PROGRAM IN INDUSTRIAL & SYSTEMS ENGINEERING
IE 8773-8774

Linearly Parameterized Bandits

by

Dr. Paat Rusmevichientong
Assistant Professor, School of Operations Research and Information Engineering
Cornell University
Ithaca, NY 14853

Wednesday, April 29, 2009
3:15 p.m. — Refreshments before the seminar
3:30 p.m. — Graduate Seminar
Room 3125* Mechanical Engineering

ABSTRACT — Motivated by applications in e-commerce and revenue management, we consider multiarmed bandit problems involving a large (possibly infinite) collection of arms, in which the expected reward of each arm is a linear function of an unknown univariate or multivariate random variable. The objective is to choose a sequence of arms to minimize the cumulative regret and Bayes risk. For the univariate case, we find that a greedy policy suffices. For the multivariate case, when the set of arms is strongly convex, we provide a lower bound and a policy that results in a matching upper bound. In the absence of strong convexity, we describe a policy whose performance is within a polylogarithmic factor from the optimal.

This is a joint work with John Tsitsiklis and Adam Mersereau.

BIO — Paat Rusmevichientong is an Assistant Professor in the School of Operations Research and Information Engineering at Cornell University. Before joining the faculty at Cornell, he worked in the data mining and personalization group at Amazon.com. His research interests include data mining, information technology, and stochastic optimization, with applications to supply chain and revenue management. He has received the NSF CAREER Award and the Cornell Sonny Yau '72 Excellence in Teaching Award.

Host: Prof. William Cooper

*change of meeting room