ADAPTIVE INDIVIDUAL ACCOUNTING WITH THE TARGET CHARGING TECHNIQUE

Alison Caulfield, Mark Chen, Peihan Liu

Department of Computer Science | Columbia University

Advised by: Rachel Cummings, Roxana Geambasu, Mathias Lecuyer, Pierre Tholoniat

Motivation

- **Individual Differential Privacy (IDP)** enables per-user privacy accounting -- ideal for workloads where users contribute unevenly to measurements.
- A prime example is **ad measurement**, where most users contribute nothing to most reports. Our group was the first to identify this efficiency opportunity and design an **IDP-based architecture** for advertising measurements (*Cookie Monster*), now forming the basis of a **W3C draft standard** that will enable privacy-preserving ad measurement without user tracking.
- As active participants in this standardization effort, we're tackling emerging theoretical challenges that affect not only ad measurement but the broader use of IDP -- most notably, **bias introduced by IDP filters in query results**, a key obstacle to achieving both privacy and utility at scale.

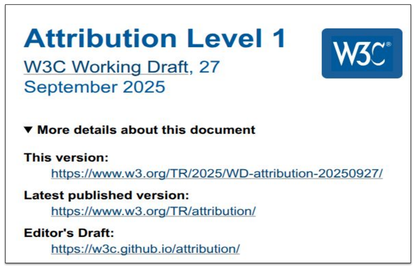

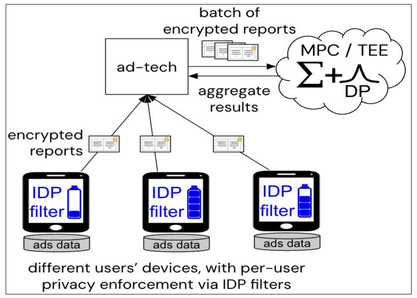

IDP Filters in Attribution API

Problem

- IDP introduces **selection bias** because users' privacy filters deplete at different rates, leading to **biased aggregate results**.
- We study the **Target-Charging Technique (TCT)** -- a generalization of the DP Sparse Vector Technique -- to **quantify IDP-induced bias** efficiently from a privacy-loss standpoint.
- TCT has been analyzed in **standard DP settings** with **limited adaptivity** and is designed to avoid privacy loss for **out-of-target results** (T).
- Our goal: **Extend and strengthen TCT** -- both algorithmically and analytically -- to support **adaptivity** and **integration with IDP filters**, enabling its use in real-world IDP systems like the **Attribution** API, where adaptivity is intrinsic.

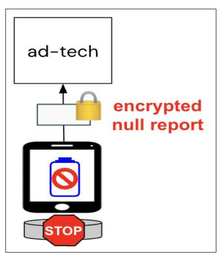

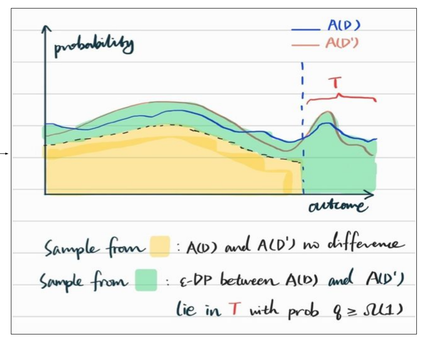

Our Work

- We are working on two complementary approaches, with plan to combine in the end. All are in progress and **we want your feedback**!

Part 1: TCT integration with Approximate IDP Filters

- TCT is **not adaptive**. It cannot handle changing datasets, on-demand queries, and adaptive budget per query.
- By integrating TCT analysis with approximate IDP filters, we achieve an algorithm that (1) **avoids privacy loss for out-of-target results**, **(2)** accounts for individual privacy loss and (3) supports at least **query adaptivity**.

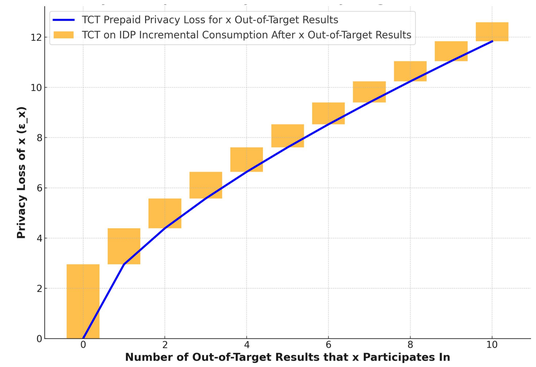

Preliminary results: TCT prepaid privacy loss vs. TXT on IDP incremental consumption. (Ignores the failure probability delta, which would increase privacy losses for TCT over IDP)

Part 2: TCT with RDP Analysis

- Why RDP for TCT? RDP filters smoothly **remove the δ arbitrary error** in the approximate IDP filters, and have cleaner compositions when considering **query and budget adaptivities** [4].
- We refine the TCT definitions (q-target etc) using the smoothed Rényi divergence for better bounds.

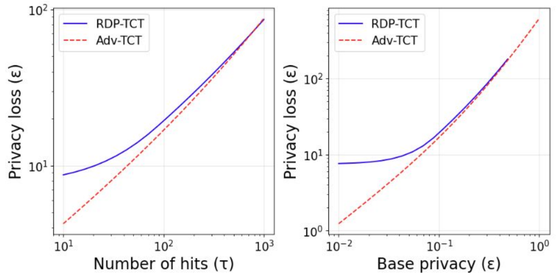

Preliminary results: RDP-TCT vs. Advanced Composition-TCT. (Assume pure-DP queries with TCT but composition with either RDP or Advanced Composition)

Part 3: Combine: TCT with individual RDP Filters

- Integrate the TCT's IDP aspect with RDP **individual filters** [5].

Related Work & Refs

- [1] Tholoniat, et al (2024, November). Cookie Monster: Efficient On-Device Budgeting for Differentially-Private Ad-Measurement Systems. In Proceedings of the ACM SIGOPS 30th Symposium on Operating Systems Principles (pp. 693-708).
- [2] Cohen, E., & Lyu, X. (2023). The target-charging technique for privacy analysis across interactive computations. Advances in Neural Information Processing Systems, 36, 62139-62168.
- [3] Kaplan, H., Mansour, Y., & Stemmer, U. (2021, July). The sparse vector technique, revisited. In Conference on Learning Theory (pp. 2747-2776). PMLR.
- [4] Lécuyer, M. (2021). Practical privacy filters and odometers with R\'enyi differential privacy and applications to differentially private deep learning. arXiv preprint arXiv:2103.01379.
- [5] Feldman, V., & Zrnic, T. (2021). Individual privacy accounting via a renyi filter. Advances in Neural Information Processing Systems, 34, 28080-28091.