Privately Releasing Statistics

NAS CNSTAT Privacy Workshop
Panel on Current Capabilities of Differential Privacy

Ashwin Machanavajjhala

ashwin@cs.duke.edu

Differential Privacy - Recap

For every pair of **Neighboring Tables**

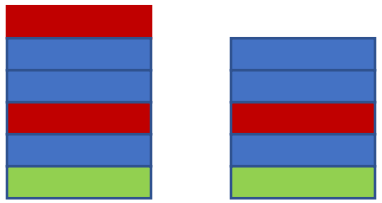

D_1 D_2

For every output

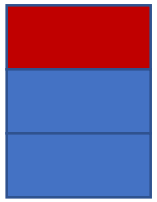

O

Should not be able to distinguish whether O was generated by D_1 or D_2

$$\log\left(\frac{\Pr[A(D_1) = O]}{\Pr[A(D_2) = O]}\right) < \varepsilon \qquad (\varepsilon>0)$$

Adding or removing a row from the input table should not significantly impact the output of the algorithm.

Differential Privacy - Recap

For every pair of **Neighboring Tables**

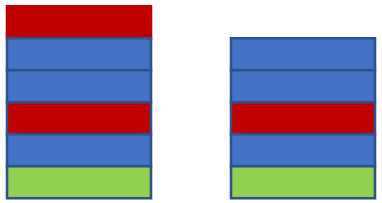

D_1 D_2

For every output

O

Should not be able to distinguish whether O was generated by D_1 or D_2

$$\log\left(\frac{\Pr[A(D_1) = O]}{\Pr[A(D_2) = O]}\right) < \varepsilon \quad (\varepsilon>0)$$

- Neighboring tables differ in one row.
- Consider all pairs of tables, and not just the actual input to the algorithm
- May depend on what must be kept secret
  - Privacy of persons
  - Privacy of households
  - Privacy of businesses

Differential Privacy - Recap

For every pair of Neighboring Tables

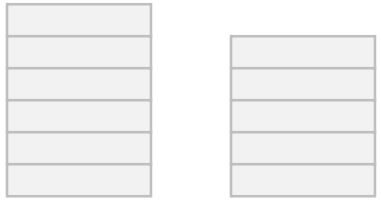

D_1 D_2

For every output

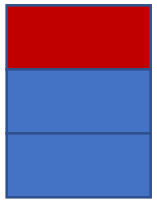

O

Should not be able to distinguish whether O was generated by D_1 or D_2

$$\log\left(\frac{\Pr[A(D_1) = O]}{\Pr[A(D_2) = O]}\right) < \varepsilon \quad (\varepsilon>0)$$

- Privacy bound must hold for all possible outputs
  - A worst case guarantee
- Outputs can be any type
  - Statistics
  - Contingency tables
  - Microdata
  - Regression parameters
  - ML models

Differential Privacy - Recap

For every pair of Neighboring Tables

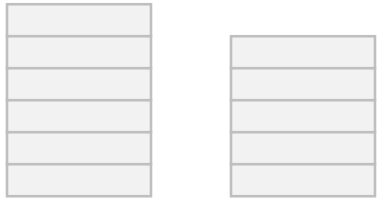

D_1 D_2

For every output

O

Should not be able to distinguish whether O was generated by D_1 or D_2

$$\log\left(\frac{\Pr[A(D_1) = O]}{\Pr[A(D_2) = O]}\right) < \varepsilon \qquad (\varepsilon>0)$$

- **Plausible Deniability:** Attacker can't tell whether input was D1 (with individual) or D2 (without individual)
- **Privacy Loss Parameter:** Larger epsilon is lesser privacy
- **Privacy Loss Budget:** Releasing multiple outputs results in additive increase in privacy loss.

Achieving differential privacy

- Techniques known for releasing outputs of several data analyses

- **Statistics and tabular summaries**
- Synthetic microdata
- Parameters of regression and statistical tests
- Machine learning models

Releasing a count: Laplace Mechanism

Add noise to each count to hide the contributions of one individual

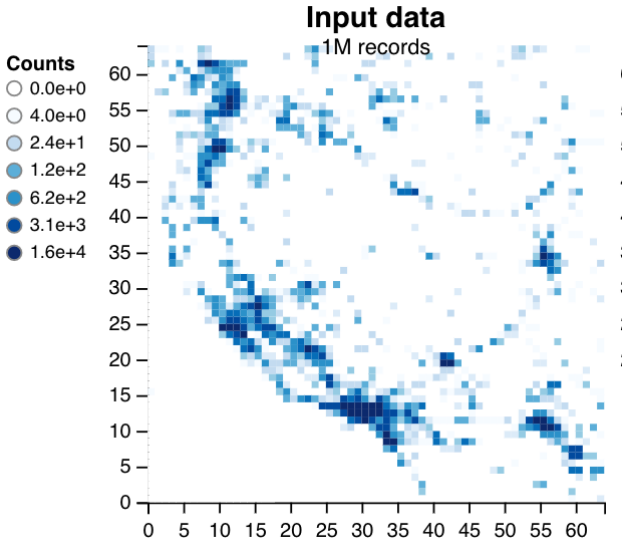

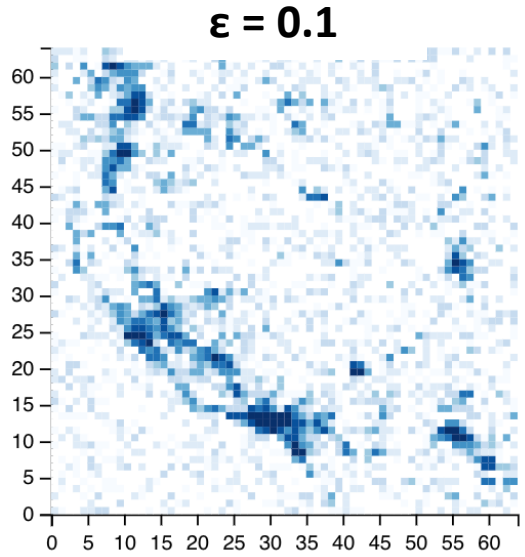

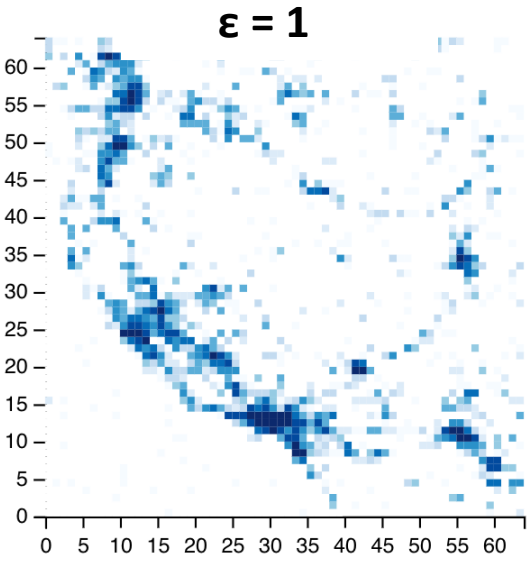

From single counts to tabular summaries

Release several contingency tables … at different levels of geography

- Total population
- … by Age
- … by Sex
- … by Age x Sex
- … by Race
- … by Age x Race
- …

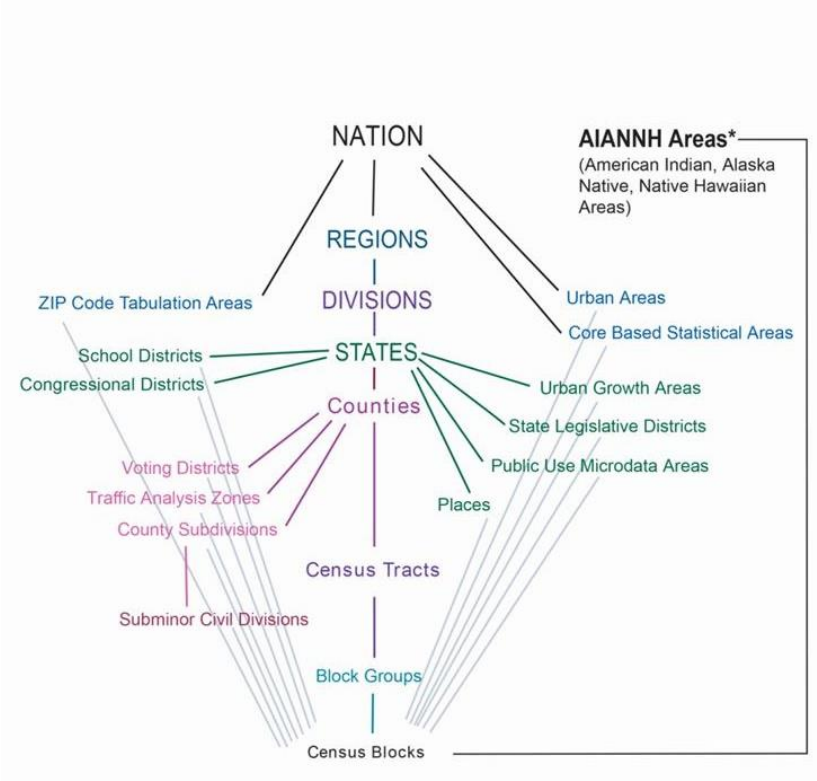

Key idea 1: Composition

- *Sequential Composition*:
  Privacy loss is additive across multiple releases …

- *Parallel Composition:*
  … unless they are run on disjoint subsets of data (e.g. across states)

- Algorithm: Use Laplace mechanism with parameter ε:
  - For each marginal (total, by age, by sex, …)
  - And for each geography (national, state, county, …)

Total privacy loss: # tables * #geo levels * ε

Challenge

- Algorithm: Use Laplace mechanism with parameter epsilon:
  - For each marginal (total, by age, by sex, …)
  - And for each geography (national, state, county, …)

- Problem: Consistency
  - Released statistics do not add up
  - E.g.: State counts do not add up to national counts

Key idea 2: Postprocessing & Inference

- *Postprocessing theorem*:
  Postprocessing the output of a DP mechanism does not degrade privacy

- Idea: Inference

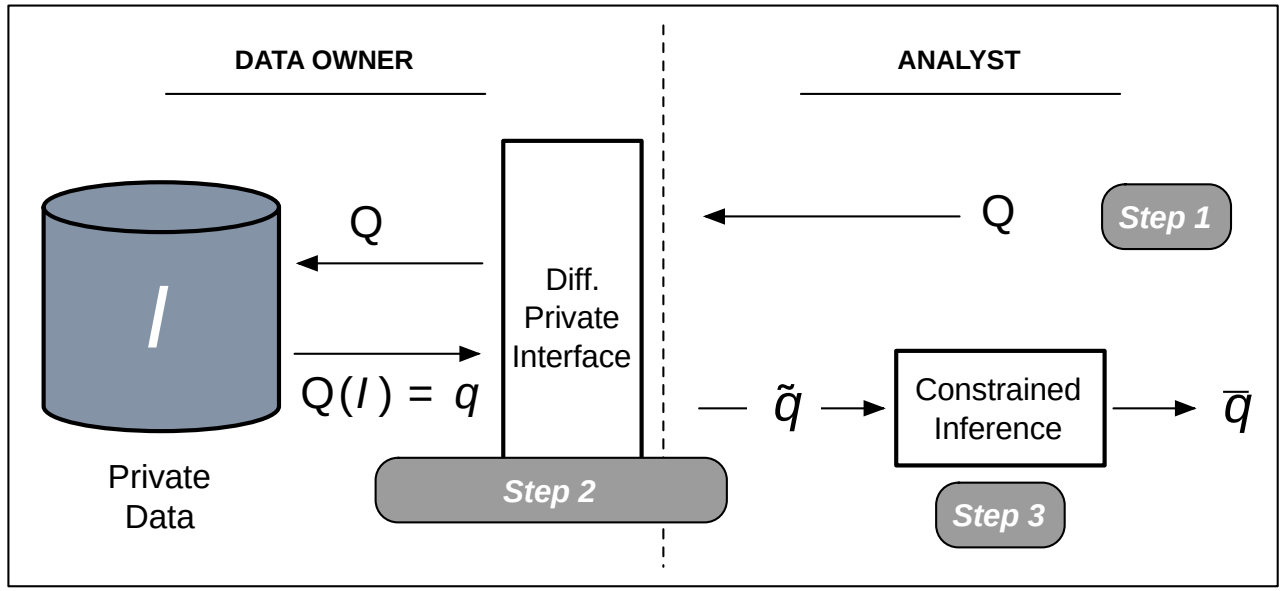

Challenge

- Algorithm: Use Laplace mechanism with parameter epsilon:
  - For each marginal (total, by age, by sex, …)
  - And for each geography (national, state, county, …)

Total privacy loss: # tables * #geo levels * ε

- Problem: Privacy loss adds up
  - Either get all the results accurately but with poor overall privacy loss
  - Or get a bounded privacy loss, but all the statistics have high error

Key idea 3:
Carefully select what to add noise to.

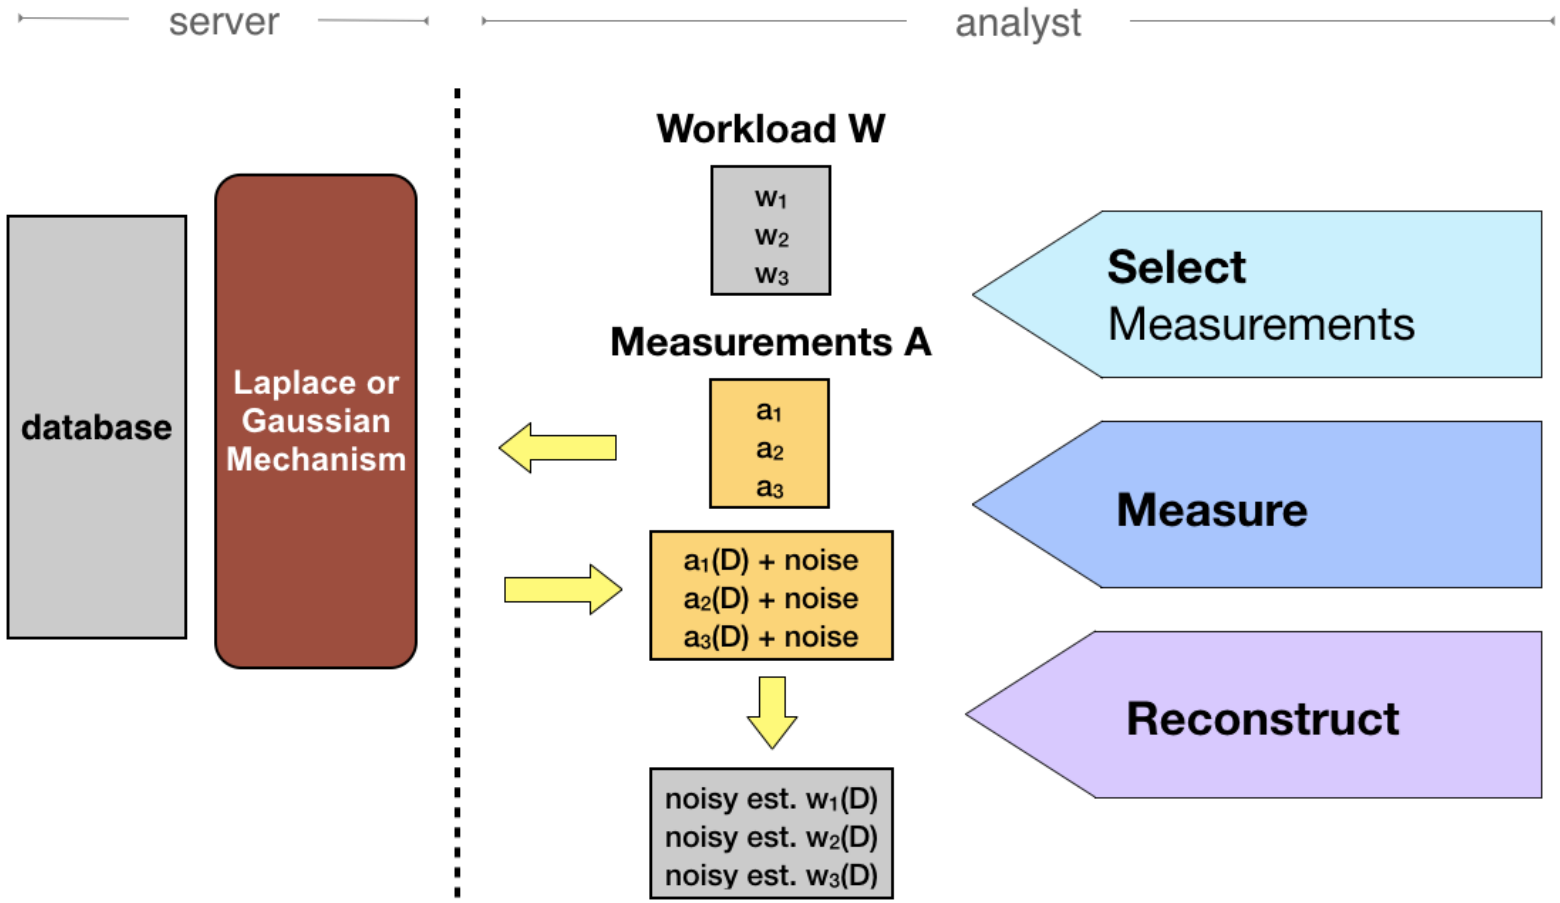

Key idea 3:
Carefully select what to add noise to.

- Total population
- … by Age
- … by Sex
- … by Age x Sex
- … by Race
- … by Age x Race

Select and Measure these tables

Reconstruct these tables using inference

Select-Measure-Reconstruct

[Hardt-Talwar 2010]

- We know tight lower bounds on the error of a set of linear queries under a fixed budget

- We know efficient methods to automatically choose the right strategy
  - K-norm Mechanism
  - Matrix Mechanism
  - **HDMM:** High Dimensional Matrix Mechanism

[Hardt-Talwar 2010]

[Li et al 2010]

[McKenna et al 2018]

In ongoing experiments with US Census Bureau products (2020 Decennial, Business Dynamics Statistics), HDMM reduced error by factors of 3x – 48x compared to baseline algorithms.

More DP algorithm design ideas

[Hardt et al 2012]
- Iteratively construct a "synthetic database" by measuring the query with most error

[Mironov 2017]
- Clever proof techniques to lower privacy loss

[Zhang et al 2014]
- Reduce the dimensionality of the data or statistics released

[Li et al 2014]
- Data dependent noise addition

[Kotsoginannis et al 2019]
- Truncating the data (Lipshitz extensions) for queries with higher sensitivity (e.g., queries with joins, counts)

Sophisticated algorithms lower error at same level of privacy

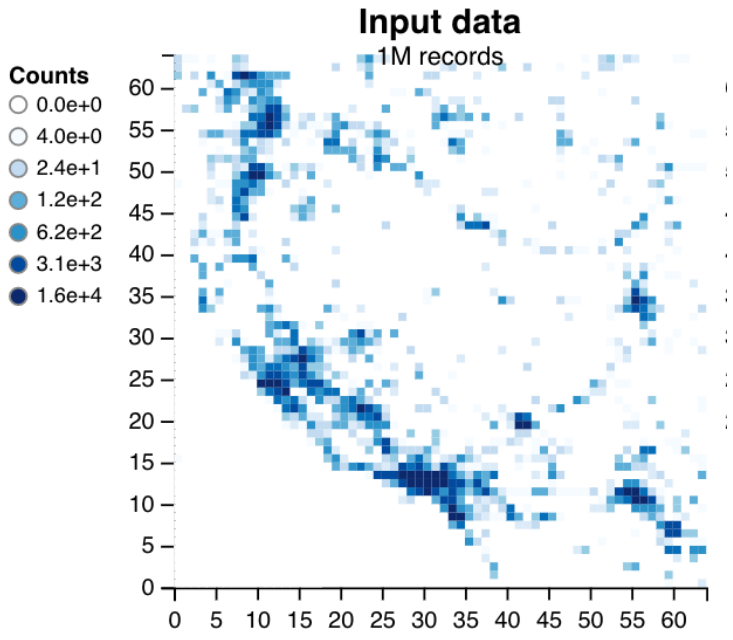

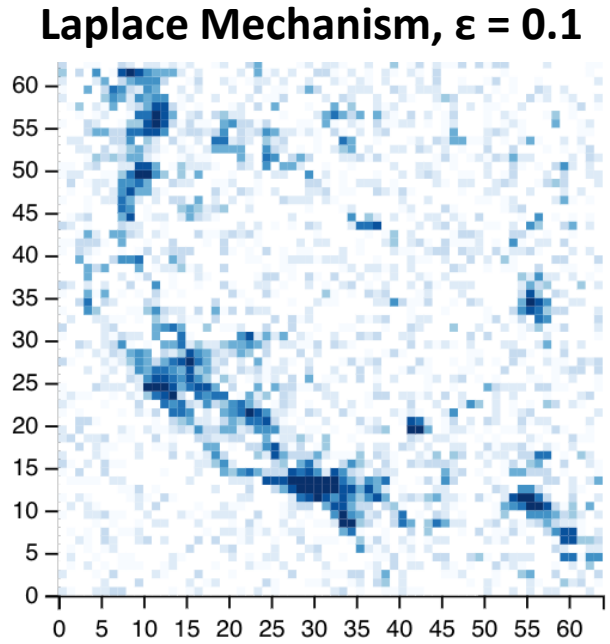

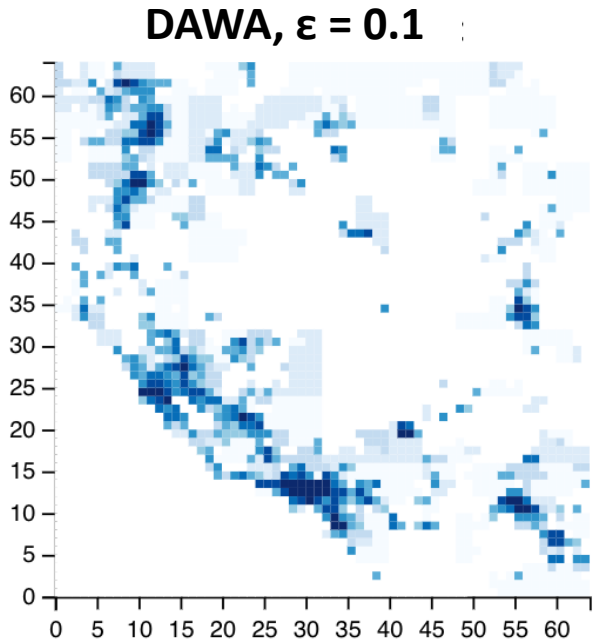

Open source DP tools for practitioners

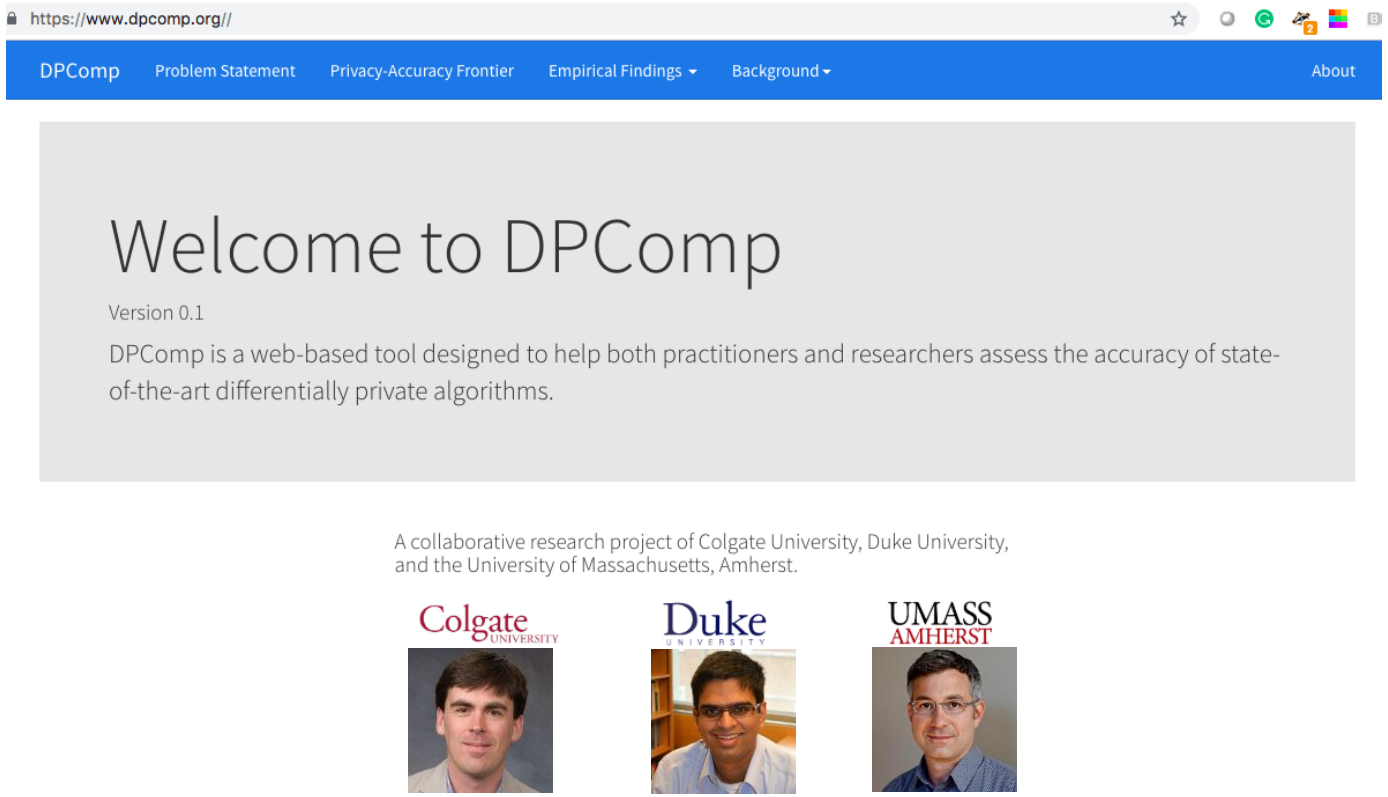

https://www.dpcomp.org//

Ektelo

Ektelo is a novel programming framework and system for implementing both existing and new privacy algorithms.

View the Project on GitHub
https://github.com/ektelo/ektelo

https://ektelo.github.io/

References

Hardt, Talwar
"*On the Geometry of Differential Privacy",* **STOC 2010**

Hardt, Ligett, McSherry
"*A Simple and Practical Algorithm for Differentially Private Data Release",* **NeurIPS 2012**

Li, Hay, Miklau, Wang
"A Data- and Workload-Aware Algorithm for Range Queries Under Differential Privacy", **VLDB 2014**

Zhang, Chen, Xu, Meng, Xie,
"*Towards accurate histogram publication under differential privacy",* **ICDM 2014**

Hay, Machanavajjhala, Miklau, Chen, Zhang,
"*Principled evaluation of differentially private algorithms using* **DPBench", SIGMOD 2016**

Mironov "*Renyi DP",* **CSF 2017**

Zhang, McKenna, Kotsogiannis, Hay, Machanavajjhala, Miklau
"Ektelo*: A framework for describing differentially private algorithms***", SIGMOD 2018**

Kotsogiannis, Tao, He, Machanavajjhala, Hay, Miklau
"PrivateSQL: *a differentially private SQL engine",* **CIDR 2019 (in submission VLDB 2019)**