Differentially Private Matroid Bandits for Online Network Topology Optimization

Youming Tao
TU Berlin
Berlin, Germany

Shuzhen Chen
Ocean University of China
Qingdao, China

Kai Li
TU Berlin
Berlin, Germany

Dongxiao Yu
Shandong University
Qingdao, China

Xiuzhen Cheng
Shandong University
Qingdao, China

Falko Dressler
TU Berlin
Berlin, Germany

ABSTRACT

Efficient and adaptive network topology design is fundamental in communication systems. A key problem is to construct low-cost structures like minimum spanning trees (MSTs) from uncertain and stochastic link measurements. However, topology and link data often reveal sensitive attributes, including location and traffic patterns, posing privacy risks. We propose Differentially Private Matroid Bandits, an online learning framework for consecutive MST construction with formal privacy guarantees. Modeling the problem as stochastic optimization over a matroid, we develop algorithms that balance exploration and exploitation while ensuring differential privacy in both central and local models. We establish regret upper bounds capturing the trade-offs among privacy, accuracy, and structural complexity. Experiments on real network data show our approach achieves near-optimal performance with strong privacy protection.

KEYWORDS

Network topology optimization, Minimum spanning tree, Matroid, bandit learning, differential privacy

ACM Reference Format:
Youming Tao, Shuzhen Chen, Kai Li, Dongxiao Yu, Xiuzhen Cheng, and Falko Dressler. 2025. Differentially Private Matroid Bandits for Online Network Topology Optimization. In *Workshop on Mobility in the Evolving Internet Architecture (MobiArch '25), November 4–8, 2025, Hong Kong, China.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3737897.3767290

Correspondence to tao@ccs-labs.org. This work was supported in part by the National Science Foundation of China (NSFC) under Grant 623B2068.

MobiArch '25, November 4–8, 2025, Hong Kong, China

ACM ISBN 979-8-4007-1974-5/25/11
https://doi.org/10.1145/3737897.3767290

1 INTRODUCTION

Efficient communication topologies are essential to networked systems, including wireless sensor networks [7, 18], ad-hoc and mobile networks [17, 26, 46], satellite constellations [21, 25], and data center interconnects [10, 36]. Among various design primitives, the minimum spanning tree (MST) serves as a key structure for enabling multi-hop routing [14], data aggregation [19], and broadcast dissemination [1]. In these systems, edge weights used for MST computation often represent application-specific connectivity metrics, such as energy consumption, signal strength, or latency in communication networks, or interaction frequency or social proximity in social networks. Revealing these weights or the resulting MST can expose sensitive structural information [4, 23, 38]. For instance, in wireless or IoT deployments (e.g., smart buildings or battlefield networks), an MST may disclose device proximity, relay nodes, or aggregation points, creating opportunities for traffic analysis or targeted disruption. In organizational graphs, spanning tree structures may reveal implicit hierarchies or access patterns; in social networks, they may highlight tightly-knit communities or influential intermediaries. Prior work has shown that even anonymized graphs are vulnerable to re-identification attacks via topological patterns [38]. These risks underscore the need for privacy-preserving topology construction algorithms, particularly in settings where network structures reflect sensitive operational data. Differential privacy (DP) [12] offers a rigorous and widely adopted framework for private data analysis. Several works have studied differentially private graph algorithms [22], including for MST computation [28, 29, 32, 34]. However, these methods are fundamentally *offline*: they assume access to the full graph and all edge weights in advance, and return a one-shot output. This assumption breaks down in practical systems, where link quality is uncertain, stochastic, or observable only through repeated interactions (e.g., via packet success rates or round-trip time measurements). In such settings, MST construction becomes an online learning problem: the system must sequentially explore edges and build increasingly accurate and efficient topologies over time. This raises a natural question: *Can we design algorithms that learn to construct minimum-cost spanning trees through sequential interactions, while ensuring rigorous differential privacy guarantees?*

Our Contribution. We address this question by studying the problem of matroid bandit optimization (MBO), focusing on MST as a canonical application. The matroid bandit framework [20] models stochastic combinatorial optimization where feasible actions form the independent sets of a matroid, a structure generalizing spanning trees, matchings, and linear independence. In our setting, each network edge corresponds to a stochastic base arm, and each feasible spanning tree corresponds to a super arm (i.e., a matroid basis). At each round, the algorithm selects a spanning tree, observes stochastic feedback on its edges, and aims to minimize cumulative regret under DP constraints. Our main contributions are: (1) We formalize private MBO problem under both central and local DP (CDP/LDP) models; (2) We develop new *anytime* privacy-preserving algorithms that achieve near-optimal regret bounds. Our methods combine the optimistic matroid maximization principle with tailored noise mechanisms and confidence bounds for matroid structures; (3) We apply our framework to online private MST construction in stochastic networks, and empirically demonstrate that accurate, private topologies can be learned over time even under strong privacy constraints.

Related Work on Matroid Bandits. The matroid bandit problem was introduced by [20], who proposed the Optimistic Matroid Maximization (OMM) algorithm using UCB1 [3] to construct confidence bounds. [39] later improved gap-dependent constants via KL-UCB [15], yielding the KL-OSM algorithm. Beyond UCB approaches, [45] developed a Thompson Sampling variant with Beta priors that matched the optimal regret rate. Other works have explored related directions, including best-arm identification [8], efficiency improvements [31], submodular maximization [30], and robustness [40, 41]. Privacy in bandit learning has been studied extensively for classical [24, 33, 35, 42, 43] and combinatorial bandits [9], but only recently considered for matroid bandits by [6]. Compared to their work, we (1) study both central and local DP models, while they focus solely on CDP; (2) provide both gap-dependent and gap-free regret bounds, whereas they only give gap-dependent bounds; and (3) introduce a hybrid-scheme-based DP-UCB strategy, showing it is as effective as their lazy-update approach.

2 PRELIMINARIES

Matoid Optimization A matroid is a pair $\mathcal{M} = (E, \mathcal{I})$, where $E = \{1, \dots, N\}$ is a finite ground set, and $\mathcal{I} \subseteq 2^E$ is a family of subsets of E, called the independent sets, satisfying the following properties: (1) **Non-emptiness:** $\emptyset \in \mathcal{I}$; (2) **Hereditary property:** if $A \in \mathcal{I}$ and $A' \subseteq A$, then $A' \in \mathcal{I}$; (3) **Augmentation property:** if $A, B \in \mathcal{I}$ and $|A| > |B|$, then there exists $a \in A \setminus B$ such that $B \cup \{a\} \in \mathcal{I}$. An independent set $A \in \mathcal{I}$ is called a **basis** of $\mathcal{M}$ if it is inclusion-maximal in $\mathcal{I}$, i.e., there exists no $A' \in \mathcal{I}$ such that $A \subset A'$. Equivalently, for any $a \in E \setminus A$, we have $A \cup \{a\} \notin \mathcal{I}$. Let $\mathcal{B}$ denote the set of all bases of $\mathcal{M}$. It is a well-known fact that all bases of a matroid have the same cardinality [27], known as the **rank** of the matroid. We denote the rank by K, so that for any $A \in \mathcal{B}$, $|A| = K$. In the classical matroid optimization problem, each element $a \in E$ is assigned a non-negative weight w_a. Let $w \in (\mathbb{R}^+)^N$ denote the weight vector. The objective is to find a basis A^* that maximizes the total weight, i.e., $A^* := \arg\max_{A \in \mathcal{B}} \sum_{a \in A} w_a$. This problem admits an efficient solution via the greedy algorithm [13].

Matroid Bandit Optimization Matroid bandit optimization (MBO) is an online and stochastic generalization of matroid optimization. An agent interacts with a matroid bandit instance sequentially over T rounds. Each element $a \in E$ is referred to as a **base arm**, associated with an unknown reward distribution $\mathcal{D}_a$, and each basis in $\mathcal{B}$ is called a **super arm**. At each round $t \in [T]$, the agent selects a super arm $A(t) := \{a_1(t), \dots, a_K(t)\} \in \mathcal{B}$. Each selected base arm $a \in A(t)$ yields stochastic feedback $x_a(t) \sim \mathcal{D}_a$, drawn i.i.d. across rounds. We assume each $\mathcal{D}_a$ is bounded in $[0, 1]$. Let $\mu_a := \mathbb{E}[x_a(t)]$ denote the mean of $\mathcal{D}_a$, which serves as the expected reward or "weight" of base arm a. We consider the semi-bandit feedback setting [2], where the agent observes the feedback from each played base arm, i.e., upon playing $A(t)$, the agent receives total reward $r(t) := \sum_{a \in A(t)} x_a(t)$ and observes the feedback set $\{(a, x_a(t)) \mid a \in A(t)\}$. The goal is to design a learning policy π that maximizes the agent's expected cumulative reward over T rounds. Equivalently, we aim to minimize the expected cumulative **regret**, defined as: $\mathcal{R}_T := \mathbb{E}[\sum_{t=1}^{T}(\sum_{a \in A^\star} x_a(t) - \sum_{a \in A(t)} x_a(t))]$, where $A^\star := \arg\max_{A \in \mathcal{B}} \sum_{a \in A} \mu_a$ is the optimal super arm, and the expectation is over all randomness involved. W.l.o.g., let $A^\star = \{a_1^\star, \dots, a_K^\star\}$, sorted so that $a_k^\star$ is the base arm with the k-th largest mean. A base arm $a \in E \setminus A^\star$ is called **sub-optimal**; any $a_k^\star \in A^\star$ is an **optimal base arm**. For any pair $(a, a_k^\star)$ with a sub-optimal and $a_k^\star$ optimal, define the **pairwise gap** as $\Delta_{a,k} := \mu_{a_k^\star} - \mu_a$. For each sub-optimal base arm a, define the index set: $\mathcal{H}_a := \{k \in [K] \mid \Delta_{a,k} > 0\}$, which collects the indices of optimal arms with strictly higher mean than a. Let $H_a := |\mathcal{H}_a|$. Define the **gap** of base arm a as the smallest non-zero pairwise gap: $\Delta_a := \Delta_{a,H_a}$.

Differential Privacy Since the feedback in each round t may reveal sensitive user information, we aim to protect the data contributed by user t. We study how to design private learning algorithms for matroid bandit optimization that satisfy differential privacy (DP). Specifically, we consider both the *central* and *local* models of DP, and provide definitions tailored to the matroid bandit setting. In the **central DP model (CDP)**, we assume the agent is trusted and has access to the true feedback from each user $t \in [T]$. Let o_t denote the observation in round t, and $o_{1:t-1}$ the sequence of observations up to round $t-1$. We require the learning algorithm π to be differentially private, as defined below.

Definition 1 (ϵ-Differentially Private Matroid Bandit Algorithm). A matroid bandit algorithm π is said to be ϵ-differentially private if, for any two observation sequences $o_{1:t-1}$ and $o'_{1:t-1}$ that differ in at most one round, and for any measurable subset $\mathcal{U} \subseteq \mathcal{B}$, it holds that: $\mathbb{P}_\pi[A(t) \in \mathcal{U} \mid o_{1:t-1}] \leq e^\epsilon \cdot \mathbb{P}_\pi[A(t) \in \mathcal{U} \mid o'_{1:t-1}]$.

This definition requires the algorithm π to be randomized such that its output distribution does not change significantly when a single observation is modified. In the **local DP model (LDP)**, the agent is not trusted to access raw feedback. Instead, each user perturbs their feedback using a local randomizer $\mathcal{R}(\cdot)$ before sending it to the agent. The agent thus receives only privatized feedback. The local DP requirement is defined as follows.

Definition 2 (Local Differential Privacy). Let $\mathcal{R} : \mathcal{X} \to \mathcal{Y}$ be a randomized mechanism. Then $\mathcal{R}$ satisfies ϵ-local differential privacy if, for any $x, x' \in \mathcal{X}$ and any measurable subset $\mathcal{U} \subseteq \mathcal{Y}$, we have: $\mathbb{P}[\mathcal{R}(x) \in \mathcal{U}] \leq e^\epsilon \cdot \mathbb{P}[\mathcal{R}(x') \in \mathcal{U}]$.

In both CDP/LDP, ϵ is known as the **privacy budget**, which quantifies the level of indistinguishability between neighboring inputs: smaller ϵ indicate stronger privacy.

Definition 3 (Laplace Mechanism [12]). Let $q : \mathcal{X}^n \to \mathbb{R}^d$ be a query function, and let $D \in \mathcal{X}^n$ be a dataset. The Laplace mechanism with privacy parameter ϵ outputs: $\mathcal{L}(D, q, \epsilon) = q(D) + (L_1, L_2, \ldots, L_d)$, where each L_i is drawn i.i.d. from the zero-mean Laplace distribution $\mathrm{Lap}(\frac{\Delta_1(q)}{\epsilon})$. Here, $\Delta_1(q)$ denotes the ℓ_1-sensitivity of q: $\Delta_1(q) := \sup_{D\sim D'} \|q(D) - q(D')\|_1$, and $D \sim D'$ indicates that D and D' differ in at most one entry. The Laplace distribution $\mathrm{Lap}(b)$ has probability density function $\frac{1}{2b}\exp(-|x|/b)$. The Laplace mechanism satisfies ϵ-DP.

3 MBO IN CDP

We present an *anytime* algorithm that achieves ϵ-DP for the matroid bandit optimization problem under the central model. That is, our algorithm does not require prior knowledge of the time horizon T and maintains privacy guarantees throughout. It is based on the *Optimistic Matroid Maximization (OMM)* principle [20], combined with the *hybrid data releasing scheme* [5] to privately estimate the cumulative feedback for each base arm. The hybrid scheme is a powerful technique for privately releasing partial sums over unbounded data streams. Originally developed for general streaming data, it has been applied to private multi-armed bandits [16, 43]; here, we extend it to the differentially private combinatorial bandit setting. The key idea is to combine the logarithmic mechanism and the binary mechanism: the former releases noisy aggregates at exponentially increasing time steps $t = 2^r$ using Laplace noise, while the latter constructs additional noisy sums between these points to support continuous estimation. In contrast, [9] adopt only the binary mechanism for the combinatorial semi-bandit setting, which requires a known time horizon T and assumes a bounded stream. Moreover, they treat each round's feedback as a K-dimensional vector and apply a single binary mechanism to the entire vector stream, splitting the privacy budget ϵ evenly across components. This leads to a regret bound with a multiplicative factor of $K\log^3 T$, due to the sequence length for each base arm effectively being t rather than its own pull number $T_a(t-1)$. We address this limitation by assigning a dedicated hybrid instance to each base arm, allocating privacy budget ϵ evenly across arms. Consequently, the sequence length for each base arm a becomes its own pull count $T_a(t-1)$. This refinement improves the previous $K\log^3 T$ regret factor to $K\min\{K, \log T\}\log T$. Building on this design, we propose **Differentially Private Optimistic Matroid Maximization (DP-OMM)**, a DP adaptation of the OMM algorithm [20]. DP-OMM maintains private upper confidence bounds (UCBs) for each base arm using its associated hybrid mechanism. A detailed description is provided in Algorithm 1. At a high level, DP-OMM operates as follows: for each base arm $a \in E$, it creates a hybrid mechanism instance Hybrid_a and relies solely on its noisy feedback sum. In each round t, the algorithm computes the UCB $U_a(t)$ for each arm, sorts all arms by UCB in descending order, and greedily selects a maximal independent set $A(t)$. It then plays $A(t)$ and updates each Hybrid_a with the observed feedback $x_a(t)$.

Theorem 1. *DP-OMM (Algorithm 1) satisfies ϵ-DP under the central model. The expected cumulative regret of DP-OMM after T rounds is bounded as follows: (1)* ***Gap-dependent bound:*** $\mathcal{R}_T \leq \widetilde{O}(\sum_{a\in\overline{A^\star}} \frac{\log T}{\Delta_a} + \frac{(N-K)K\min\{K,\log T\}\log T}{\epsilon})$; *(2)* ***Gap-free bound:*** $\mathcal{R}_T \leq \widetilde{O}(\sqrt{(N-K)KT\log T} + \frac{(N-K)K\min\{K,\log T\}\log T}{\epsilon})$.

4 MBO IN LDP

We now study the problem of designing an ϵ-differentially private algorithm for matroid bandit optimization under the local DP model. As discussed earlier, in the local DP setting the agent is untrusted, and each user t must add noise to their own feedback before sending it to the agent. We adopt the Laplace mechanism for local randomization and

Algorithm 1: DP-OMM

Input: Time horizon T, privacy parameter ϵ
Output: Arm choices $A(t)$ for $t = 1, \ldots, T$

Construct N hybrid mechanisms for base arms denoted by $\mathsf{Hybrid}_1, \ldots, \mathsf{Hybrid}_N$, each with privacy budget $\epsilon_0 = \frac{\epsilon}{K}$;
for *each* $a \in I$ **do**
  Sample a feedback $x_a(0) \sim \mathcal{D}_a$;
  Insert $x_a(0)$ into Hybrid_a;
  Initialize $T_a(0) \leftarrow 1$;
for $t \leftarrow 1$ **to** T **do**
  for *each* $a \in E$ **do**
    Query Hybrid_a to obtain noisy feedback sum $\widetilde{S}_a(t-1)$;
    Compute the empirical mean: $\widetilde{w}_a(t-1) \leftarrow \widetilde{S}_a(t-1)/T_a(t-1)$;
    Compute the upper confidence bound:
    $U_a(t) \leftarrow \widetilde{w}_a(t-1) + \sqrt{\frac{2\log t}{T_a(t-1)}} + \frac{6\sqrt{8}\log(t)\log(T_a(t-1))}{\epsilon_0 \cdot T_a(t-1)}$;
  Initialize $A(t) \leftarrow \emptyset$;
  Sort the base arms in descending order of their upper confidence bounds, i.e., $U_{a_1}(t) \geq \cdots \geq U_{a_N}(t)$;
  for $i \leftarrow 1$ **to** N **do**
    if $A(t) \cup \{a_i\} \in \mathcal{I}$ **then**
      Update $A(t) \leftarrow A(t) \cup \{a_i\}$;
  Play the super arm $A(t)$ and obtain the reward $\sum_{a \in A(t)} x_a(t)$;
  for *each* $a \in A(t)$ **do**
    Insert the observed feedback $x_a(t)$ into Hybrid_a;
    Update $T_a(t) \leftarrow T_a(t-1) + 1$;

Algorithm 2: LDP-OMM

Input: Time horizon T, privacy parameter ϵ
Output: Arm choices $A(t)$ for $t = 1, \ldots, T$

$\epsilon_0 \leftarrow \frac{\epsilon}{K}$;
for *each* $a \in E$ **do**
  Sample a feedback $x_a(0) \sim \mathcal{D}_a$;
  Add Laplace noise to the feedback: $\widetilde{w}_a(0) \leftarrow x_a(0) + \nu$, where $\nu \sim \mathrm{Lap}(\frac{1}{\epsilon_0})$;
  Initialize $T_a(0) \leftarrow 1$;
for $t \leftarrow 1$ **to** T **do**
  for *each* $a \in E$ **do**
    Compute the upper confidence bound:
    $U_a(t) \leftarrow \widetilde{w}_a(t-1) + \sqrt{\frac{2\log t}{T_a(t-1)}} + \max\left\{\frac{1}{\epsilon_0}\sqrt{\frac{2\log(2t^4)}{T_a(t-1)}}, \frac{2\log(2t^4)}{\epsilon_0 T_a(t-1)}\right\}$;
  Initialize $A(t) \leftarrow \emptyset$;
  Sort the base arms in descending order of their upper confidence bounds, i.e., $U_{a_1}(t) \geq \cdots \geq U_{a_N}(t)$;
  for $i \leftarrow 1$ **to** N **do**
    if $A(t) \cup \{a_i\} \in \mathcal{I}$ **then**
      Update $A(t) \leftarrow A(t) \cup \{a_i\}$;
  Play the super arm $A(t)$ and obtain the reward $\sum_{a \in A(t)} x_a(t)$;
  for *each* $a \in A(t)$ **do**
    Add Laplace noise to the feedback: $x'_a(t) = x_a(t) + \nu$, where $\nu \sim \mathrm{Lap}(\frac{1}{\epsilon_0})$;
    Agent observes $x'_a(t)$ of each $a \in A(t)$;
    Update $T_a(t) \leftarrow T_a(t-1) + 1$;
    Update the empirical mean:
    $\widetilde{w}_a(t) \leftarrow \frac{T_a(t-1) \cdot \widetilde{w}_a(t-1) + x'_a(t)}{T_a(t)}$;

propose the **Locally Differentially Private Optimistic Matroid Maximization (LDP-OMM)** algorithm, presented in Algorithm 2. LDP-OMM extends the non-private OMM framework [20] to the local DP setting via a UCB-based strategy. The overall structure parallels Algorithm 1, but with a key distinction: LDP-OMM maintains private confidence intervals for each base arm using independently perturbed feedback, rather than aggregating noisy empirical means. Adding Laplace noise to each observed feedback introduces additional variance in mean estimation. Thus, the confidence intervals must be carefully adjusted to account for this noise. Prior results on Laplace concentration (e.g., [44, Lemma 1]) require a sample size n exceeding $\Omega\big(\log \frac{2}{\delta}\big)$, which is unsuitable for online settings where n starts from 1 and grows incrementally. To address this, we establish the following fine-grained concentration for Laplace random variables:

Lemma 1. *Let $X_1, \ldots, X_n \sim \mathrm{Lap}(b)$ be i.i.d. samples. Then, for any $n \geq 1$, with probability at least $1 - \delta$, we have:* $|\frac{1}{n}\sum_{i=1}^{n} X_i| \leq \max\{2b\sqrt{2\log(2/\delta)/n},\ 4b\log(2/\delta)/n\}$.

This result is used to construct the confidence intervals of LDP-OMM, which provides a valid confidence bound for all $n \in \mathbb{N}^+$. When $n > 2\log\frac{2}{\delta}$, the bound recovers the standard $O(\sqrt{\log(1/\delta)/n})$ rate as in [44]. When n is small, the bound transitions smoothly to $O(\log(1/\delta)/n)$, which is consistent with the classical tail bound for Laplace variables.

Theorem 2. *LDP-OMM (Algorithm 2) satisfies ϵ-DP under the local model. The expected cumulative regret of DP-OMM after T rounds is bounded as follows: (1)* ***Gap-dependent bound:*** $\mathcal{R}_T \leq O(\frac{K^2}{\epsilon^2}\sum_{a\in\overline{A^\star}}\frac{\log T}{\Delta_{a,H_a}})$*; (2)* ***Gap-free bound:*** $\mathcal{R}_T \leq O(\frac{K}{\epsilon}\sqrt{(N-K)KT\log T})$

5 EXPERIMENTS

We evaluate the performance of our proposed DP/LDP-OMM algorithms on real-world datasets. As a baseline, we use the non-private OMM algorithm [20]. The task considered is network routing, formulated as a stochastic optimization problem over a graphic matroid, using data from the RocketFuel dataset [37]. We assume that the routing network forms a spanning tree, and the Internet Service Provider (ISP) aims to learn a spanning tree that minimizes the expected total latency across all selected edges. The ground set E corresponds to the set of all graph edges, and a subset is considered independent if it contains no cycles. The edge latency at round t is defined as $x_e(t) := \mu_e - 1 + \theta(t)$, where μ_e is the mean latency of edge e from the RocketFuel dataset, and $\theta(t)$ is sampled from an exponential distribution with rate parameter 1. Since edge latencies in ISP networks often reflect physical distances between nodes [11], they may reveal users' location privacy, especially if an adversary has access to auxiliary structural information. This motivates the need to protect edge feedback during the learning process. In our experiments, the privacy budget ϵ is varied over $\{0.25, 0.5, 1.0, 2.0\}$ for CDP and $\{0.5, 1.0, 2.0, 4.0\}$ for LDP. Cumulative regret is used as the evaluation metric. For each configuration, we conduct at least 100 independent runs and report the average results. The performance of the baseline OMM algorithm is shown in all figures as a dashed line. Experiments are conducted on two ISP topologies of different sizes from the dataset. As shown in Figure 1, both DP/LDP-OMM achieve sublinear cumulative regret under various privacy levels, significantly outperforming the non-private baseline in terms of privacy protection. DP-OMM consistently yields lower regret than LDP-OMM, which is expected given that LDP enforces stricter privacy constraints. However, the performance gap narrows as ϵ increases, indicating that LDP-OMM can approach the performance of DP-OMM when the privacy budget is sufficiently large. These results demonstrate the efficacy of our algorithms for DP-MBO under both CDP/LDP.

6 CONCLUSION

We presented differentially private algorithms for matroid bandit optimization, addressing the challenge of learning efficient network topologies under rigorous privacy constraints. We designed algorithms that achieve strong theoretical guarantees and empirical performance. Our results demonstrate

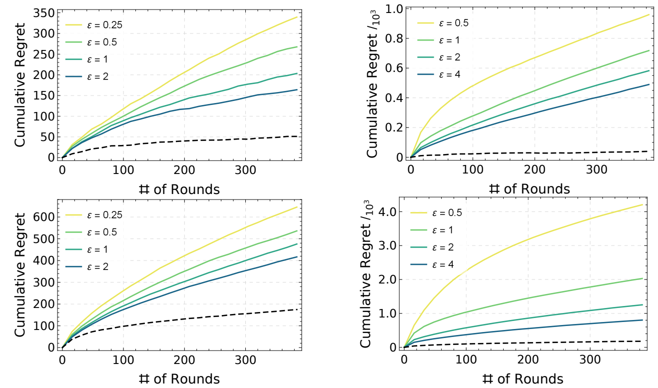

Figure 1: Cumulative regret of ϵ-DP-OMM (left column) and ϵ-LDP-OMM (right column) with baseline OMM on ISP networks. Top row: 108 nodes and 153 edges. Bottom row: 87 nodes and 161 edges.

that meaningful utility can be retained even under strict privacy requirements, enabling practical deployment of private topology optimization in sensitive networked systems.

REFERENCES

[1] Christoph Ambühl. 2005. An optimal bound for the MST algorithm to compute energy efficient broadcast trees in wireless networks. In *32nd International Colloquium on Automata, Languages, and Programming (ICALP 2005)*. Springer, Lisbon, Portugal, 1139–1150.

[2] Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. 2014. Regret in online combinatorial optimization. *Mathematics of Operations Research* 39, 1 (2014), 31–45.

[3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47 (2002), 235–256.

[4] Marc Barthélemy. 2011. Spatial networks. *Elsevier Physics Reports* 499, 1-3 (2011), 1–101.

[5] T.-H.Hubert Chan, Elaine Shi, and Dawn Song. 2011. Private and Continual Release of Statistics. *ACM Transactions on Information and System Security* 14, 3 (11 2011).

[6] Kushagra Chandak, Bingshan Hu, and Nidhi Hegde. 2023. Differentially private algorithms for efficient online matroid optimization. In *2nd Conference on Lifelong Learning Agents (CoLLAs 2023)*. PMLR, Montréal, Canada, 66–88.

[7] Arkadev Chattopadhyay, Jaikumar Radhakrishnan, and Atri Rudra. 2014. Topology matters in communication. In *55th IEEE Annual Symposium on Foundations of Computer Science*. IEEE, Philadelphia, PA, 631–640.

[8] Lijie Chen, Anupam Gupta, and Jian Li. 2016. Pure exploration of multi-armed bandit under matroid constraints. In *29th Annual Conference on Learning Theory (COLT 2016)*. PMLR, New York City, NY, 647–669.

[9] Xiaoyu Chen, Kai Zheng, Zixin Zhou, Yunchang Yang, Wei Chen, and Liwei Wang. 2020. (Locally) differentially private combinatorial semi-bandits. In *37th International Conference on Machine Learning (ICML 2020)*. PMLR, Virtual Conference, 1757–1767.

[10] Zina Chkirbene, Rachid Hadjidj, Sebti Foufou, and Ridha Hamila. 2020. LaScaDa: A novel scalable topology for data center network. *IEEE/ACM Transactions on Networking* 28, 5 (10 2020), 2051–2064.

[11] Baek-Young Choi, Sue Moon, Zhi-Li Zhang, Konstantina Papagiannaki, and Christophe Diot. 2007. Analysis of point-to-point packet delay in an operational network. *Elsevier Computer Networks* 51, 13 (2007),

3812–3827.

[12] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *3rd Theory of Cryptography Conference (TCC 2006)*. Springer, New York City, NY, 265–284.

[13] Jack Edmonds. 1971. Matroids and the greedy algorithm. *Mathematical Programming* 1 (1971), 127–136.

[14] Silin Feng, Jiajun Chen, and Zhiwei Zhao. 2022. Cost effective routing in large-scale multi-hop LoRa networks. In *41st IEEE International Conference on Computer Communications (INFOCOM 2022), 2nd IEEE INFOCOM Workshop on Networking Algorithms (WNA 2022)*. Virtual Conference, 1–6.

[15] Aurélien Garivier and Olivier Cappé. 2011. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *24th annual conference on learning theory (COLT 2011)*. PMLR, Budapest, Hungary, 359–376.

[16] Bingshan Hu, Zhiming Huang, and Nishant A Mehta. 2021. *Near-optimal algorithms for private online learning in a stochastic environment.* cs.LG 2102.07929. arXiv.

[17] Andy An-Kai Jeng and Rong-Hong Jan. 2011. Adaptive topology control for mobile ad hoc networks. *IEEE Transactions on Parallel and Distributed Systems* 22, 12 (2011), 1953–1960.

[18] Erez Kantor, Zvi Lotker, Merav Parter, and David Peleg. 2011. The topology of wireless communication. In *43rd annual ACM symposium on Theory of computing*. ACM, San Jose, CA, 383–392.

[19] Bhaskar Krishnamachari, Deborah Estrin, and Stephen Wicker. 2002. The Impact of Data Aggregation in Wireless Sensor Networks. In *22nd IEEE International Conference on Distributed Computing Systems (ICDCS 2002), International Workshop Distributed Event Based System (DEBS 2002)*. IEEE, Vienna, Austria.

[20] Branislav Kveton, Zheng Wen, Azin Ashkan, Hoda Eydgahi, and Brian Eriksson. 2014. Matroid bandits: fast combinatorial optimization with learning. In *13th Conference on Uncertainty in Artificial Intelligence (UAI 2014)*. AUAI Press, Québec, Canada, 420–429.

[21] Siyuan Li, Dong Ye, Zhaowei Sun, Jianqiao Zhang, and Weichao Zhong. 2023. Collision-free flocking control for satellite cluster with constrained spanning tree topology and communication delay. *IEEE Trans. Aerospace Electron. Systems* 59, 4 (2023), 4134–4146.

[22] Yang Li, Michael Purcell, Thierry Rakotoarivelo, David Smith, Thilina Ranbaduge, and Kee Siong Ng. 2023. Private graph data release: A survey. *Comput. Surveys* 55, 11 (2023), 1–39.

[23] Charalampos Mavroforakis, Richard Garcia-Lebron, Ioannis Koutis, and Evimaria Terzi. 2015. Spanning edge centrality: Large-scale computation and applications. In *24th international conference on world wide web (WWW 2015)*. ACM, Florence, Italy, 732–742.

[24] Nikita Mishra and Abhradeep Thakurta. 2015. (Nearly) optimal differentially private stochastic multi-arm bandits. In *31st Conference on Uncertainty in Artificial Intelligence (UAI 2015)*. AUAI Press, Amsterdam, Netherlands, 592–601.

[25] Andrea Nardin, Juan A Fraire, and Fabio Dovis. 2022. Contact plan design for GNSS constellations: A case study with optical intersatellite links. *IEEE Trans. Aerospace Electron. Systems* 58, 3 (2022), 1981–1995.

[26] Hiroki Nishiyama, Thuan Ngo, Nirwan Ansari, and Nei Kato. 2012. On minimizing the impact of mobility on topology control in mobile ad hoc networks. *IEEE Transactions on Wireless Communications* 11, 3 (2012), 1158–1166.

[27] James G Oxley. 2006. *Matroid theory.* OUP, Oxford, United Kingdom.

[28] Rasmus Pagh and Lukas Retschmeier. 2024. *Faster private minimum spanning trees.* cs.DS 2408.06997. arXiv.

[29] Rasmus Pagh, Lukas Retschmeier, Hao Wu, and Hanwen Zhang. 2025. Optimal Bounds for Private Minimum Spanning Trees via Input Perturbation. *Proceedings of the ACM on Management of Data (PACMMOD)* 3, 2 (2025), 1–26.

[30] Orestis Papadigenopoulos and Constantine Caramanis. 2021. Recurrent submodular welfare and matroid blocking semi-bandits. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, Vol. 34. Curran Associates Inc., Virtual Conference, 23334–23346.

[31] Pierre Perrault, Vianney Perchet, and Michal Valko. 2019. Exploiting structure of uncertainty for efficient matroid semi-bandits. In *36th International Conference on Machine Learning (ICML 2019)*. Long Beach, CA, 5123–5132.

[32] Rafaël Pinot. 2018. *Minimum spanning tree release under differential privacy constraints.* cs.CR 1801.06423. arXiv.

[33] Touqir Sajed and Or Sheffet. 2019. An optimal private stochastic-mab algorithm based on optimal private stopping rule. In *36th International Conference on Machine Learning (ICML 2019)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.), Vol. 97. PMLR, Long Beach, CA, 5579–5588.

[34] Adam Sealfon. 2016. Shortest paths and distances with differential privacy. In *35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS 2016)*. ACM, San Francisco, CA, 29–41.

[35] Roshan Shariff and Or Sheffet. 2018. Differentially private contextual linear bandits. In *32nd International Conference on Neural Information Processing Systems (NIPS 2018)*. Curran Associates Inc., Montréal, Canada, 4301–4311.

[36] Ankit Singla, P Brighten Godfrey, and Alexandra Kolla. 2014. High throughput data center topology design. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2014)*. USENIX, Seattle, WA, 29–41.

[37] Neil Spring, Ratul Mahajan, and David Wetherall. 2002. Measuring ISP topologies with Rocketfuel. *ACM SIGCOMM Computer Communication Review* 32, 4 (2002), 133–145.

[38] Mudhakar Srivatsa and Mike Hicks. 2012. Deanonymizing mobility traces: Using social network as a side-channel. In *19th ACM Conference on Computer and Communications Security (CCS 2012)*. ACM, Raleigh, NC, 628–637.

[39] Mohammad Sadegh Talebi and Alexandre Proutiere. 2016. An optimal algorithm for stochastic matroid bandit optimization. In *15th International Conference on Autonomous Agents & Multiagent Systems (AAMAS 2016)*. IFAAMAS, Singapore, Singapore, 548–556.

[40] Youming Tao, Xiuzhen Cheng, Falko Dressler, Zhipeng Cai, and Dongxiao Yu. 2024. Robust Matroid Bandit Optimization against Adversarial Contamination. In *30th International Computing and Combinatorics Conference (COCOON 2024)*. Springer, Shanghai, China, 326–347.

[41] Youming Tao, Xiuzhen Cheng, Falko Dressler, Zhipeng Cai, and Dongxiao Yu. 2025. Robust Matroid Bandit Optimization: Near-Optimal Rates under Adversarial Contamination. *Theoretical Computer Science* 1051 (10 2025), 115416.

[42] Youming Tao, Yulian Wu, Peng Zhao, and Di Wang. 2022. Optimal Rates of (Locally) Differentially Private Heavy-tailed Multi-Armed Bandits. In *25th International Conference on Artificial Intelligence and Statistics (AISTATS 2022)*, Vol. 151. PMLR, Valencia, Spain.

[43] Aristide Tossou and Christos Dimitrakakis. 2016. Algorithms for differentially private multi-armed bandits. In *30th AAAI Conference on Artificial Intelligence*, Vol. 30. AAAI, Phoenix, AZ.

[44] Di Wang, Marco Gaboardi, and Jinhui Xu. 2018. Empirical risk minimization in non-interactive local differential privacy revisited. In *32nd International Conference on Neural Information Processing Systems (NIPS 2018)*. Curran Associates Inc., Montréal, Canada.

[45] Siwei Wang and Wei Chen. 2018. Thompson sampling for combinatorial semi-bandits. In *35th International Conference on Machine Learning (ICML 2018)*. PMLR, Stockholm, Sweden, 5114–5122.

[46] Jie Wu and Fei Dai. 2006. Mobility-sensitive topology control in mobile ad hoc networks. *IEEE Transactions on Parallel and Distributed Systems* 17, 6 (2006), 522–535.