SecureMCMR: Privacy-Preserving Computation Outsourcing for MapReduce Applications

Lindsey Kennard (PhD Student) and Ana Milanova (PI)

Department of Computer Science

Rensselaer Polytechnic Institute

Problem Statement: Use untrusted clouds (e.g., AWS, Google) to run MapReduce applications, while preserving (1) privacy of data and (2) efficiency of computation

Related Work: CryptDB (Popa et al., SOSP'11), Monomi (Tu et al., VLDB'13), MrCrypt (Tetali et al., OOPSLA'13), SecureMR (Milanova et al., Poster at PPoPP'18, HotSoS'18)

Key Tools: **L**inearly **H**omomorphic **E**ncryption (LHE), **R**andomized **E**ncoding (RE), and **O**rder **P**reserving **E**ncryption (OPE); Program Analysis

Key Problems

- Protocol design. LHE-based primitives are insufficient:
  - `x + y, c*x, c*x + y`, etc. ☑
  - $x_1*y_1 + \ldots + x_n*y_n$ ☒
  - $(x_1+\ldots + x_n) > y$ ☒

  (`x` and `y`'s are ciphertexts, `c` is plaintext)
- Program analysis and transformation of MapReduce programs w.r.t. these protocols
- Privacy: formal security analysis
- Efficiency: optimizations, goal is ≈5x overhead over plaintext

Scientific Impact

- Results have applications to Secure Multi-Party Computation (MPC), a growing alternative to computation outsourcing
- Upcoming paper in CCS'19 on program analysis for protocol selection for MPC
- Optimizations, essentially automatic parallelization, relevant to MPC
- Towards MPC-based outsourcing for MapReduce applications

Key Results and Contributions

- SecureMCMR: Clouds *A* & *B* collaboratively execute program *P*
  - Invariants: *A* sees LHE-ciphertexts, *B* sees blinded plaintexts
  - RE-based protocols for multiplication and comparison
  - Reasoning about OPE-security of program *P* using Adversary Advantage
- Classification of ≈40 standard MapReduce benchmarks
  - Majority are Secure or OPE-Secure
- Experiments on Google Cloud and AWS on 3 MapReduce benchmarks
  - 3x to 5x overhead over plaintext execution

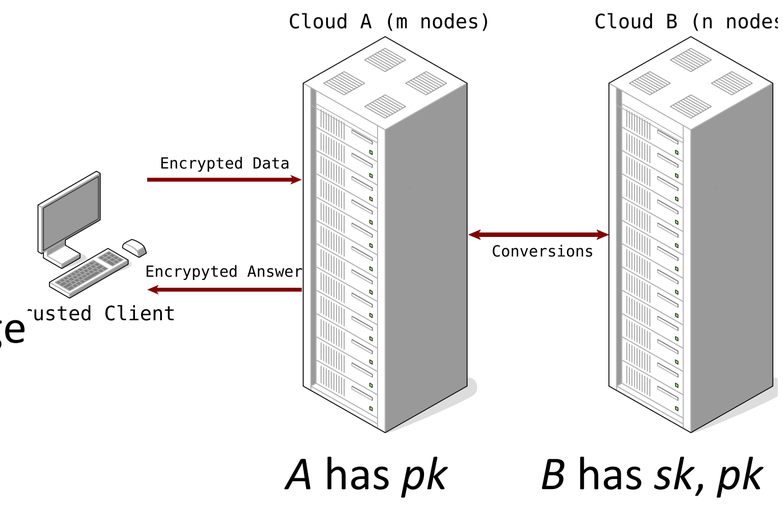

Impact on Society

- Large amounts of data
- Inexpensive and powerful cloud infrastructures (e.g., AWS, Google Cloud)
- Security remains an obstacle to the adoption of computation outsourcing

Education and Outreach

- PhD students and PI are underrepresented minority or female
- PI teaches undergraduate

1. Modern Binary Exploitation
2. Program Analysis

Potential Impact

- Better security of computation outsourcing
- Wider adoption of computation outsourcing
- New applications in medical research