A constant approximation algorithm for the *a priori* TSP

David Shmoys

In the traveling salesman problem, one is given N, a set of points, and for each pair of points in N, one is also given the distance between them, where we assume that these are symmetric and satisfy the triangle inequality; the aim is to find a tour τ through all points in N that minimizes the total length $c(\tau)$ of the tour. In the *a priori* TSP, one is also given a probability distribution Π over the subsets $A \subseteq N$ of so-called active sets. For each subset A, each tour τ induces a tour τ_A by "shortcutting" those points not in A; we let $c(\tau_A)$ denote the length of the resulting tour of the points in A. In the *a priori* TSP, we measure the quality of a tour τ by computing the expected length with respect to a random choice of A drawn according to Π, $E_A[c(\tau_A)]$; the aim is to compute a tour that minimizes this expectation. Let τ^* denote an optimal *a priori* tour. We consider the case in which Π is specified by giving independent activation probabilities for each point in N. We give a simple randomized 4-approximation algorithm for this problem, and then derandomize this algorithm to give a deterministic 8-approximation algorithm; that is, in polynomial time, we compute a tour τ such that $E_A[c(\tau_A)] \leq 8E_A[c(\tau^*_A)]$.

This is joint work with Kunal Talwar.