abstract

COMPUTER SCIENCE/DISCRETE MATH, I
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

We will describe definitions and algorithmic results for the ``census problem''. Informally, in a census individual respondents give private information to a trusted (and trustworthy) party, who publishes a sanitized version of the data. There are two fundamentally conflicting requirements: privacy for the respondents and utility of the sanitized data. Unlike in the study of secure function evaluation, in which privacy is preserved to the extent possible given a specific functionality goal, in the census problem privacy is paramount; intuitively, things that cannot be learned ``safely'' should not be learned at all.

The definition of privacy and the requirements for a safe sanitization are important contributions of this work. Our definition of privacy formalizes the notion of protection from being brought to the attention of others -- one's privacy is maintained to the extent that one blends in with the crowd. Our definition of a safe sanitization emulates the definition of semantic security for a cryptosystem and says, roughly, that an adversary given access to the sanitized data is not much more able to compromise privacy than an adversary who is not given access to the sanitization.

Joint work with Shuchi Chawla, Frank McSherry, Adam Smith, and Hoeteck Wee.