DISTRIBUTED DIFFERENTIAL PRIVACY IN MULTI-ARMED BANDITS

Anonymous authors
Paper under double-blind review

ABSTRACT

We consider the standard K-armed bandit problem under a distributed trust model of differential privacy (DP), which enables to guarantee privacy without a trustworthy server. Under this trust model, previous work on private bandits largely focus on achieving privacy using a shuffle protocol, where a batch of users data are randomly permuted before sending to a central server. This protocol achieves (ε, δ) or approximate-DP guarantee by sacrificing an additive $O\left(\frac{K \log T \sqrt{\log(1/\delta)}}{\varepsilon}\right)$ factor in T-step cumulative regret. In contrast, the optimal privacy cost to achieve a stronger $(\varepsilon, 0)$ or pure-DP guarantee under the widely used central trust model is only $\Theta\left(\frac{K \log T}{\varepsilon}\right)$, where, however, a trusted server is required. In this work, we aim to obtain a pure-DP guarantee under distributed trust model while sacrificing no more regret than that under the central trust model. We achieve this by designing a generic bandit algorithm based on successive arm elimination, where privacy is guaranteed by corrupting rewards with an equivalent discrete Laplace noise ensured by a secure computation protocol. We also show that our algorithm, when instantiated with Skellam noise and the secure protocol, ensures *Rényi differential privacy* – a stronger notion than approximate DP – under distributed trust model with a privacy cost of $O\left(\frac{K\sqrt{\log T}}{\varepsilon}\right)$. Our theoretical findings are corroborated by numerical evaluations on both synthetic and real-world data.

1 INTRODUCTION

The multi-armed bandit (MAB) problem provides a simple but powerful framework for sequential decision-making under uncertainty with bandit feedback, which has attracted a wide range of practical applications such as online advertising (Abe et al., 2003), product recommendations (Li et al., 2010), clinical trials (Tewari & Murphy, 2017), to name a few. Along with its broad applicability, however, there is an increasing concern of privacy risk in MAB due to its intrinsic dependence on users' feedback, which could leak users' sensitive information (Pan et al., 2019).

To alleviate the above concern, the notion of *differential privacy*, introduced by Dwork et al. (2006) in the field of computer science theory, has recently been adopted to design privacy-preserving bandit algorithms (see, e.g., Mishra & Thakurta (2015); Tossou & Dimitrakakis (2016); Shariff & Sheffet (2018)). Differential privacy (DP) provides a principled way to mathematically prove privacy guarantees against adversaries with arbitrary auxiliary information about users. To achieve this, a differentially private bandit algorithm typically relies on a well-tuned random noise to obscure each user's contribution to the output, depending on privacy levels ε, δ – smaller values lead to stronger protection but also suffer worse utility (i.e., regret). For example, the central server of a recommendation system can use random noise to perturb its statistics on each item after receiving feedback (i.e., clicks/ratings) from users. This is often termed as *central model* (Dwork et al., 2014), since the central server has the trust of its users and hence has a direct access to their raw data. Under this model, an optimal private MAB algorithm with a pure DP guarantee (i.e., when $\delta = 0$) is proposed in Sajed & Sheffet (2019), which only incurs an *additive* $O\left(\frac{K \log T}{\varepsilon}\right)$ term in the cumulative regret compared to the standard setting when privacy is not sought after (Auer, 2002). However, this high trust model is not always feasible in practice since users may not be willing to share their raw data directly to the server. This motivates to employ a *local model* (Kasiviswanathan

Table 1: Best-known performance of private MAB under different privacy models (K = number of arms, T = time horizon, Δ_a= reward gap of arm a w.r.t. best arm, $\varepsilon, \delta, \alpha$ = privacy parameters)

Trust Model	Privacy Guarantee	Best-Known Regret Bounds
Central	$(\varepsilon, 0)$-DP	$\Theta\left(\sum_{a\in[K]:\Delta_a>0} \frac{\log T}{\Delta_a} + \frac{K \log T}{\varepsilon}\right)$ (Sajed & Sheffet, 2019)
Local	$(\varepsilon, 0)$-DP	$\Theta\left(\frac{1}{\varepsilon^2} \sum_{a\in[K]:\Delta_a>0} \frac{\log T}{\Delta_a}\right)$ (Ren et al., 2020)
Distributed	(ε, δ)-DP	$O\left(\sum_{a:\Delta_a>0} \frac{\log T}{\Delta_a} + \frac{K \log T \sqrt{\log \frac{1}{\delta}}}{\varepsilon}\right)$ (Tenenbaum et al., 2021)
Distributed	$(\varepsilon, 0)$-DP	$\Theta\left(\sum_{a\in[K]:\Delta_a>0} \frac{\log T}{\Delta_a} + \frac{K \log T}{\varepsilon}\right)$ (Theorem 1)
Distributed	$O(\alpha, \frac{\alpha\varepsilon^2}{2})$-RDP	$O\left(\sum_{a\in[K]:\Delta_a>0} \frac{\log T}{\Delta_a} + \frac{K \sqrt{\log T}}{\varepsilon}\right)$ (Theorem 2)

et al., 2011) of trust, where DP is achieved without a trusted server as each user perturbs her data prior to sharing with the server. This ensures a stronger privacy protection, but leads to a high cost in utility due to large aggregated noise from all users. As shown in Ren et al. (2020), under the local model, private MAB algorithms have to incur a *multiplicative* $1/\varepsilon^2$ factor in the regret rather than the additive one in the central model.

In attempts to recover the same utility of central model while without a trustworthy server like the local model, an intermediate DP trust model called *distributed model* has gained an increasing interest, especially in the context of (federated) supervised learning (Kairouz et al., 2021b; Agarwal et al., 2021; Kairouz et al., 2021a; Girgis et al., 2021; Lowy & Razaviyayn, 2021). Under this model, each user first perturbs her data via a local randomizer, and then sends the randomized data to a *secure* computation function. This secure function can be leveraged to guarantee privacy through aggregated noise from distributed users. There are two popular secure computation functions: *secure aggregation* (Bonawitz et al., 2017) and *secure shuffling* (Bittau et al., 2017). The former often relies on cryptographic primitives to securely aggregate users' data so that the central server only learns the aggregated result, while the latter securely shuffle users' messages to hide their source. To the best of our knowledge, distributed DP model is far less studied in online learning as compared to supervised learning, with only known results for standard K-armed bandits in Tenenbaum et al. (2021), where secure shuffling is adopted. Despite being pioneer work, the results obtained in this paper have several limitations: (i) The privacy guarantee is obtained only for approximate DP ($\delta > 0$) – a stronger pure DP ($\delta = 0$) guarantee is not achieved; (ii) The cost of privacy is a multiplicative $\sqrt{\log(1/\delta)}$ factor away from that of central model, leading to a higher regret bound; (iii) The secure protocol works only for binary rewards (or communication intensive for real rewards).[1]

Our contributions. In this work, we design the *first* communication-efficient MAB algorithm that satisfies pure DP in the distributed model while attaining the same regret bound as in the central model (see Table 1). We overcome several key challenges that arise in the design and analysis of distributed DP algorithms for bandits. We now list the challenges and our proposed solutions below.

(a) Private and communication efficient algorithm design. Secure aggregation (SecAgg) works only in the integer domain due to an inherent modular operation (Bonawitz et al., 2017). Hence, leveraging this in bandits to achieve distributed DP with *real* rewards needs adopting data quantization, *discrete privacy noise* and modular summation arithmetic in the algorithm design. To this end, we take a batch version of the successive arm elimination technique as a building block of our algorithm, and on top of it, employ a privacy protocol tailored to discrete privacy noise and modular operation (see Algorithm 1). Instantiating the protocol at each user with *Pólya* random noise, we ensure that our algorithm satisfies pure DP in the distributed model. Moreover, the communication bits per-user scale only logarithmicaly with the number of participating users in each batch.

(b) Regret analysis under pure DP with SecAgg. While our pure DP guarantee exploits known results for discrete Laplace mechanism, the utility analysis gets challenging due to modular clipping of SecAgg. In fact, in supervised learning, no known convergence rate exists for SGD under pure DP with SecAgg (although the same is well-known under central model). This is because modular

[1]For more general linear bandits under distributed DP via shuffling, see Chowdhury & Zhou (2022b); Garcelon et al. (2022), which also have similar limitations.

clipping makes gradient estimates biased, and hence, standard convergence guarantees using unbiased estimates do not hold. In bandits, however, we work with zeroth order observations to build estimates of arms' rewards, and require high-confidence tight tail bounds for the estimates to analyse convergence. To this end, relying on tail properties of discrete Laplace and a careful analysis of modular operation, we prove a sublinear regret rate of our algorithm, which matches the optimal one in the central model, and thus, achieves the optimal rate under pure DP (see Theorem 1).

(c) Improved regret bound under RDP. While our main focus were to design the first bandit algorithm with pure distributed DP that achieves the same regret rate under central model, our template protocol is general enough to obtain different privacy guarantees by tuning the noise at each user. We demonstrate this by achieving *Rényi differential privacy* (RDP) (Mironov, 2017) using a Skellam random noise. RDP is a weaker notion of privacy compared to pure DP, but it is still stronger than approximate DP. It also provides a tighter privacy accounting for composition compared to approximate DP. This is particularly useful for bandit algorithms, when users may participate in multiple rounds, necessitating the need for privacy composition. Hence, we focus on RDP with SecAgg and and show that a tighter regret bound compared to pure DP can be achieved (see Theorem 2) by proving novel tail-bound for Skellam distribution. We support our theoretical findings with extensive numerical evaluation over bandit instances generated from both synthetic and real-life data.

Finally, our analysis technique is also general enough to recover best-known regrets under central and local DP models while only using *discrete* privacy noise (see Appendix H). This is important in practice since continuous Laplace noise might leak privacy on finite computers due to floating point arithmetic (Mironov, 2012), which is a drawback of existing central and local DP MAB algorithms.

2 PRELIMINARIES

In this section, we formally introduce the distributed differential privacy model in bandits. Before that we recall the learning paradigm in multi-armed bandits and basic differential privacy definitions.

Learning model and regret in MAB. At each time slot $t \in [T] := \{1, \ldots, T\}$, the agent (e.g., recommender system) selects an arm $a \in [K]$ (e.g., an advertisement), recommends it to a new user t and obtains an i.i.d reward r_t (e.g., a rating indicating how much she likes it), which is sampled from a distribution over $[0, 1]$ with mean given by μ_a. Let $a^* := \operatorname{argmax}_{a\in[K]} \mu_a$ be the arm with the highest mean and denote $\mu^* := \mu_{a^*}$ for simplicity. Let $\Delta_a := \mu^* - \mu_a$ be the gap of the expected reward between the optimal arm a^* and any other arm a. Further, let $N_a(t)$ be the total number of times that arm a has been recommended to first t users. The goal of the agent is to maximize its total reward, or equivalently to minimize the cumulative expected pseudo-regret, defined as

$$\mathbb{E}\left[\text{Reg}(T)\right] := T \cdot \mu^* - \mathbb{E}\left[\sum\nolimits_{t=1}^{T} r_t\right] = \mathbb{E}\left[\sum\nolimits_{a\in[K]} \Delta_a N_a(T)\right].$$

Differential privacy. Let $\mathcal{D} = [0, 1]$ be the data universe, and $n \in \mathbb{N}$ be the number of *unique* users. we say $D, D' \in \mathcal{D}^n$ are neighboring datasets if they only differ in one user's reward D_i for some $i \in [n]$. We have the following standard definition of differential privacy (Dwork et al., 2006).

Definition 1 (Differential Privacy)**.** *For $\varepsilon, \delta > 0$, a randomized mechanism $\mathcal{M}$ satisfies (ε, δ)-DP if for all neighboring datasets D, D' and all events $\mathcal{E}$ in the range of $\mathcal{M}$, we have*

$$\mathbb{P}\left[\mathcal{M}(D) \in \mathcal{E}\right] \leq e^{\varepsilon} \cdot \mathbb{P}\left[\mathcal{M}(D') \in \mathcal{E}\right] + \delta.$$

The special case of $(\varepsilon, 0)$-DP is often referred to as *pure differential privacy*, whereas, for $\delta > 0$, (ε, δ)-DP is referred to as *approximate differential privacy*. We also consider a related notion of privacy called Rényi differential privacy (RDP) Mironov (2017), which allows for a tighter composition compared to approximate differential privacy.

Definition 2 (Rényi Differential Privacy)**.** *For $\alpha > 1$, a randomized mechanism $\mathcal{M}$ satisfies $(\alpha, \varepsilon(\alpha))$-RDP if for all neighboring datasets D, D', we have $D_\alpha(\mathcal{M}(D), \mathcal{M}(D')) \leq \varepsilon(\alpha)$, where $D_\alpha(P, Q)$ is the Rényi divergence (of order α) of the distribution P from the distribution Q, and is given by $D_\alpha(P, Q) := \frac{1}{\alpha-1} \log\left(\mathbb{E}_{x\sim Q}\left[\left(\frac{P(x)}{Q(x)}\right)^{\alpha}\right]\right)$.*

Distributed differential privacy. A distributed bandit learning protocol $\mathcal{P} = (\mathcal{R}, \mathcal{S}, \mathcal{A})$ consists of three parts: (i) a (local) randomizer $\mathcal{R}$ at each user's side, (ii) an intermediate secure protocol

$\mathcal{S}$, and (iii) an analyzer $\mathcal{A}$ at the central server. Each user i first locally apply the randomizer $\mathcal{R}$ on its raw data (i.e., reward) D_i, and sends the randomized data to a secure computation protocol $\mathcal{S}$ (e.g., secure aggregation or shuffling). This intermediate secure protocol $\mathcal{S}$ takes a batch of users' randomized data and generates inputs to the central server, which utilizes an analyzer $\mathcal{A}$ to compute the output (e.g., action) using received messages from $\mathcal{S}$.

The secure computation protocol $\mathcal{S}$ has two main variations: *secure shuffling* and *secure aggregation*. Both of them essentially work with a batch of users' randomized data and guarantee that the central server cannot infer any individual's data while the total noise in the inputs to the analyzer provides a high privacy level. To adapt both into our MAB protocol, it is natural to divide participating users into batches. For each batch $b \in [B]$ with n_b users, the outputs of $\mathcal{S}$ is given by $\mathcal{S} \circ \mathcal{R}^{n_b}(D) := \mathcal{S}(\mathcal{R}(D_1), \ldots, \mathcal{R}(D_{n_b}))$. The goal is to guarantee that the the view of all B batches' outputs satisfy DP. To this end, we define a (composite) mechanism

$$\mathcal{M}_{\mathcal{P}} = (\mathcal{S} \circ \mathcal{R}^{n_1}, \ldots, \mathcal{S} \circ \mathcal{R}^{n_B}),$$

where each individual mechanism $\mathcal{S} \circ \mathcal{R}^{n_b}$ operates on n_b users' rewards, i.e., on a dataset from $\mathcal{D}^{n_b}$. With this notation, we have the following definition of distributed differential privacy.

Definition 3 (Distributed DP)**.** *A protocol $\mathcal{P} = (\mathcal{R}, \mathcal{S}, \mathcal{A})$ is said to satisfy DP (or RDP) in the distributed model if the mechanism $\mathcal{M}_{\mathcal{P}}$ satisfies Definition 1 (or Definition 2).*

In the central DP model, the privacy burden lies with a central server (in particular, analyzer $\mathcal{A}$), which needs to inject necessary random noise to achieve privacy. On the other hand, in the local DP model, each user's data is privatized by local randomizer $\mathcal{R}$. In contrast, in the distributed DP model, privacy without a trusted central server is achieved by ensuring that the inputs to the analyzer $\mathcal{A}$ already satisfy differential privacy. Specifically, by properly designing the intermediate protocol $\mathcal{S}$ and the noise level in the randomizer $\mathcal{R}$, one can ensure that the final added noise in the aggregated data over a batch of users matches the noise that would have otherwise been added in the central model by the trusted server. Through this, distributed DP model provides the possibility to achieve the same level of utility as the central model while without a trustworthy central server.

3 A Generic Algorithm for Private Bandits

In this section, we propose a generic algorithmic framework (Algorithm 1) for multi-armed bandits under the distributed privacy model.

Batch-based successive arm elimination. Our algorithm builds upon the classic idea of successive arm elimination (Even-Dar et al., 2006) with the additional incorporation of batches and a black-box protocol $\mathcal{P} = (\mathcal{R}, \mathcal{S}, \mathcal{A})$ to achieve distributed differential privacy. It divides the time horizon T into batches of exponentially increasing size and eliminates sub-optimal arms successively. To this end, for each active arm a at batch b, it first prescribes arm a to a *batch* of $l(b) = 2^b$ *new* users.[2] After pulling the prescribed action a, each user applies the local randomizer $\mathcal{R}$ to her reward and sends the randomized reward to the intermediary function $\mathcal{S}$, which runs a secure computation protocol (e.g., secure aggregation or secure shuffling) over the total $l(b)$ number of randomized rewards. Then, upon receiving the outputs of $\mathcal{S}$, the server applies the analyzer $\mathcal{A}$ to compute the the sum of rewards for batch b when pulling arm a (i.e., $R_a(b)$), which in turn gives the new mean estimate $\widehat{\mu}_a(b)$ of arm a after being divided by the total pulls $l(b)$. Then, upper and lower confidence bounds, $\text{UCB}_a(b)$ and $\text{LCB}_a(b)$, respectively, are computed around the mean estimate $\widehat{\mu}_a(b)$ with a properly chosen confidence width $\beta(b)$. Finally, after the iteration over all active arms in batch b (denoted by the set $\Phi(b)$), it adopts the standard arm elimination criterion to remove all obviously sub-optimal arms, i.e., it removes an arm a from $\Phi(b)$ if $\text{UCB}_a(b)$ falls below $\text{LCB}_{a'}(b)$ of any other arm $a' \in \Phi(b)$. It now only remains to design a distributed DP protocol $\mathcal{P}$, which is detailed as follows.

There is one key difference between our algorithm and the VB-SDP-AE algorithm in Tenenbaum et al. (2021). At the start of a batch, VB-SDP-AE uses all the past data to compute reward estimates. In contrast, we adopt the idea of *forgetting* and use only the data of the last completed batch to design the estimates. This small but important tweak allows us to accumulate "less" noise in reward estimates compared to VB-SDP-AE, which is key to achieve optimal regret.

[2] That is, if there are p active arms, a total of $p \cdot 2^b$ rounds will be consumed at batch b.

Algorithm 1 Private Batch-Based Successive Arm Elimination

Parameters: # arms K, Time horizon T, privacy level $\varepsilon > 0$, Confidence radii $\{\beta(b)\}_{b\geq 1}$
Initialize: Batch count $b=1$, Active arm set $\Phi(b)=\{1,\ldots,K\}$, Estimate $\widehat{\mu}_a(1)=0,\ \forall a\in[K]$
for batch $b=1,2,\ldots$ **do**
  Set batch size $l(b)=2^b$
  for each active arm $a \in \Phi(b)$ **do**
    for each new user i from 1 to $l(b)$ **do**
      Pull arm a and generate reward $r_a^i(b)$
      Send randomized data $y_a^i(b)=\mathcal{R}(r_a^i(b))$ to $\mathcal{S}$ // randomizer
      If total number of pulls reaches T, **exit**
    end for
    Send messages $\widehat{y}_a(b)=\mathcal{S}(\{y_a^i(b)\}_{1\leq i\leq l(b)})$ to $\mathcal{A}$ // secure computation
    Compute the sum of rewards $R_a(b)=\mathcal{A}(\widehat{y}_a(b))$ // analyzer
    Compute mean estimate $\widehat{\mu}_a(b)=R_a(b)/l(b)$
    Compute confidence bounds $\mathrm{UCB}_a(b)=\widehat{\mu}_a(b)+\beta(b)$ and $\mathrm{LCB}_a(b)=\widehat{\mu}_a(b)-\beta(b)$
  end for
  Update active set of arms: $\Phi(b+1)=\left\{a\in\Phi(b) : \mathrm{UCB}_a(b)\geq \max_{a'\in\Phi(b)} \mathrm{LCB}_{a'}(b)\right\}$
end for
Subroutine: Local Randomizer $\mathcal{R}$ (**Input:** $x_i \in [0,1]$, **Output:** y_i)
Require: precision $g \in \mathbb{N}$, modulo $m \in \mathbb{N}$, batch size $n \in \mathbb{N}$, privacy level ε
Encode x_i as $\widehat{x}_i = \lfloor x_i g \rfloor + \mathbf{Ber}(x_i g - \lfloor x_i g \rfloor)$
Generate discrete noise η_i (depending on n, ε, g) // random noise generator
Add noise and modulo clip $y_i = (\widehat{x}_i + \eta_i) \bmod m$
Subroutine: Secure Aggregation $\mathcal{S}$ (**Input:** $y_1,\ldots,y_n$, **Output:** $\widehat{y}$)
Require: modulo $m \in \mathbb{N}$
Securely compute $\widehat{y} = (\sum_{i=1}^n y_i) \bmod m$ // black-box function
Subroutine: Analyzer $\mathcal{A}$ (**Input:** $\widehat{y}$, **Output:** z)
Require: precision $g \in \mathbb{N}$, modulo $m \in \mathbb{N}$, batch size $n \in \mathbb{N}$, accuracy level $\tau \in \mathbb{R}$
if $\widehat{y} > ng + \tau$ **then**
  set $z = (\widehat{y} - m)/g$ // correction for underflow
else set $z = \widehat{y}/g$

Distributed DP protocol via discrete privacy noise. Inspired by Balle et al. (2020); Cheu & Yan (2021), we provide a general template protocol $\mathcal{P}$ for the distributed DP model, which relies only on *discrete privacy noise*.

Local randomizer $\mathcal{R}$ receives each user i's real-valued data x_i and encodes it as an integer via fixed-point encoding with precision $g > 0$ and randomized rounding. Then, it generates a discrete noise, which depends on the specific privacy-regret trade-off requirement (to be discussed later under specific mechanisms). Next, it adds the random noise to the encoded reward, clips the sum with modulo $m \in \mathbb{N}$ and sends the final integer y_i as input to secure computation function $\mathcal{S}$.

We mainly focus on secure aggregation (SecAgg) for $\mathcal{S}$ here.[3] SecAgg is treated as a black-box function as in previous work on supervised learning (Kairouz et al., 2021a), which implements the following procedure: given n users and their randomized messages $y_i \in \mathbb{Z}_m$ (i.e., integer in $\{0,1,\ldots,m-1\}$) obtained via $\mathcal{R}$, the SecAgg function $\mathcal{S}$ *securely* computes the modular sum of the n messages, $\widehat{y} = (\sum_{i=1}^n y_i) \bmod m$, while revealing no further information on individual messages to a potential attacker, ensuring that it is *perfectly secure*. Details of engineering implementations of SecAgg is beyond the scope of this paper, see Appendix G for a brief discussion on this.

The job of analyzer $\mathcal{A}$ is to compute the sum of rewards within a batch as accurately as possible. It uses an accuracy parameter $\tau \in \mathbb{R}$ and g to correct for possible underflow due to modular operation and bias due to encoding. To sum it up, the end goal of our protocol $\mathcal{P} = (\mathcal{R}, \mathcal{S}, \mathcal{A})$ is to ensure that it provides the required privacy protection while guaranteeing an output $z \approx \sum_{i=1}^n x_i$ with high probability, which is the key to our privacy and regret analysis in the following sections.

[3]Instead of SecAgg, we can also use secure shuffling as $\mathcal{S}$ (see Appendix D.2), since each has advantages over the other. The high-level idea is same for both techniques; i.e., to ensure that after receiving messages from $\mathcal{S}$, the server cannot distinguish each individual's message.

4 ACHIEVING PURE DP IN THE DISTRIBUTED MODEL

In this section, we show that Algorithm 1 achieves pure-DP in the distributed DP model via secure aggregation. To do so, we need to carefully determine the amount of (discrete) noise in $\mathcal{R}$ so that the total noise in a batch provides $(\varepsilon, 0)$-DP. One natural choice is the discrete Laplace noise.

Definition 4 (Discrete Laplace Distribution)**.** *Let $b > 0$. A random variable X has a discrete Laplace distribution with scale parameter b, denoted by $\boldsymbol{Lap}_{\mathbb{Z}}(b)$, if it has a p.m.f. given by*

$$\forall x \in \mathbb{Z}, \quad \mathbb{P}[X = x] = \frac{e^{1/b} - 1}{e^{1/b} + 1} \cdot e^{-|x|/b}.$$

A key property of discrete Laplace that we will use is its *infinite divisibility*, which allows us to simulate it in a distributed way (Goryczka & Xiong, 2015, Theorem 5.1).

Fact 1 (Infinite Divisibility of Discrete Laplace)**.** *A random variable X has a Pólya distribution with parameters $r > 0, \beta \in [0, 1]$, denoted by $\boldsymbol{Pólya}(r, \beta)$, if it has a p.m.f. given by*

$$\forall x \in \mathbb{N}, \quad \mathbb{P}[X = x] = \frac{\Gamma(x + r)}{x!\Gamma(r)} \beta^x (1 - \beta)^r.$$

Now, for any $n \in \mathbb{N}$, let $\{\gamma_i^+, \gamma_i^-\}_{i \in [n]}$ be $2n$ i.i.d samples[4] from $\boldsymbol{Pólya}(1/n, e^{-1/b})$, then the random variable $\sum_{i=1}^n (\gamma_i^+ - \gamma_i^-)$ is distributed as $\boldsymbol{Lap}_{\mathbb{Z}}(b)$.

Armed with the above fact and the properties of discrete Laplace noise (see Fact 3 in Appendix K), we are able to obtain the following main theorem, which shows that the same regret as in the central model is achieved under the distributed model via SecAgg.

Theorem 1 (Pure-DP via SecAgg)**.** *Fix $\varepsilon > 0$ and $T \in \mathbb{N}$. For each batch b, let noise for the i-th user in the batch be $\eta_i = \gamma_i^+ - \gamma_i^-$, where $\gamma_i^+, \gamma_i^- \overset{i.i.d.}{\sim} \boldsymbol{Pólya}(1/n, e^{-\varepsilon/g})$, set $n = l(b), g = \lceil \varepsilon \sqrt{n} \rceil$, $\tau = \lceil \frac{g}{\varepsilon} \log(2T) \rceil$ and $m = ng + 2\tau + 1$. Then, Algorithm 1 achieves $(\varepsilon, 0)$-DP in the distributed model. Moreover, setting $\beta(b) = O\left(\sqrt{\frac{\log(|\Phi(b)| b^2 T)}{2l(b)}} + \frac{2\log(|\Phi(b)| b^2 T)}{\varepsilon l(b)}\right)$, it enjoys expected regret*

$$\mathbb{E}[Reg(T)] = O\left(\sum\nolimits_{a \in [K]: \Delta_a > 0} \frac{\log T}{\Delta_a} + \frac{K \log T}{\varepsilon}\right).$$

Theorem 1 achieves optimal regret under pure DP. Theorem 1 achieves the same regret bound as the one achieved in Sajed & Sheffet (2019) under the central trust model with *continuous* Laplace noise. Moreover, it matches the lower bound obtained under pure DP in Shariff & Sheffet (2018), indicating the bound is indeed tight. Note that, we achieve this rate under distributed trust model – a stronger notion of privacy protection than the central model – while using only discrete noise.

Communication bits. Algorithm 1 needs to communicate $O(\log m)$ bits per user to the secure protocol $\mathcal{S}$, i.e., communicating bits scales logarithmically with the batch size. In contrast, the number of communication bits required in existing distributed DP bandit algorithms that work with real-valued rewards (as we consider here) scale polynomially with the batch size (Chowdhury & Zhou, 2022b; Garcelon et al., 2022).

Remark 1 (Pure DP via Secure Shuffling)**.** *It turns out that one can achieve same privacy and regret guarantees (orderwise) using a* relaxed *SecAgg protocol. Building on this result, we also establish pure DP under shuffling while again maintaining the same regret bound as the central model (see Theorem 3 in Appendix D.2). This improves the state-of-the-art result for MAB with shuffling (Tenenbaum et al., 2021) in terms of both privacy and regret.*

5 ACHIEVING RDP IN THE DISTRIBUTED MODEL

A natural question to ask is whether one can get a better regret performance by sacrificing a small amount of privacy. We consider the notion of RDP (see Definition 2), which is a weaker notion of privacy than pure DP. However, it avoids the possible catastrophic privacy failure in approximate DP, and also provides a tighter privacy accounting for composition (Mironov, 2017).

[4]One can sample from **Pólya** as follows. First, sample $\lambda \sim$ **Gamma**$(r, \beta/(1-\beta))$ and then use it to sample $X \sim$ **Poisson**(λ), which is known to follow **Pólya**(r, β) distribution (Goryczka & Xiong, 2015).

To achieve RDP guarantee using discrete noise, we consider the Skellam distribution – which has recently been introduced in private federated learning (Agarwal et al., 2021). A key challenge in the regret analysis of our bandit algorithm is to characterize the tail property of Skellam distribution. This is different from federated learning, where characterizing the variance renders sufficient. In Proposition 1, we prove that Skellam has sub-exponential tails, which not only is the key to our regret analysis, but could also be of independent interest. Below is the formal definition of Skellam.

Definition 5 (Skellam Distribution). *A random variable X has a Skellam distribution with mean μ and variance σ^2, denoted by $\boldsymbol{Sk}(\mu, \sigma^2)$, if it has a probability mass function given by*

$$\forall x \in \mathbb{Z}, \quad \mathbb{P}\left[X = x\right] = e^{-\sigma^2} I_{x-\mu}(\sigma^2),$$

where $I_\nu(\cdot)$ is the modified Bessel function of the first kind.

To sample from Skellam distribution, one can rely on existing procedures for Poisson samples. This is because if $X = N_1 - N_2$, where $N_1, N_2 \overset{\text{i.i.d.}}{\sim} \textbf{Poisson}(\sigma^2/2)$, then X is $\textbf{Sk}(0, \sigma^2)$ distributed. Moreover, due to this fact, Skellam is *closed* under summation, i.e., if $X_1 \sim \textbf{Sk}(\mu_1, \sigma_1^2)$ and $X_2 \sim \textbf{Sk}(\mu_2, \sigma_2^2)$, then $X_1 + X_2 \sim \textbf{Sk}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$.

Proposition 1 (Sub-exponential Tail of Skellam). *Let $X \sim \boldsymbol{Sk}(0, \sigma^2)$. Then, X is $(2\sigma^2, \frac{\sqrt{2}}{2})$-sub-exponential. Hence, for any $p \in (0, 1]$, with probability at least $1 - p$,*

$$|X| \leq 2\sigma\sqrt{\log(2/p)} + \sqrt{2}\log(2/p).$$

With the above result, we can establish the following privacy and regret guarantee of Algorithm 1.

Theorem 2 (RDP via SecAgg). *Fix $\varepsilon > 0$, $T \in \mathbb{N}$ and a scaling factor $s \geq 1$. For each batch b, let noise for the i-th user be $\eta_i \sim \boldsymbol{Sk}(0, \frac{g^2}{n\varepsilon^2})$, set $n = l(b), g = \lceil s\varepsilon\sqrt{n}\rceil$, $\tau = \lceil \frac{2g}{\varepsilon}\sqrt{\log(2T)} + \sqrt{2}\log(2T)\rceil$ and $m = ng + 2\tau + 1$. Then, Algorithm 1 achieves $(\alpha, \widehat{\varepsilon}(\alpha))$-RDP in the distributed model for all $\alpha = 2, 3, \ldots$, with $\widehat{\varepsilon}(\alpha) = \frac{\alpha\varepsilon^2}{2} + \min\left\{\frac{(2\alpha-1)\varepsilon^2}{4s^2} + \frac{3\varepsilon}{2s^3}, \frac{3\varepsilon^2}{2s}\right\}$. Moreover, setting $\beta(b) = O\left(\sqrt{\frac{\log(|\Phi(b)|b^2T)}{2l(b)}} + \frac{(1+1/s)\log(|\Phi(b)|b^2T)}{\varepsilon l(b)}\right)$, it enjoys the expected regret*

$$\mathbb{E}\left[\textit{Reg}(T)\right] = O\left(\sum\nolimits_{a\in[K]:\Delta_a>0} \frac{\log T}{\Delta_a} + \frac{K\sqrt{\log T}}{\varepsilon} + \frac{K\log T}{s\varepsilon}\right).$$

Privacy-Regret-Communication Trade-off. Observe that the scaling factor s allows us to achieve different trade-offs. If s increases, both privacy and regret performances improve. In fact, for a sufficiently large value of s, the third term in the regret bound becomes sufficiently small, and we obtain an improved regret bound compared to Theorem 1. Moreover, the RDP privacy guarantee improves to $\widehat{\varepsilon}(\alpha) \approx \frac{\alpha\varepsilon^2}{2}$, which is the standard RDP rate for Gaussian mechanism (Mironov, 2017). However, a larger s leads to an increase of communicating bits per user, but only grows logarithmically, since Algorithm 1 needs to communicate $O(\log m)$ bits to the secure protocol $\mathcal{S}$. See Appendix B for further discussion on the communication cost.

RDP to Approximate DP. To shed more insight on Theorem 2, we convert our RDP guarantee to approximate DP for a sufficiently large s. It holds that under the setup of Theorem 2, for sufficiently large s, one can achieve $(O(\varepsilon), \delta)$-DP with regret $O\big(\sum_{a:\Delta_a>0} \frac{\log T}{\Delta_a} + \frac{K\sqrt{\log T\log(1/\delta)}}{\varepsilon}\big)$ (via Lemma 10 in Appendix K). Implication of this conversion is three-fold. **First**, this regret bound is $O(\sqrt{\log T})$ factor tighter than that achieved by Tenenbaum et al. (2021) using a shuffle protocol. **Second**, it yields a better regret performance compared to the bound achieved under $(\varepsilon, 0)$-DP in Theorem 1 when the privacy budget $\delta > 1/T$. This observation is consistent with the fact that a weaker privacy guarantee typically warrants a better utility bound. **Third**, this conversion via RDP also yields a gain of $O(\sqrt{\log(1/\delta)})$ in the regret when dealing with privacy composition (e.g., when participating users across different batches are *not unique*) as compared to Tenenbaum et al. (2021) that only relies on approximate DP (see Appendix I for details). This results from the fact that RDP provides a tighter composition compared to approximate DP.

Remark 2 (Achieving RDP with discrete Gaussian). *One can also achieve RDP using discrete Gaussian noise (Canonne et al., 2020). Here, we work with Skellam noise since it is closed under summation and enjoys efficient sampling procedure as opposed to discrete Gaussian (Agarwal et al., 2021). Nevertheless, as a proof of flexibility of our proposed framework, we show in Appendix F that Algorithm 1 with discrete Gaussian noise can guarantee RDP with a similar regret bound.*

6 KEY TECHNIQUES: OVERVIEW

Now, we provide an overview of the key techniques behind our privacy and regret guarantees. We show that the results of Theorem 1 and 2 can be obtained via a clean generic analytical framework, which not only covers the analysis of distributed pure DP/RDP with SecAgg, but also offers a unified view of private MAB under central, local and distributed DP models.

As in many private learning algorithms, the key is to characterize the impact of added privacy noise on the utility. In our case, this reduces to capturing the tail behavior of total noise $n_a(b) := R_a(b) - \sum_{i=1}^{l(b)} r_a^i(b)$ added at each batch b for each active arm a. The following lemma gives a generic regret bound of our algorithm under mild tail assumptions on $n_a(b)$.

Lemma 1 (Generic regret)**.** *Let there exist constants $\sigma, h > 0$ such that, with probability $\geq 1-p$, $|n_a(b)| \leq \mathcal{N} := O\left(\sigma\sqrt{\log(KT/p)} + h\log(KT/p)\right)$ for all $b \geq 1$, $a \in [K]$. Then, setting confidence radius $\beta(b) = O\left(\sqrt{\log(KT/p)/l(b)} + \mathcal{N}/l(b)\right)$ and $p = 1/T$, Algorithm 1 enjoys expected regret*

$$\mathbb{E}\left[Reg(T)\right] = O\Big(\sum\nolimits_{a\in[K]:\Delta_a>0} \frac{\log T}{\Delta_a} + K\sigma\sqrt{\log T} + Kh\log T\Big).$$

An acute reader may note that the bound $\mathcal{N}$ on the noise is the tail bound of sub-exponential distribution and it reduces to the bound for sub-Gaussian tail if $h = 0$. Our SecAgg protocol $\mathcal{P}$ with discrete Laplace noise (as in Theorem 1) satisfy this bound with $\sigma = \sqrt{2}/\varepsilon, h = 1/\varepsilon$. Similarly, our protocol with Skellam noise (as in Theorem 2) satisfy this bound with $\sigma = O(1/\varepsilon), h = 1/(s\varepsilon)$. Therefore, we can build on the above general result to directly obtain our regret bounds. In the following, we present the high-level idea behind privacy and regret analysis in distributed DP model.

Privacy. For distributed DP, by definition, the view of the server during the entire algorithm needs to be private. Since each user only contributes once,[5] by parallel-composition of DP, it suffices to ensure that each view $\widehat{y}_a(b)$ (line 11 in Algorithm 1) is private. To this end, under SecAgg, the distribution of $\widehat{y}_a(b)$ can be simulated via $(\sum_i y_i) \bmod m$, which further reduces to $(\sum_i \widehat{x}_i + \eta_i) \bmod m$ by the distributive property of modular sum. Now, consider a mechanism $\mathcal{M}$ that accepts an input dataset $\{\widehat{x}_i\}_i$ and outputs $\sum_i(\widehat{x}_i + \eta_i)$. By post-processing, it suffices to show that $\mathcal{M}$ satisfies pure DP or RDP. To this end, the variance σ_{tot}^2 of the total noise $\sum_i \eta_i$ needs to scale with the sensitivity of $\sum_i \widehat{x}_i$. Thus, each user within a batch only needs to add a proper noise with variance of σ_{tot}^2/n. Finally, by the particular distribution properties of the noise, one can show that $\mathcal{M}$ is pure DP or RDP, and hence, obtain the privacy guarantees.

Regret. Thanks to Lemma 1, we only need to focus on the tail of $n_a(b)$. To this end, fix any batch b and arm a. We have $\widehat{y} = \widehat{y}_a(b)$, $x_i = r_a^i(b)$, $n = l(b)$ for $\mathcal{P}$ in Algorithm 1 and we need to establish that with probability at least $1-p$, for some σ and h,

$$\left|\mathcal{A}(\widehat{y}) - \sum\nolimits_i x_i\right| \leq O\left(\sigma\sqrt{\log(1/p)} + h\log(1/p)\right). \tag{1}$$

To get the bound, inspired by Balle et al. (2020); Cheu & Yan (2021), we divide the LHS into Term (i) $= |\mathcal{A}(\widehat{y}) - \sum_i \widehat{x}_i/g|$ and Term (ii) $= |\sum_i \widehat{x}_i/g - \sum_i x_i|$, where Term (i) captures the error due to privacy noise and modular operation, while Term (ii) captures the error due to random rounding. In particular, Term (ii) can be easily bounded via sub-Gaussian tail since the noise is bounded. Term (i) needs care for the possible underflow due to modular operation by considering two different cases (see line 28-30 in Algorithm 1). In both cases, one can show that Term (i) is upper bounded by τ/g with high probability, where τ is the tail bound on the total privacy noise $\sum_i \eta_i$. Thus, depending on particular privacy noise and parameter choices, one can find σ and h such that equation 1 holds, and hence, obtain the corresponding regret bound by Lemma 1.

Remark 3. *As a by-product of our generic analysis technique, Algorithm 1 and privacy protocol $\mathcal{P}$ along with Lemma 1 provide a new and structured way to design and analyze private MAB algorithms under central and local models with discrete private noise (see Appendix H for details). This enables us to reap the benefits of working with discrete noise (e.g., finite-computer representations, bit communications) in all three trust models (central, local and distributed).*

[5]This assumption that users only contributes once is adopted in nearly all previous works for privacy analysis in bandits. We also provide a privacy analysis for returning users via RDP, see Appendix I.

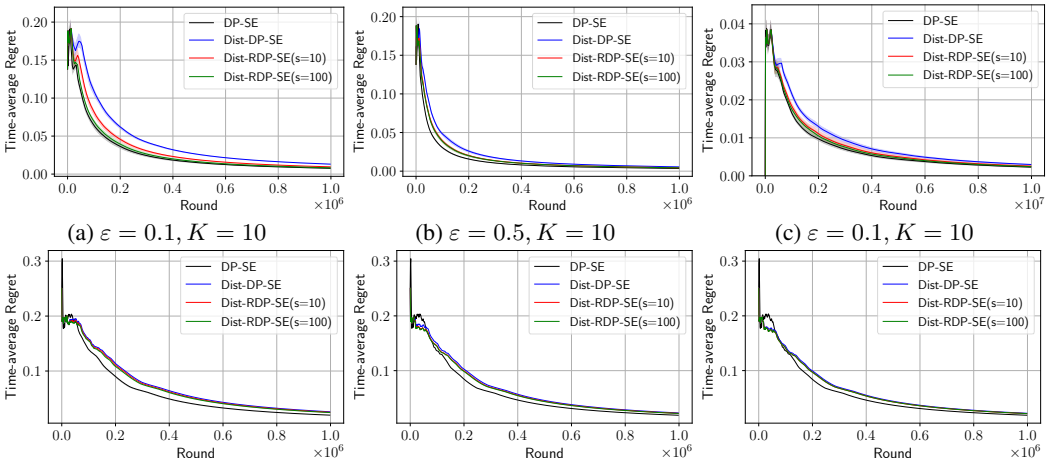

(a) $\varepsilon = 0.1, K = 10$ (b) $\varepsilon = 0.5, K = 10$ (c) $\varepsilon = 0.1, K = 10$

(d) $\varepsilon = 1, K = 50$ (e) $\varepsilon = 5, K = 50$ (f) $\varepsilon = 10, K = 50$

Figure 1: Comparison of time-average regret for Dist-DP-SE, Dist-RDP-SE, and DP-SE. **Top:** Synthetic Gaussian bandit instances with (a, b) large reward gap (easy instance) and (c) small reward gap (hard instance). **Bottom:** Bandit instances generated from MSLR-WEB10K learning to rank dataset.

7 Simulation Results

We empirically evaluate the regret performance of our successive elimination scheme with SecAgg protocol (Algorithm 1) under distributed trust model, which we abbreviate as Dist-DP-SE and Dist-RDP-SE when the randomizer $\mathcal{R}$ is instantiated with Pólya noise (for pure DP) and Skellam noise (for RDP), respectively. We compare them with the DP-SE algorithm of Sajed & Sheffet (2019) that achieves optimal regret under pure DP in the central model, but works only with *continuous* Laplace noise. We fix confidence level $p = 0.1$ and study comparative performances under varying privacy levels ($\varepsilon < 1$ for synthetic data, $\varepsilon \geq 1$ for real data). We plot time-average regret $\mathrm{Reg}(T)/T$ in Figure 1 by averaging results over 20 randomly generated bandit instances.

Bandit instances. In the *top panel*, similar to Vaswani et al. (2020), we consider easy and hard MAB instances with $K = 10$ arms: in the former, arm means are sampled uniformly in $[0.25, 0.75]$, while in the latter, those are sampled in $[0.45, 0.55]$. We consider real rewards – sampled from Gaussian distribution with aforementioned means and projected to $[0, 1]$. In the *bottom panel*, we generate bandit instances from Microsoft Learning to Rank dataset MSLR-WEB10K (Qin & Liu, 2013). The dataset consists of 1,200,192 rows and 138 columns, where each row corresponds to a query-url pair. The first column is relevance label 0, 1, ... , 4 of the pair, which we take as rewards. The second column denotes the query id, and the rest 136 columns denote contexts of a query-url pair. We cluster the data by running K-means algorithm with $K = 50$. We treat each cluster as a bandit arm with mean reward as the empirical mean of the individual ratings in the cluster. This way, we obtain a bandit setting with number of arms $K = 50$.

Observations. We observe that as T becomes large, the regret performance of Dist-DP-SE matches the regret of DP-SE. The slight gap in small T regime is the cost that we pay to achieve distributed privacy using discrete noise without access to a trusted server (for higher ε value, this gap is even smaller). In addition, we find that a relatively small scaling factor ($s = 10$) provides a considerable gain in regret under RDP compared to pure DP, especially when ε is small (i.e., when the cost of privacy is not dominated by the non-private part of regret). The experimental findings are consistent with our theoretical results. Here, we note that our simulations are proof-of-concept only and we did not tune any hyperparameters. More details and additional plots are given in Appendix J.

Concluding remarks. We show that MAB under distributed trust model can achieve pure DP while maintaining the same regret under central model. In addition, RDP is also achieved in MAB under distributed trust model for the first time. Both results are obtained via a unified algorithm design and performance analysis. More importantly, our work also opens the door to a promising and interesting research direction – private online learning with distributed DP guarantees, including contextual bandits and reinforcement learning.

REFERENCES

Naoki Abe, Alan W Biermann, and Philip M Long. Reinforcement learning with immediate rewards and linear hypotheses. *Algorithmica*, 37(4):263–293, 2003.

Naman Agarwal and Karan Singh. The price of differential privacy for online learning. In *International Conference on Machine Learning*, pp. 32–40. PMLR, 2017.

Naman Agarwal, Peter Kairouz, and Ziyu Liu. The skellam mechanism for differentially private federated learning. *Advances in Neural Information Processing Systems*, 34, 2021.

Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.

Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. The privacy blanket of the shuffle model. In *Annual International Cryptology Conference*, pp. 638–667. Springer, 2019.

Borja Balle, James Bell, Adria Gascón, and Kobbi Nissim. Private summation in the multi-message shuffle model. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 657–676, 2020.

Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In *Annual International Cryptology Conference*, pp. 451–468. Springer, 2008.

James Henry Bell, Kallista A Bonawitz, Adrià Gascón, Tancrède Lepoint, and Mariana Raykova. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1253–1269, 2020.

Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pp. 441–459, 2017.

Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pp. 635–658. Springer, 2016.

Clément L Canonne, Gautam Kamath, and Thomas Steinke. The discrete gaussian for differential privacy. *Advances in Neural Information Processing Systems*, 33:15676–15688, 2020.

TH Hubert Chan, Elaine Shi, and Dawn Song. Optimal lower bound for differentially private multi-party aggregation. In *European Symposium on Algorithms*, pp. 277–288. Springer, 2012.

Wei-Ning Chen, Christopher A Choquette Choo, Peter Kairouz, and Ananda Theertha Suresh. The fundamental price of secure aggregation in differentially private federated learning. In *International Conference on Machine Learning*, pp. 3056–3089. PMLR, 2022a.

Wei-Ning Chen, Ayfer Ozgur, and Peter Kairouz. The poisson binomial mechanism for unbiased federated learning with secure aggregation. In *International Conference on Machine Learning*, pp. 3490–3506. PMLR, 2022b.

Xiaoyu Chen, Kai Zheng, Zixin Zhou, Yunchang Yang, Wei Chen, and Liwei Wang. (locally) differentially private combinatorial semi-bandits. In *International Conference on Machine Learning*, pp. 1757–1767. PMLR, 2020.

Albert Cheu and Chao Yan. Pure differential privacy from secure intermediaries. *arXiv preprint arXiv:2112.10032*, 2021.

Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 375–403. Springer, 2019.

Albert Cheu, Matthew Joseph, Jieming Mao, and Binghui Peng. Shuffle private stochastic convex optimization. *arXiv preprint arXiv:2106.09805*, 2021.

Sayak Ray Chowdhury and Xingyu Zhou. Differentially private regret minimization in episodic markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 6375–6383, 2022a.

Sayak Ray Chowdhury and Xingyu Zhou. Shuffle private linear contextual bandits. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 3984–4009. PMLR, 17–23 Jul 2022b.

Sayak Ray Chowdhury, Xingyu Zhou, and Ness Shroff. Adaptive control of differentially private linear quadratic systems. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pp. 485–490. IEEE, 2021.

Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, Naval Research Lab Washington DC, 2004.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006.

Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.

Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2468–2479. SIAM, 2019.

Eyal Even-Dar, Shie Mannor, Yishay Mansour, and Sridhar Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.

Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 954–964. IEEE, 2022.

Evrard Garcelon, Vianney Perchet, Ciara Pike-Burke, and Matteo Pirotta. Local differential privacy for regret minimization in reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

Evrard Garcelon, Kamalika Chaudhuri, Vianney Perchet, and Matteo Pirotta. Privacy amplification via shuffling for linear contextual bandits. In *International Conference on Algorithmic Learning Theory*, pp. 381–407. PMLR, 2022.

Badih Ghazi, Noah Golowich, Ravi Kumar, Pasin Manurangsi, Rasmus Pagh, and Ameya Velingker. Pure differentially private summation from anonymous messages. *arXiv preprint arXiv:2002.01919*, 2020a.

Badih Ghazi, Pasin Manurangsi, Rasmus Pagh, and Ameya Velingker. Private aggregation from fewer anonymous messages. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 798–827. Springer, 2020b.

Antonious Girgis, Deepesh Data, Suhas Diggavi, Peter Kairouz, and Ananda Theertha Suresh. Shuffled model of differential privacy in federated learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 2521–2529. PMLR, 2021.

Slawomir Goryczka and Li Xiong. A comprehensive comparison of multiparty secure additions with differential privacy. *IEEE transactions on dependable and secure computing*, 14(5):463–477, 2015.

Christopher Hillar and Andre Wibisono. Maximum entropy distributions on graphs. *arXiv preprint arXiv:1301.3321*, 2013.

Peter Kairouz, Ziyu Liu, and Thomas Steinke. The distributed discrete gaussian mechanism for federated learning with secure aggregation. In *International Conference on Machine Learning*, pp. 5201–5212. PMLR, 2021a.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021b.

Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pp. 661–670, 2010.

Chonghua Liao, Jiafan He, and Quanquan Gu. Locally differentially private reinforcement learning for linear mixture markov decision processes. *arXiv preprint arXiv:2110.10133*, 2021.

Andrew Lowy and Meisam Razaviyayn. Private federated learning without a trusted server: Optimal algorithms for convex losses. *arXiv preprint arXiv:2106.09779*, 2021.

Paul Luyo, Evrard Garcelon, Alessandro Lazaric, and Matteo Pirotta. Differentially private exploration in reinforcement learning with linear representation. *arXiv preprint arXiv:2112.01585*, 2021.

Ilya Mironov. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pp. 650–661, 2012.

Ilya Mironov. Rényi differential privacy. In *IEEE 30th computer security foundations symposium (CSF)*, pp. 263–275. IEEE, 2017.

Nikita Mishra and Abhradeep Thakurta. (nearly) optimal differentially private stochastic multi-arm bandits. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pp. 592–601, 2015.

Xinlei Pan, Weiyao Wang, Xiaoshuai Zhang, Bo Li, Jinfeng Yi, and Dawn Song. How you act tells a lot: Privacy-leaking attack on deep reinforcement learning. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 368–376, 2019.

Tao Qin and Tie-Yan Liu. Introducing LETOR 4.0 datasets. *CoRR*, abs/1306.2597, 2013. URL http://arxiv.org/abs/1306.2597.

Wenbo Ren, Xingyu Zhou, Jia Liu, and Ness B Shroff. Multi-armed bandits with local differential privacy. *arXiv preprint arXiv:2007.03121*, 2020.

Touqir Sajed and Or Sheffet. An optimal private stochastic-mab algorithm based on optimal private stopping rule. In *International Conference on Machine Learning*, pp. 5579–5588. PMLR, 2019.

Roshan Shariff and Or Sheffet. Differentially private contextual linear bandits. *Advances in Neural Information Processing Systems*, 31, 2018.

Youming Tao, Yulian Wu, Peng Zhao, and Di Wang. Optimal rates of (locally) differentially private heavy-tailed multi-armed bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 1546–1574. PMLR, 2022.

Jay Tenenbaum, Haim Kaplan, Yishay Mansour, and Uri Stemmer. Differentially private multi-armed bandits in the shuffle model. *Advances in Neural Information Processing Systems*, 34, 2021.

Ambuj Tewari and Susan A Murphy. From ads to interventions: Contextual bandits in mobile health. In *Mobile Health*, pp. 495–517. Springer, 2017.

Aristide Charles Yedia Tossou and Christos Dimitrakakis. Achieving privacy in the adversarial multi-armed bandit. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.

Aristide CY Tossou and Christos Dimitrakakis. Algorithms for differentially private multi-armed bandits. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

Sharan Vaswani, Abbas Mehrabian, Audrey Durand, and Branislav Kveton. Old dog learns new tricks: Randomized ucb for bandit problems. In *International Conference on Artificial Intelligence and Statistics*, pp. 1988–1998. PMLR, 2020.

Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

Giuseppe Vietri, Borja Balle, Akshay Krishnamurthy, and Steven Wu. Private reinforcement learning with pac and regret guarantees. In *International Conference on Machine Learning*, pp. 9754–9764. PMLR, 2020.

Kun Wang, Jing Dong, Baoxiang Wang, and Shuai Li. Cascading bandit under differential privacy. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4418–4422. IEEE, 2022.

Huiming Zhang and Song Xi Chen. Concentration inequalities for statistical inference. *arXiv preprint arXiv:2011.02258*, 2020.

Kai Zheng, Tianle Cai, Weiran Huang, Zhenguo Li, and Liwei Wang. Locally differentially private (contextual) bandits learning. *Advances in Neural Information Processing Systems*, 33:12300–12310, 2020.

Xingyu Zhou. Differentially private reinforcement learning with linear function approximation. *Proc. ACM Meas. Anal. Comput. Syst.*, 6(1), Feb 2022.

Xingyu Zhou and Jian Tan. Local differential privacy for bayesian optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):11152–11159, May 2021.