Reducing the download complexity of Private Information Retrieval schemes

Simon Blackburn
Royal Holloway University of London, UK

Abstract. In the classical Private Information Retrieval (PIR) model, a user stores some data on a collection of servers and later wishes to retrieve a single bit of that data. This should be done without any individual server knowing which bit is being retrieved. The aim is to design a protocol that minimises the amount of communication between the user and servers. Variations of this model have been studied recently, motivated by applications in distributed storage. I will give a brief introduction to the area, and I will present some recent constructions and open problems. No knowledge of cryptography or distributed storage is required. This is joint work with Tuvi Etzion (Technion) and Maura Paterson (Birkbeck).