EE 223: Stochastic Estimation and Control	**Spring 2007**

Lecture 28 — May 03

Lecturer: Venkat Anantharam *Scribe: none*

This lecture was not scribed.

We discussed the notion of *regret* in stochastic adaptive control. It refers to the loss in performance associated to the lack of knowledge of the underlying model. We discussed the analysis of regret carried out in the papers of Agrawal, Teneketzis, and Anantharam, IEEE T-AC, March 89, pp. 258 -267 and Dec 89, pp. 1249 -1259. Controlled finite state, finite parameter, finite action space models were studied in these papers, both in the i.i.d. case and in the Markov case. This includes examples like the coin tossing example discussed in the preceding lecture. For a uniformly good scheme, the regret for each parameter value that has another parameter value posing an identification conflict with it must grow logarithmically in time; further there is a scheme that simultaneously achieves this lower bound for all parameter values (the coefficient of the logarithmic growth of regret for each parameter value is also precisely characterized in terms of certain Kullback Liebler distances from the statistics under the parameter and those under parameters that pose an identification conflict with the underlying parameter).

We then briefly discussed the Witsenhausen counterexample, described in Witsenhausen, SIAM Journal on Control, Vol. 6, No. 1, 1968, pp. 131-147. The optimality of linear controllers even in a simple linear quadratic control problem with Gaussian noise breaks down when the control is decentralized. The importance of *signaling* between the distributed controllers becomes apparent from the analysis of this example.