Differential Privacy for DAP

Junye Chen, Audra McMillan, Christopher Patton,
Kunal Talwar, Shan Wang

IETF 118 – PPM

DP Motivation - from IETF117

- Keeping the measurements private (as DAP does) may not be enough: **the aggregate result may leak (bits of) an individual measurement**
    - Motivating example: Average height of a group of people with or without an especially tall (or short) individual
- Differential Privacy (DP): the aggregate result (or, more generally, the adversary's view) **should not change significantly if any one measurement is replaced by another**
    - Achieved by adding noise to:
        - the measurements by the Clients; and/or
        - the aggregate shares by the Aggregators.

DP Background

- DP is a class of definitions, e.g., ε-DP, (ε, δ)-DP, Rényi-DP, and each of them can be the preferable one depending on the application.
- DP is in the eye of the beholder: what DP guarantee you get against a particular adversary is a function of what information is available to that adversary. Hence, we need to define trust models.

New draft: draft-wang-ppm-differential-privacy-00

- Choose a class of DP notions that are suitable for DAP, e.g., pure ε-DP, approximate (ε, δ)-DP.
- Define various trust models that we aim to achieve DP in.
- Refine interfaces for “DP mechanisms”.
- Refine interfaces for “DP policies” that are implemented with DP mechanisms and composed with VDAFs.
- Describe concrete use cases, e.g., Histogram, with DP achieved by different DP policies.

Our audiences

- DAP deployments that want a “cookbook” for making their applications differentially private.
- DP researchers and domain experts.

Standardize DP Definitions

- ε-DP: ε describes the privacy loss of observing the aggregate result, when there is a change in the batch of measurements. Smaller ε means stronger privacy.
- (ε, δ)-DP: relaxes ε-DP by a small δ, which describes the probability of information leakage. Allowing for a small δ can allow randomized algorithms to add less noise. Smaller δ means stronger privacy.
- We note there are other DP definitions that we haven’t accounted for in the first version of the draft.

Trust Models

- Goal: Design DP policies that account for attackers that control the network and corrupt parties in DAP.
- We define three increasingly pessimistic trust models:
  - One-Aggregator-Most-Clients (OAMC)
    - Same trust model as Core DAP when all Clients are honest.
  - One-Aggregator-One-Client (OAOC)
  - One-Client (OC)

Trust Models - Hedging

- Hedging: Achieve some degree of privacy when an optimistic trust model's assumptions turn out to be false.
- For example: a DP policy achieves ideal ε in OAMC trust model. But if deployment turns out to be OAOC, then it's more desirable for the DP policy to maintain some weaker DP guarantee of ε'.

OAMC, ideal ε

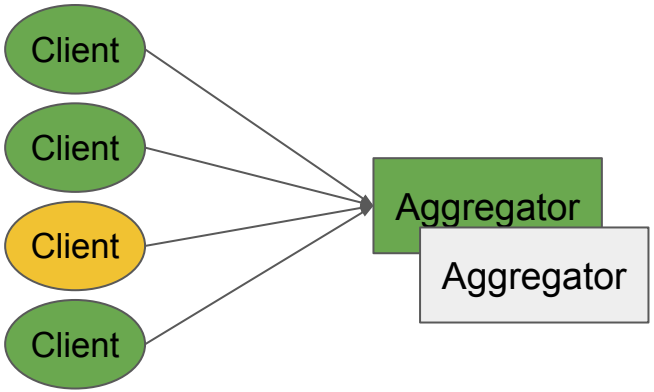

OAOC, ε' >> ε

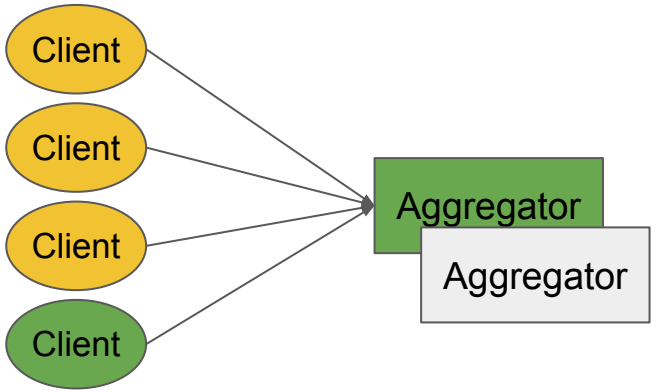

DP Mechanisms

- A DP mechanism is responsible for sampling noise with parameters derived based on the target DP.
- Examples:
  - Discrete Laplace [CKS'20]
  - Discrete Gaussian [CKS'20]
  - Symmetric RAPPOR [EPK'14, MJTB+'22]
- We want to standardize DP mechanisms to prevent implementation bugs that break DP [CSVW'22, JMRO'22, Mir'12].

[EPK'14] Erlingsson, Ú., Pihur, V., and A. Korolova, "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response", 2014
[MJTB+'22] McMillan, A., Javidbakht, O., Talwar, K., Briggs, E. "Private Federated Statistics in an Interactive Setting", 2022
[CKS'20] Canonne, C. L., Kamath, G., and T. Steinke, "The Discrete Gaussian for Differential Privacy", 2020
[CSVW'22] Casacuberta et al. "Widespread Underestimation of Sensitivity in Differentially Private Libraries and How to Fix It." CCS 2022
[JMRO'22] Jin et al. "Are We There Yet? Timing and Floating-Point Attacks on Differential Privacy Systems." IEEE S&P 2022
[Mir'12] Mironov. "On Significance of the Least Significant Bits For Differential Privacy." ACM CCS 2012

DP Policies

- A DP policy is implemented with DP mechanisms to endow VDAFs with DP.

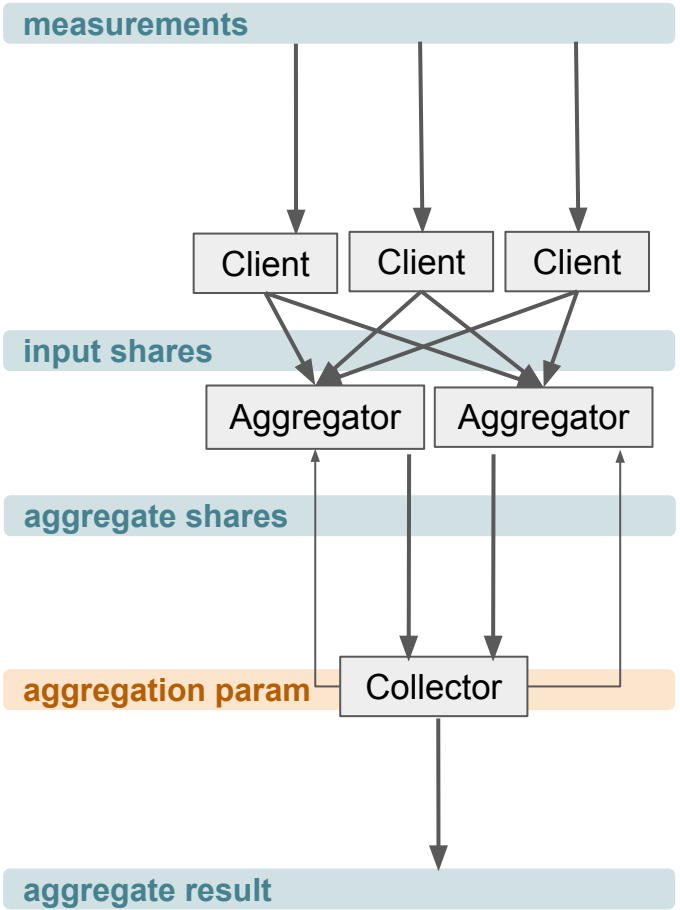

DP Policies

- A DP policy is implemented with DP mechanisms to endow VDAFs with DP.
- It requires applying DP mechanisms by Clients and/or Aggregators, and debiasing aggregate result by the Collector.

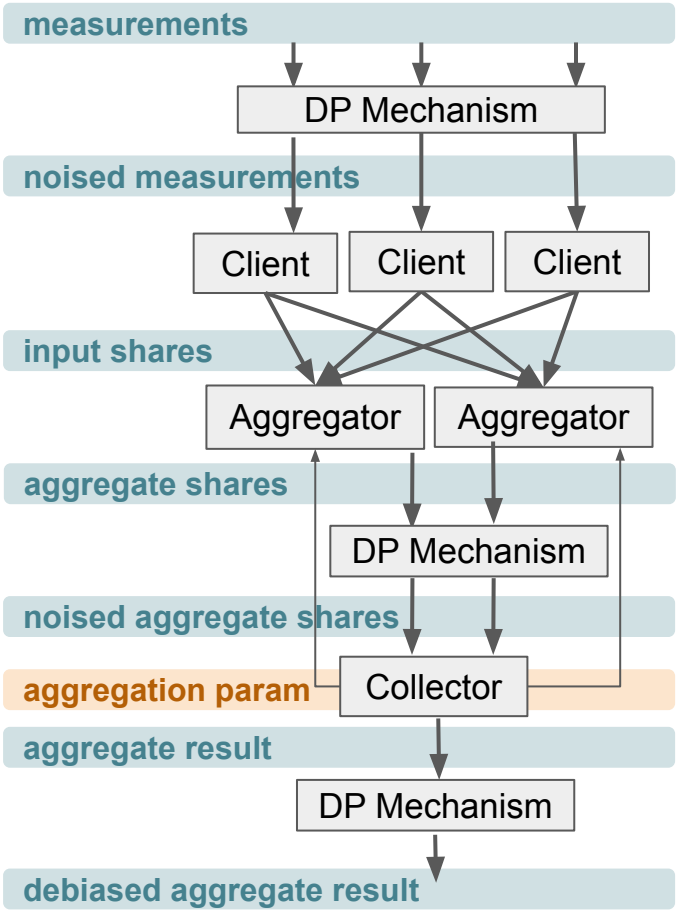

Use Case: Collecting histogram

- Goal: Achieve (ε, δ)-DP on collecting histogram, where each Client submits an one-hot vector.

	Policy 1: Pure Client Randomization	Policy 2: Pure Aggregator Randomization
Target Trust Model	OAMC	OAOC
DP Mechanism	Symmetric RAPPOR [EPK'14, MJTB+'22] from each honest Client	Discrete Gaussian [CKS'20, BW'18] from each honest Aggregator
VDAF	Prio3MultiHotHistogram*	Prio3Histogram

Table 1: DP Policies for Histogram.
*We note Prio3MultiHotHistogram is a private VDAF.

[EPK'14] Erlingsson, Ú., Pihur, V., and A. Korolova, "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response", 2014
[MJTB+'22] McMillan, A., Javidbakht, O., Talwar, K., E. Briggs. "Private Federated Statistics in an Interactive Setting", 2022
[CKS'20] Canonne, C. L., Kamath, G., and T. Steinke, "The Discrete Gaussian for Differential Privacy", 2020
[BW'18] Balle, B. and Y. Wang, "Improving the Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising", 2018

Use Case: Collecting histogram - Utility

- Either DP policy has utility advantage in different settings of (ε, δ)-DP.
- Noise is doubled in the policy with pure Aggregator randomization.

ε	δ	Standard Deviation of Pure Client Randomization	Standard Deviation of Pure Aggregator Randomization (two Aggregators)
0.32	1e-9	26.14	33.09
0.91	1e-9	12.28	12.08
1.53	1e-9	9.59	7.35

Table 2: Utility of DP policies in different (ε, δ)-DP.
Lower standard deviation means better utility.

Future Work

- Work out the implementation details of DP mechanism.
- Figure out if there are other quantitative and qualitative criteria to evaluate DP policies.
- Figure out if it's worth discussing MPC protocols [KKLVH'23] for Aggregators to collectively add noise.
- More concrete use cases.

[KKLVH'23] Keeler D., Komlo C., Lepert E., Veitch S., and X. He. "DPrio: Efficient Differential Privacy with High Utility for Prio", 2023

Questions

We feel this work is important and that PPM is well-positioned to take it on.

1. Is this work useful?
2. Is the draft scoped properly? Any suggestions?
3. Should PPM adopt this draft?