Private Inference Control for Aggregate Database Queries

Geetha Jagannathan and Rebecca N. Wright, Rutgers University

INFERENCE CONTROL FOR PRIVATE QUERIES

Enables a client to perform aggregate database queries subject to a server's inference control policies:

- Client learns only the result of queries that pass the inference control rule, and nothing else about the database.

- Database learns nothing about the queries, including whether they pass.

Applies to statistical database queries as well as more general multi-argument queries.

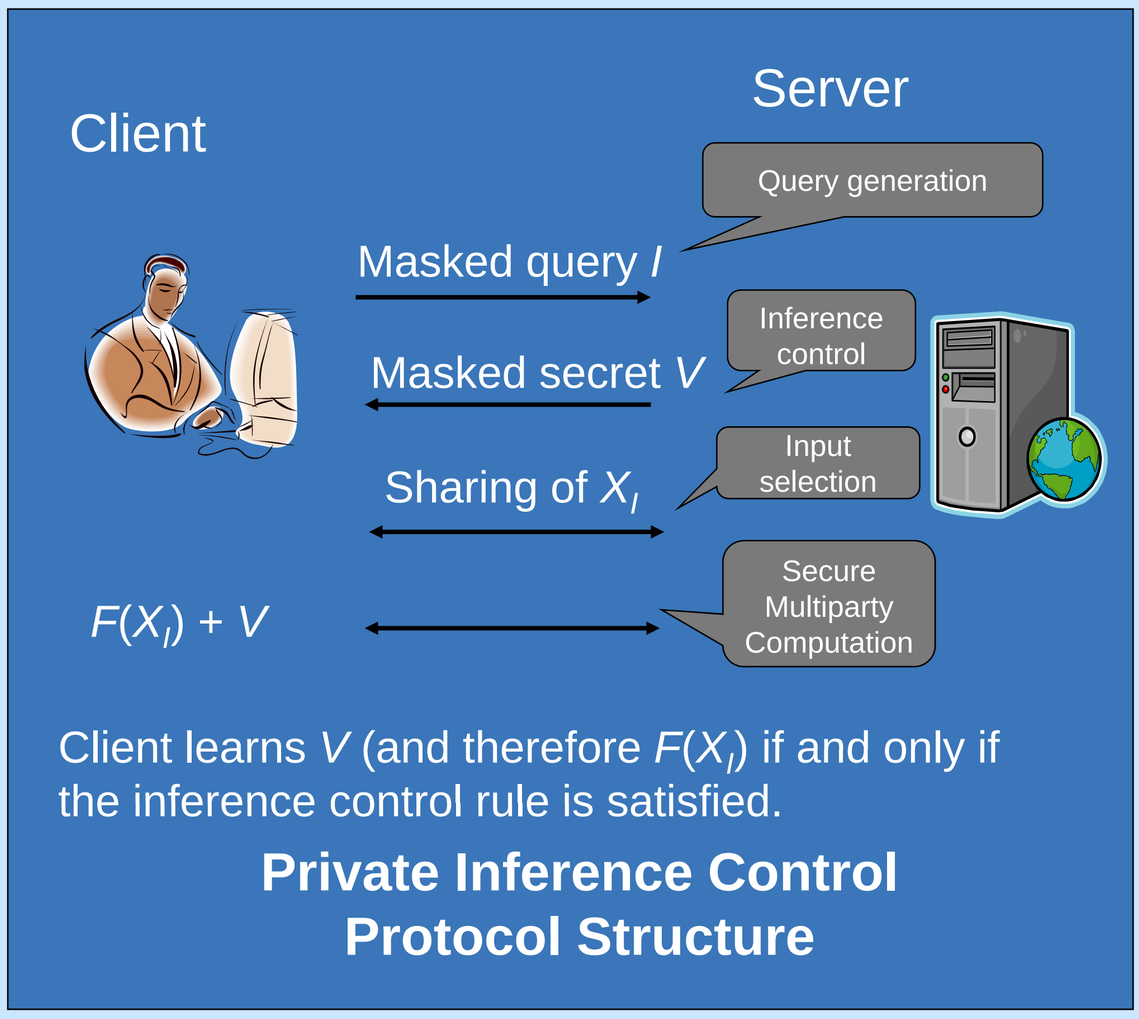

Client learns V (and therefore $F(X_I)$ if and only if the inference control rule is satisfied.

Private Inference Control Protocol Structure

Approach and Impact

New approach

- Extends earlier inference control work of Woodruff and Staddon (2004) to the more realistic case of aggregate queries.
- Solutions are more efficient than generic MPC.

Research Impact

- We provide several solutions that tradeoff communication and computation complexity depending on the number of queries and the size of each query.

We consider two specific inference control rules:

Strict ICR: A query is allowed if its indices do not intersect with the indices of any previous query.

Relaxed ICR: A query is allowed if the cardinality of the intersection of the indices with all previous queries is less than a specified threshold t.

We provide several cryptographic protocols to satisfy these inference control rules and the privacy requirements:

Protocol 1). Uses homomorphic encryption and oblivious polynomial evaluation.

Protocol 2). More efficient when there are fewer queries, but queries involve more indices.

Protocol 3). Reduced communication when there are more queries.

National Science Foundation
WHERE DISCOVERIES BEGIN

NSF Cyber Trust Principal Investigators Meeting
March 16-18, 2008
New Haven, CT

RUTGERS