Practical Private Information Retrieval

Massachusetts Institute of Technology

Challenge:

- Private information retrieval (PIR): Query a database while hiding query from the server
- **Problem:** PIR is expensive in server-side computation
- **Goal 1:** Reduce server-side cost of PIR, in theory and practice.
- **Goal 2:** Use PIR to protect privacy in large-scale systems.

Solution:

- **Offline/online PIR:** Push the heavy computation into an offline phase, done in advance (USENIX Security 2021)
- **PIR with low amortized cost:** Construct PIR schemes for which the average per-query cost is small (Eurocrypt 2022)
- **PIR from lightweight primitives:** Use lattice-based crypto to reduce the concrete costs (In progress)

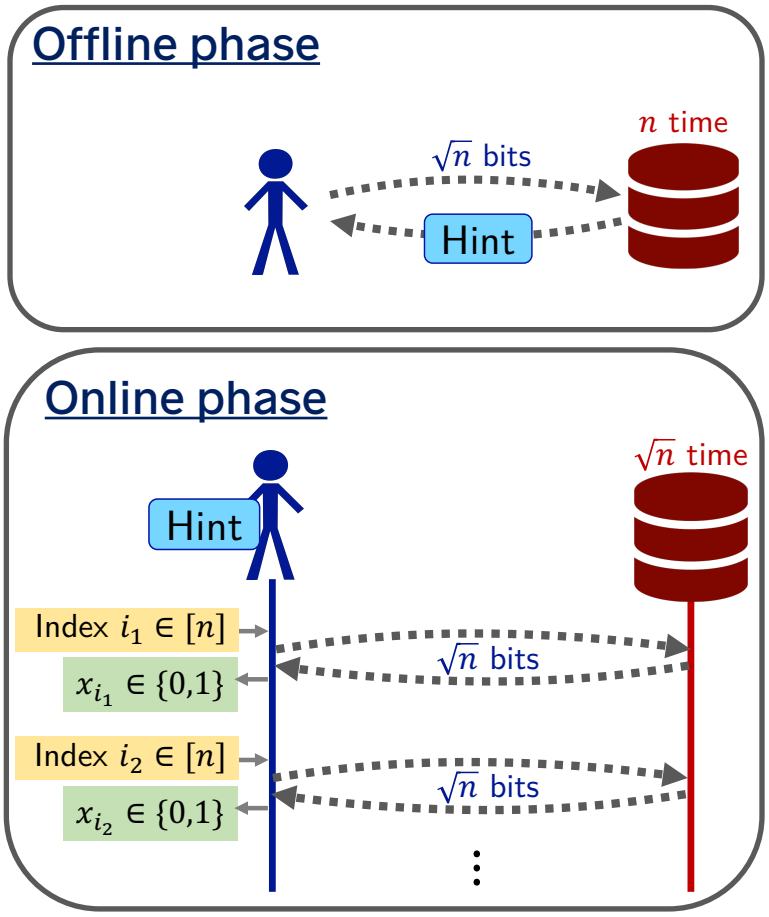

Figure: We give the first single-server PIR schemes in which the client can make a sequence of adaptive queries to an n-bit database at amortized server cost $\leq n$ per query.

Scientific impact:

- Use new models (offline/online, preprocessing, ...) to circumvent old impossibility results.
- Understand true cost of private database lookups
- Develop techniques with applications to other areas of cryptography (ORAM, MPC, etc.)

Broader impact:

- **Goal:** Perform web search without revealing query to search engine
- Collaborations ongoing with major tech cos. to apply PIR in web context (private ad retrieval, etc.)
- New undergrad and grad courses at MIT in cryptography and security, featuring PIR and other privacy tools

PI: Henry Corrigan-Gibbs (MIT)
Joint work with Alexandra Henzinger (MIT) and Dmitry Kogan (Stanford & Fordefi)
SaTC Core Small, Award #: CNS-2054869