abstract

COMPUTER SCIENCE/DISCRETE MATH, II
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

One of the simplest and fundamental distributed computing problems involving noise is the ``noisy broadcast problem'': There are n processors each holding a bit, and there is a special processor called receiver. Processors act according to a protocol which proceeds in time steps. At each time step some processor broadcasts a bit (all other processors receive the bit) which is a function of its input and bits from broadcasts in previous steps. (This is similar to the number-in-hand model of multiparty communication complexity.) However in our setting communication is noisy, which means that the bit a processor receives is flipped with a small probability \epsilon independently of all other events. The goal is for the receiver to learn the input bits of all processors with probability >2/3 for any input. The problem is how many broadcasts are needed? Gallager (1988) gave a protocol with O(n log log n) broadcasts. This result remained unimproved and no lower bounds better than the trivial \Omega(n) were known.

In this talk we show that Gallager's protocol is optimal. To this end, we introduce a new model of noisy computation, that may be interesting in its own right, called ``generalized noisy decision trees''. Decision trees in this model can query *any* Boolean function of a noisy copy of the input (obtained by flipping each bit independently with a small noise probability). We show lower bounds in this model, which then give lower bounds for the noisy broadcast problem.

This work is joint with Guy Kindler and Michael Saks.