$

It is not surprising that the expected congestion for the randomized rounding procedure is at most the congestion for the original flow. The random variables X_i are defined to be independent so that we use concentration bounds and argue that with high probability the rounding procedure finds a routing with low congestion. The analysis of the rounding procedure splits into two cases corresponding to high and low congestions.