PANDA: Private Anonymous Data Access

Ariel Hamlin, Northeastern University

We consider a scenario where a server holds a huge database that it wants to make accessible to a large group of clients. After an initial setup phase, clients should be able to read arbitrary locations in the database while maintaining privacy (the server does not learn which locations are being read) and anonymity (the server does not learn which client is performing each read). This should hold even if the server colludes with a subset of the clients. Moreover, the run-time of both the server and the client during each read operation should be low, ideally only poly-logarithmic in the size of the database and the number of clients. We call this notion Private Anonymous Data Access (PANDA).

PANDA simultaneously combines aspects of Private Information Retrieval (PIR) and Oblivious RAM (ORAM). PIR has no initial setup, and allows anybody to privately and anonymously access a public database, but the server's run-time is linear in the data size. On the other hand, ORAM achieves poly-logarithmic server run-time, but requires an initial setup after which only a single client with a secret key can access the database. The goal of PANDA is to get the best of both worlds: allow many clients to privately and anonymously access the database as in PIR, while having an efficient server as in ORAM.

In this work, we construct bounded-collusion PANDA schemes, where the efficiency scales linearly with a bound on the number of corrupted clients that can collude with the server, but is otherwise poly-logarithmic in the data size and the total number of clients. Our solution relies on standard assumptions, namely the existence of fully homomorphic encryption, and combines techniques from both PIR and ORAM. We also extend PANDA to settings where clients can write to the database.