Modern Secure Multi-Party Computation Techniques

Marina Blanton

Department of Computer Science and Engineering
University at Buffalo

CSE 501 Fall 2024

Secure Computation

- Secure multi-party computation enables parties to jointly evaluate a function on their private inputs without disclosing them
  - each participant holds their own private data
  - participants jointly perform the computation on cryptographically protected private data
  - private data can only leave the owner after applying proper cryptographic protection
  - only the outcome is revealed to the intended output recipients

The Famous Millionaires Problem

- Alice and Bob would like to determine who is richer without revealing their worth to each other

Alice
private x

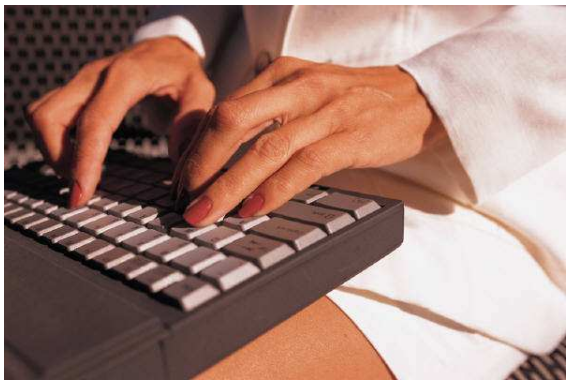

$\longrightarrow$
$\longleftarrow$
$\longrightarrow$
$\longleftarrow$
$\longrightarrow$

Bob
private y

output $x < y$

Secure Computation

- This is very different from traditional ways of computing:
  - obtain access to private data and promise to comply with data usage requirements
  - or private data is accessible to software, but the software doesn't let users to "see" the data

Secure Computation

- This is very different from traditional ways of computing:
  - obtain access to private data and promise to comply with data usage requirements
  - or private data is accessible to software, but the software doesn't let users to "see" the data
- Contrast this with secure computation:
  - nothing about private data is recoverable by others throughout the computation
  - observed information is the same as if the computation was performed by a trusted third party

Secure Multi-Party Computation

- Secure computation can take many forms
  - secure two-party computation
    - e.g., evaluating predisposition to a genetic disease

  - secure multi-party computation
    - e.g., determining the best treatment for a rare condition by multiple hospitals

  - secure computation outsourcing
    - e.g., offloading image processing to a cloud

Secure Multi-Party Computation

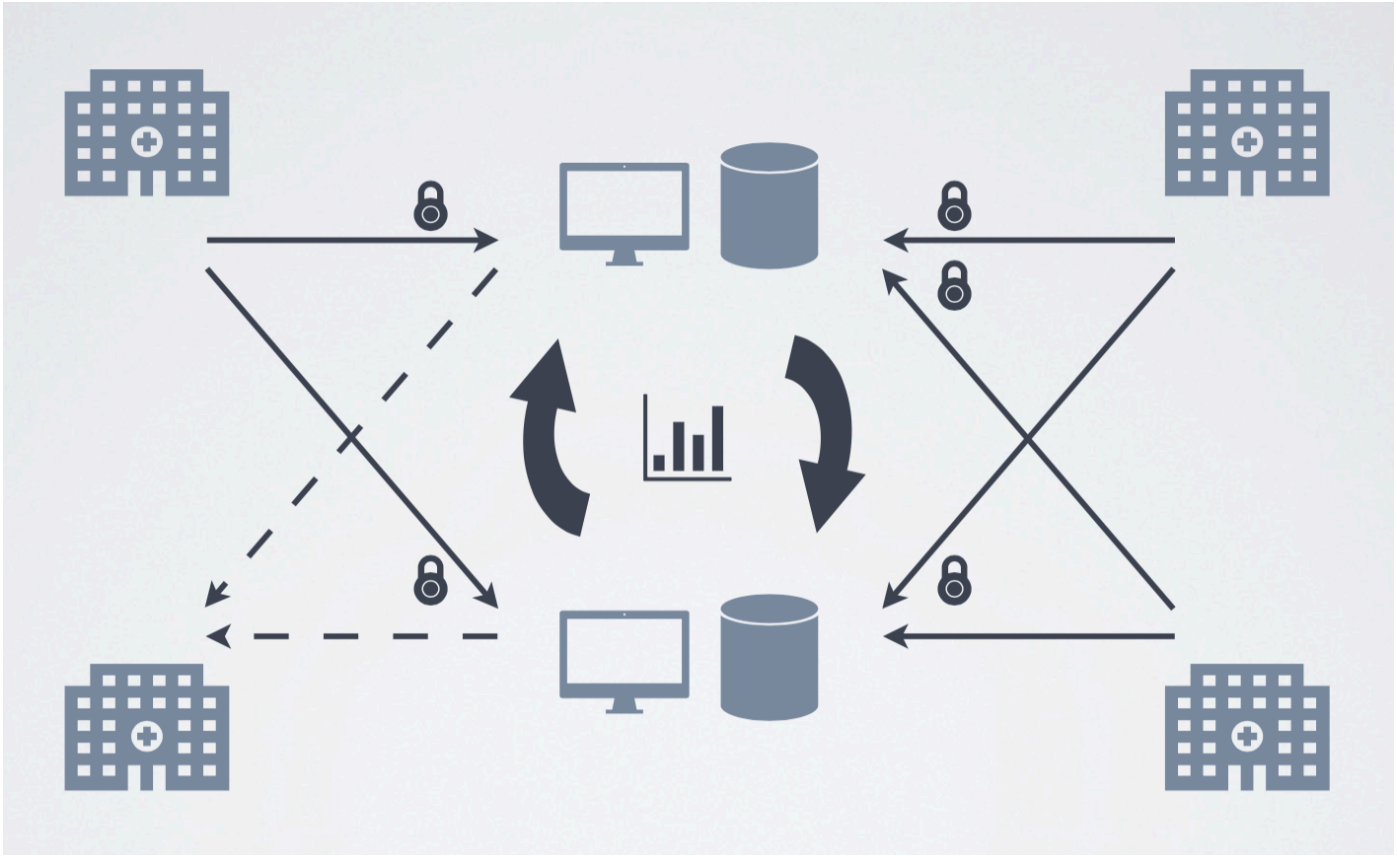

Data Protection Techniques

- There are different ways to protect private values in a way that we can compute with protected data
  - garbled circuit evaluation
  - secret sharing
  - homomorphic encryption

Data Protection Techniques

- There are different ways to protect private values in a way that we can compute with protected data
  - garbled circuit evaluation
  - secret sharing
  - homomorphic encryption
- With (n, t) secret sharing
  - a secret s is divided into n shares $s_1, s_2, \ldots, s_n$
  - access to $\leq t$ shares reveals nothing about s
  - access to $\geq t + 1$ shares allows for s to be reconstructed

Secret Sharing

- Example: Additive secret sharing
  - $t = n - 1$, i.e., all n shares are required for reconstruction
  - let $n = 2$ and our secret be $0 \leq x < N$
  - choose random r from $\mathbb{Z}_N$ and set the first share $x_1 = r$
  - compute the second share $x_2 = (x - r) \bmod N$
  - to reconstruct, compute $x = (x_1 + x_2) \bmod N$
  - for example
    - let $N = 10$ and $x = 3$
    - suppose we choose random $x_1 = 5$
    - we compute $x_2 = 3 - 5 \bmod 10 = 8$

Security of Secret Sharing

- Unlike encryption, secret sharing is unbreakable
  - secret sharing enjoys information theoretic security and achieves perfect secrecy
  - intuitively, in this example, x_1 is uniformly random and is independent of the secret
  - similarly, x_2 is uniform and there is not enough information to guess the secret
- More interesting and useful are threshold secret sharing with flexible t
  - honest majority setting with $t < n/2$ and
  - dishonest majority setting $t < n$

Threshold Secret Sharing

- One example is Shamir secret sharing
  - a secret is encoded as a polynomial of degree t with random coefficients
  - each share corresponds to the evaluation of the polynomial on a new point
  - secret reconstruction uses polynomial interpolation with $\geq t + 1$ points
- Another example is replicated secret sharing
  - additive secret sharing where a party stores > 1 share
  - a new share is created for each set of parties of size t and is given to the parties not in the set

Computing on Shares

- Computing on shares can often be expressed as evaluating an arithmetic circuit
  - evaluation of addition gates is performed locally on the shares
    - all of the above secret sharing schemes are linear
    - given secret-shared $[a_1], \ldots, [a_k]$, the parties can locally evaluate a linear combination $c_1[a_1] + \cdots c_k[a_k]$
  - evaluation of multiplication gates $[a] \cdot [b]$ requires communication
    - typical logic is to compute a partial product locally and re-share
    - the goal is to minimize communication and local work
    - different types of secret sharing offer performance tradeoffs

General-Purpose Computation

- Any functionality can be expressed as an arithmetic circuit
  - this doesn't mean that we can evaluate general-purpose programs efficiently
- In practice, we
  - extend the set of elementary operations with additional gates, most importantly, random value generation
  - design custom efficient protocols for non-arithmetic operations such as comparisons
  - minimize multiplicative depth or round complexity

Security Notions

- There are two standard threat models:
  - passive (or semi-honest) security
    - the computational parties follow the protocol but attempt to gather unauthorized information
    - communication channels are publically observable
  - active (or malicious) security
    - the computational parties can arbitrarily deviate from the prescribed protocol
    - detection of protocol deviation incurs additional cost
    - one example is dual execution on $[x]$ and $[r \cdot x]$ values and checking the branches for consistency

Real-World Uses

- Danish sugar beet auction
  - country-wide sealed bid auction that cryptographically protected farmer's bids
- Tax fraud detection in Estonia
  - citizens pay taxes on goods acquired abroad
  - the government can compute taxes without access to purchases
- Correlation between student employment and college graduation
  - employment information was extracted from tax records
  - it was combined with university records

Real-World Uses

- City of Boston gender wage gap study
  - analysis of wages by gender and race in the Greater Boston area
  - significantly larger participation the second year
  - provided valuable insights for decision makers
- Effectiveness of Google ads
  - Google knows on what ads a user clicked
  - the store knows how much a user spent at the store
  - this allows for determining effectiveness of ads
- Improving effectiveness of public transport
  - trip records were privately analyzed to improve public transport utilization

Current Research Directions

- Secure computation with certified inputs
  - security definitions place no constraints on inputs
  - some applications require authentic inputs
  - current research project: compliance with sustainability policies
- Understanding information disclosure from the output
  - security definitions require no information disclosure during the computation
  - but they don't impose constraints on disclosure from the output
  - current research project: we measure this for different statistical functions
- Performance optimizations