Appendix A Additional Preliminaries

A.1 Convexity and Lipschitz Continuity

For a set X we define its diameter $D_X = \sup_{x,y \in X} \|x-y\|_2$. A set $X \subseteq \mathbb{R}^d$ is a *convex set* if for any $x, y \in X$ and any $\lambda \in [0,1]$, $\lambda x + (1-\lambda)y \in X$. For a function $f : X \to \mathbb{R}$, a *subgradient* of f at a point y, denoted $\nabla f(y)$, is a vector $g \in \mathbb{R}^d$ such that $f(x) - f(y) \geq g^\top (x-y)$ for all $x \in X$. The *subdifferential* of f at y, denoted $\partial f(y)$, is the set of all subgradients of f at y.

Definition 5 (Strongly convex function). *Let $X \subseteq \mathbb{R}^d$ be a convex set. A function $f : X \to \mathbb{R}$ is H-*strongly convex *for $H \geq 0$ if, $f(x) \geq f(y) + \nabla f(y)^\top (x-y) + \frac{H}{2}||x-y||_2^2$ for all $x, y \in X$. If $H = 0$, we say that f is* convex.[2]

Note that every strongly convex function is also convex. For convex f, the subdifferential at every point always exists and is a closed convex set.

Definition 6 (Lipschitz function). *A function $f : X \to \mathbb{R}$ is L-*Lipschitz continuous *with respect to a norm $|| \cdot ||$ if $|f(x) - f(y)| \leq L||x-y||$ for every $x, y \in X$.*

Lemma 9 gives an equivalence between Lipschiptzness of a convex function and properties of that function's subgradients.

Lemma 9 ([25]). *Let $f : X \to \mathbb{R}$ be a convex function. Then f is L-Lipschitz over X with respect to norm $|| \cdot ||$ if and only if for all $x \in X$ and for all $\nabla f(x) \in \partial f(x)$ we have that $||\nabla f(x)||_* \leq L$, where $|| \cdot ||_*$ denotes the dual norm of $|| \cdot ||$.*

Throughout the paper, we will say that a function f is L-Lipschitz to indicate that f is L-Lipschitz with respect to the L_2 norm $|| \cdot ||_2$, unless otherwise stated. We also note that the L_2 norm is self-dual: $(|| \cdot ||_2)_* = || \cdot ||_2$ [24].

A.2 Tree-Based Aggregation Protocol (TBAP)

The Tree-Based Aggregation Protocol is a tool for maintaining differentially private partial sums of vectors arriving in an online sequence. At each time t, TBAP outputs a noisy sum of the input vectors up to time t. This algorithm was first introduced by Chan et al. [7] and Dwork et al. [10], and adapted in its current form by Smith and Thakurta [26].

The algorithm works by maintaining a complete binary tree, where the d-dimensional input vectors are stored in the leaf nodes, and internal nodes in the tree store a noisy sum of all leaves in their sub-tree. At each time t, TBAP receives input z_t and updates the value of the t-th leaf node to be z_t. The algorithm also updates the value of each internal node affected by this change to be the updated sum plus noise drawn according to a high-dimensional analog of Laplace noise. The algorithm then outputs a noisy partial sum v_t of the nodes in the tree that approximately sum to z_t.

The sum at each internal node is $(\epsilon / \log_2 T)$-differentially private, and by construction each z_t affects only $\log_2 T$ nodes of the tree. By the *composition property* of differential privacy [9], the entire tree is ϵ-differentially private (Theorem 10).

Theorem 10 ([7, 10]). TBAP$(\{z_i\}_{i=1}^T, \mu, \epsilon)$ *is ϵ-differentially private for any $\mu > 0$ and any sequence of vectors $z_1, \ldots, z_T$ that each have L_2 norm at most μ.*

In addition to being private, TBAP also provides partial sums $v_t = \sum_{i=1}^t z_t$ that are accurate (with respect to the L_2 norm) up to additive $O(\frac{d\mu \log^2 T}{\epsilon})$. This is because the L_2 norm of the noise at each node is Gamma distributed with standard deviation $O(\frac{\sqrt{d}\mu \log T}{\epsilon})$, and each partial sum is computed using at most $\log T$ nodes in the tree.

[2] This is equivalent to the more commonly used definition that f is convex if for any $\lambda \in [0,1]$ and for any $x, y \in X$, $\lambda f(x) + (1-\lambda) f(y) \geq f(\lambda x + (1-\lambda) y)$.

Algorithm 4 Tree Based Aggregation Protocol: TBAP($\{z_i\}_{i=1}^T, \mu, \epsilon$)

Input: Online sequence of vectors $z_1, ..., z_T \in \mathbb{R}^d$, $\mu : L_2-$norm bound on each z_i, privacy parameter ϵ.
Output: Sequence of noisy partial sums $v_1, \ldots, v_n \in \mathbb{R}$
Initialize a binary tree A of size $2^{\lceil \log_2 T \rceil + 1} - 1$ with leaves $z_1, ..., z_T$
for t = 1, . . . , T **do**
  Accept z_t from the data stream.
  Let $P = \{z_t \to \cdots \to root\}$ be the path from z_t to the root.
  procedure TREE UPDATE
    Let Λ be the first node in P that is left-child in A. Let $P_\Lambda = \{z_t \to \cdots \to \Lambda\}$.
    for all nodes α in path P **do**
      $\alpha \leftarrow \alpha + z_t$
      if $\alpha \in P_\Lambda$ **then** $\alpha \leftarrow \alpha + \gamma$ where $\gamma \in \mathbb{R}^d$ is sampled by $\Pr[\gamma = \hat{\gamma}] \propto e^{-\frac{\|\hat{\gamma}\|_2 \epsilon}{\mu(\lceil \log_2 T \rceil + 1)}}$
      end if
    end for
  end procedure
  procedure OUTPUT PRIVATE PARTIAL SUM
    Initialize vector $v_t \in \mathbb{R}^d$ to zero. Let b be a $(\lceil \log_2 T \rceil + 1)$-bit binary representation of t.
    for $i = 1, \ldots, [\log_2 T + 1]$ **do**
      if bit $b_i = 1$ **then**
        if i-th node in P (denoted $P(i)$) is the left child in A, **then** $v \leftarrow v + P(i)$
        else $v_t \leftarrow v_t$+left sibling $P(i)$
        end if
      end if
    end for
    return noisy partial sum v_t
  end procedure
end for

A.3 The Cost of Privacy in Online Convex Optimization

Theorem 5. *Let $\{\hat{x}_t\}_{t=1}^T$ be the non private iterates of RFTL and let $\{x_t\}_{t=1}^T$ be the private iterates i.e. $x_{t+1} = \arg\min_{x \in X} v_t^\top x + \frac{H}{2}||x||^2$ where v_t is the private partial sum computed using $TBAP\{\nabla f_t(x_t), L, \epsilon\}$. It holds that*

$$\mathbb{E}[\sum_{t=1}^T f_t(x_t)] \leq \mathbb{E}[\sum_{t=1}^T f_t(\hat{x}_t)] + \frac{4nL^2T\ln^{1.5}(T)}{\epsilon H}.$$

Where the expectation is taken with respect to the randomness of TBAP.

Our proof follows a similar structure as that of Lemma 8 of [26]. However, we analyze a different algorithm, so the proof details are different.

Proof. Let $J_t = v_t^\top x + \frac{H}{2}||x||^2$. Let $\xi_t = v_t - \sum_{\tau=1}^t \nabla f_t(x_t)$ be the noise added by TBAP to $\sum_{\tau=1}^t \nabla f_t(x_t)$. Notice that $x_{t+1} = \arg\min_{x \in \mathcal{K}} J_t(x) + \xi_t^\top x$ and $\hat{x}_t = \arg\min_{x \in \mathcal{K}} J_t(x)$. Since J_t is $\frac{H}{2}$-strongly convex we have that

$$||\hat{x}_{t+1} - x_{t+1}|| \leq \frac{2||\xi_t||}{H}.$$

Since each ξ_t is formed in TBAP by adding at most $\lceil \ln(T) + 1 \rceil$ vectors with norms drawn from a Gamma distribution with scale n and shape $\frac{(\lceil \ln(T)+1 \rceil)G}{\epsilon}$ we can upper bound $\mathbb{E}[||\xi_t||]$ by $\frac{4nG\ln^{1.5}(T)}{\epsilon}$.

Since f_t is L-Lipschitz continuous we have that

$$\mathbb{E}[\sum_{t=1}^{T} f_t(x_t)] \le \mathbb{E}[\sum_{t=1}^{T} f_t(\hat{x}_t)] + \mathbb{E}[L\sum_{t=1}^{T} \frac{2||\xi_t||}{H}]$$
$$\le \mathbb{E}[\sum_{t=1}^{T} f_t(\hat{x}_t)] + \frac{4nL^2T\ln^{1.5}(T)}{\epsilon H}.$$

□

Appendix B Omitted Proofs

B.1 Full Information Setting

Theorem 6 (Privacy guarantee)**.** SubmodPRFTL($\{f_i\}_{i=1}^T, M, H, L, [n], \epsilon$) *is ϵ-differentially private for any sequence of functions $f_1, \dots, f_T$ with bounded range $[-M, M]$ and for any $M, H, L, n, T > 0$.*

Proof. By Theorem 10 we know that the output of TBAP, $\{v_t\}_{t=1}^T$, is ϵ-differentially private. By Theorem 3 we get that the sequence $\{x_t\}_{t=1}^T$ is ϵ-differentially private since the procedure $x_{t+1} \leftarrow \arg\min_{x\in K} v_t^\top x + \frac{H}{2}||x||_2^2$ is simply post-processing of the v_t's. Computing the output $\{S_t\}_{t=1}^T$ is further post-processing of the sequence $\{x_t\}_{t=1}^T$, and Theorem 3 again yields the result. □

Theorem 7 (Regret guarantee)**.** SubmodPRFTL($\{f_i\}_{i=1}^T, M, H, L, [n], \epsilon$) *run with $H = M\sqrt{T}$ and $||\nabla \hat{f}_t|| \le L = 4M$ for any sequence of submodular functions $f_1, \dots, f_T : 2^{[n]} \to [-M, M]$ for any $M, n, T > 0$ guarantees,*

$$\mathbb{E}[\text{Regret}(T)] \le O\left(\frac{nM^2\ln^{1.5}(T)\sqrt{T}}{\epsilon}\right),$$

where the expectation is taken over the randomness of TBAP *and the sampling procedure used to choose S_t.*

Proof. Let $E_{TBAP}[\cdot]$ be the expectation taken with respect to the randomness of TBAP. Notice that $\mathbb{E}[\sum_{t=1}^T f_t(S_t)] = \mathbb{E}_{TBAP}[\mathbb{E}[\sum_{t=1}^T f_t(S_t)|TBAP]] = \mathbb{E}_{TBAP}[\mathbb{E}[\sum_{t=1}^{T-1} f_t(S_t)|TBAP]] + \mathbb{E}_{TBAP,\tau_1,\dots,\tau_{T-1}}[\mathbb{E}[f_T(S_T)|TBAP, \tau_1, ..., \tau_{T-1}]] = \mathbb{E}_{TBAP}[\mathbb{E}[\sum_{t=1}^{T-1} f_t(S_t)|TBAP]] + \mathbb{E}_{TBAP}[\hat{f}_T(x_T)]$ by definition of $\hat{f}$. Repeating the argument $T-1$ more times we get $\mathbb{E}[\sum_{t=1}^T f_t(S_t)] = \mathbb{E}_{TBAP}[\sum_{t=1}^T \hat{f}_t(\hat{x}_t)]$. Now,

$$\mathbb{E}[\sum_{t=1}^{T} f_t(S_t) - \min_{S\subseteq[n]} \sum_{t=1}^{T} f_t(S)]$$
$$\le \mathbb{E}[\sum_{t=1}^{T} f_t(S_t) - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{f}_t(x)]$$
$$= \mathbb{E}_{TBAP}[\sum_{t=1}^{T} \hat{f}_t(x_t) - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{f}_t(x)]$$
$$\le \mathbb{E}_{TBAP}[\sum_{t=1}^{T} \hat{f}_t(\hat{x}_t) - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{f}_t(x)] + \frac{4nL^2T\ln^{1.5}(T)}{\epsilon H} \quad \text{(by Theorem 5)}$$
$$\le \frac{2}{H}\sum_{t=1}^{T} ||\nabla \hat{f}_t(\hat{x}_t)||^2 + \frac{H}{2}[||x||^2 - ||\hat{x}_1||^2] + \frac{4nL^2T\ln^{1.5}(T)}{\epsilon H} \quad \text{(by Theorem 4)}$$
$$\le \frac{2TL^2}{H} + \frac{Hn}{2} + \frac{4nL^2T\ln^{1.5}(T)}{\epsilon H}$$

Plugging in the bound on L from Lemma 2 and choosing $H = M\sqrt{T}$ yields the result. □

B.2 Bandit Setting

Theorem 8 (Privacy guarantee)**.** BanditSubmodPRFTL($\{f_i\}_{i=1}^T, M, H, L, [n], \epsilon, \gamma$) *is ϵ-differentially private for any sequence of functions $f_1, \dots, f_T$ with bounded range $[-M, M]$ and for any $M, H, L, n, T, \gamma > 0$.*

Proof. By Theorem 10 we know that the output of TBAP, $\{v_t\}_{t=1}^T$, is ϵ-differentially private. Notice that BanditSubmodPRFTL is running PRFTL on regularized functions $\hat{g}_t^\top x + \frac{H}{2}\|x\|^2$ thus by the same reasoning as in Theorem 6, the sequence $\{x_t\}_{t=1}^T$ is ϵ-differentially private since the procedure $x_{t+1} \leftarrow \arg\min_{x\in K} v_t^\top x + \frac{H}{2}\|x\|_2^2$ is simply post-processing of the v_t's. Since $\{S_t\}_{t=1}^T$ is post-processing on the sequence $\{x_t\}_{t=1}^T$, applying Theorem 3 again completes the proof. ☐

Lemma 3. *Let $\gamma \in (0,1)$. The random vector $\hat{g}_t$ computed in* BanditSubmodPRFTL *is an unbiased estimate of a subgradient of the Lovasz extension $\hat{f}_t$ of submodular f_t, evaluated at point x_t. That is,*

$$\mathbb{E}\left[\hat{g}_t \mid x_t\right] = \nabla \hat{f}_t(x_t).$$

Proof. Notice that conditioned on the randomness up to $t-1$

$$\hat{g}_t = \begin{cases} -\frac{1}{\rho_0} f_t(B_0) e(\pi^{-1}(1)) & \text{with probability } \rho_0 \\ \frac{2}{\rho_i} f_t(B_i) e(\pi^{-1}(i)) & \text{with probability } \frac{\rho_i}{2} \text{ for } 1 \le i \le n-1 \\ -\frac{2}{\rho_i} f_t(B_i) e(\pi^{-1}(i+1)) & \text{with probability } \frac{\rho_i}{2} \text{ for } 1 \le i \le n-1 \\ \frac{1}{\rho_n} f_t(B_n) e(\pi^{-1}(n)) & \text{with probability } \rho_n \end{cases}$$

Therefore

$$\begin{aligned} \mathbb{E}_t[\hat{g}_t] &= \rho_0[-\frac{1}{\rho_0} f_t(B_0) e(\pi^{-1}(1))] + \frac{\rho_1}{2}[\frac{2}{\rho_1} f_t(B_1) e(\pi^{-1}(1)) - \frac{2}{\rho_1} f_t(B_1) e(\pi^{-1}(2))] \\ &+ ... + \frac{\rho_{n-1}}{2}[\frac{2}{\rho_{n-1}} f_t(B_{n-1}) e(\pi^{-1}(n-1)) - \frac{2}{\rho_{n-1}} f_t(B_{n-1}) e(\pi^{-1}(n))] + \rho_n[\frac{1}{\rho_n} f_t(B_n) e(\pi^{-1}(n))] \\ &= [f_t(B_1) - f_t(B_0)] e(\pi^{-1}(1)) + [f_t(B_2) - f_t(B_1)] e(\pi^{-1}(2)) + ... + [f_t(B_n) - f_t(B_{n-1})] e(\pi^{-1}(n)) \end{aligned}$$

This means that $\mathbb{E}_t[\hat{g}_t](\pi^{-1}(i)) = f(B_i) - f_t(B_{i-1})$ for $i = 1, ..., n$. This concludes the proof since $\mathbb{E}_t[\hat{g}_t](i) = \mathbb{E}_t[\hat{g}_t](\pi^{-1}[\pi(i)]) = f_t(B_{\pi(i)}) - f_t(B_{\pi(i)-1}) = g_t(i)$ for $i = 1, ..., n$. ☐

Lemma 4. *The random vector $\hat{g}_t$ computed in* BanditSubmodPRFTL *satisfies the following bound on its expected L_2-norm,*

$$\mathbb{E}\left[\|\hat{g}_t\|^2\right] \le \frac{16M^2n^2}{\gamma},$$

where the expectation is taken over the algorithm's internal randomness up to time t.

Proof.

$$\begin{aligned} \mathbb{E}_t[||\hat{g}_t||^2] &= \rho_0[-\frac{1}{\rho_0} f_t(B_0)]^2 + \sum_{i=1}^{n-1} \frac{\rho_i}{2}[(\frac{2}{\rho_i} f_t(B_i))^2 + (-\frac{2}{\rho_i}) f_t(B_i)^2] + \rho_n[\frac{1}{\rho_n} f_t(B_n)^2] \\ &\le 4M^2 \sum_{i=0}^{n} \frac{1}{\rho_i} \\ &= 4M^2 \sum_{i=0}^{n} \frac{1}{(1-\gamma)\mu_i + \gamma/(n+1)} \\ &= \sum_{i=0}^{n} \frac{n+1}{(1-\gamma)\mu_i(n+1) + \gamma} \end{aligned}$$

$$\leq \frac{4M^2(n+1)^2}{\gamma}$$
$$\leq \frac{16M^2n^2}{\gamma}$$

The second to last inequality holds as long as $\gamma \leq 1$ which will be ensured when we choose the parameters of the algorithm. □

Lemma 5 ([14])**.** *For any submodular function $f_t : [n] \to [-M, M]$, let x_t and S_t be the corresponding iterates and sets as defined in* BanditSubmodPRFTL*, then $\mathbb{E}[f_t(S_t)] \leq \mathbb{E}[\hat{f}_t(x_t)] + 2\gamma M$. Where the expectation is taken with respect to all the randomness of the algorithm.*

The proof is identical to that of [14]. We present it here for completeness. Let $\mathbb{E}_t$ be the expectation with respect to the randomness of the algorithm in round t conditioned on the history up to time $t-1$.

Proof. We know $\mathbb{E}_t[f_t(S_t)] = \sum_{i=0}^{n} \rho_i f_t(B_i)$ and $\hat{f}_t(x_t) = \sum_{i=0}^{n} \mu_i f(B_i)$. Therefore,

$$\begin{aligned}
\mathbb{E}_t[f_t(S_t)] - \hat{f}_t(x_t) &= \sum_{i=0}^{n} (\rho_i - \mu_i) f_t(B_i) \\
&\leq \gamma \sum_{i=0}^{n} \left[\frac{1}{n+1} + \mu_i\right] |f_t(B_i)| \\
&= \gamma \left(\frac{n}{n+1} + 1\right) M \\
&\leq 2\gamma M.
\end{aligned}$$

Taking expectation with respect to the randomness up to time $t-1$ yields the result. □

Lemma 6. *Let $\{\hat{g}_t\}_{t=1}^{T}$ be the sequence of one point gradient estimates generated by* BanditSubmodPRFTL*($\{f_i\}_{i=1}^{T}, M, H, L, [n], \epsilon, \gamma$). Then,*

$$\mathbb{E}\left[\min_{x \in \mathcal{K}} \sum_{t=1}^{T} \hat{g}_t^\top x\right] \leq \mathbb{E}\left[\min_{x \in \mathcal{K}} \sum_{t=1}^{T} \nabla \hat{f}_t^\top x\right] + \frac{8Mn\sqrt{T}}{\sqrt{\gamma}},$$

where the expectation is taken with respect to all the randomness of the algorithm.

Proof. Define $\alpha_t = \nabla \hat{f}_t - \hat{g}_t$. Notice that with probability 1

$$\begin{aligned}
&|\sum_{t=1}^{T} \hat{g}_t^\top x - \sum_{t=1}^{T} \nabla \hat{f}_t^\top x| \\
&\leq ||x||_2 || \sum_{t=1}^{T} \alpha_t ||_2 \quad \text{(by Cauchy Schwartz)}
\end{aligned}$$

Therefore, with probability 1

$$\min_{x \in \mathcal{K}} \sum_{t=1}^{T} \hat{g}_t^\top x \leq \min_{x \in \mathcal{K}} \sum_{t=1}^{T} \nabla \hat{f}_t(x_t)^\top x + || \sum_{t=1}^{T} \alpha_t ||_2. \tag{1}$$

The previous ensures that our regret bound holds against adaptive adversaries.

We next proceed to bound $\mathbb{E}\left[\|\sum_{t=1}^{T}\alpha_t\|_2\right]^2$. By Lemma 10 stated below, $\mathbb{E}[\alpha_t^\top \alpha_{t'}] = 0$ for $t \neq t'$.

$$
\begin{aligned}
\mathbb{E}\left[\left\|\sum_{t=1}^{T}\alpha_t\right\|_2\right]^2 &\leq \mathbb{E}\left[\left\|\sum_{t=1}^{T}\alpha_t\right\|_2^2\right] \qquad \text{(by Jensen's inequality)}\\
&= \sum_{t=1}^{T}\mathbb{E}\left[\|\alpha_t\|_2^2\right] + 2\sum_{t<t'}\mathbb{E}\left[\alpha_t^\top \alpha_{t'}\right]\\
&= \sum_{t=1}^{T}\mathbb{E}\left[\|\nabla \hat{f}_t(x_t) - \hat{g}_t\|_2^2\right]\\
&\leq \sum_{t=1}^{T}\mathbb{E}\left[2\|\nabla \hat{f}_t(x_t)\|_2^2 + 2\|\hat{g}_t\|_2^2\right]\\
&\leq 4T\cdot\frac{16M^2n^2}{\gamma}
\end{aligned}
$$

where the last line follows from Lemma 4, and the fact that if $\|\hat{g}_t\|_2 \leq G$ then $\|\nabla \hat{f}_t(x_t)\|_2 \leq G$ by Jensen's inequality. Taking expectation on both sides of equation 1 yields the result $\square$

The following lemma was asserted without proof in [26]. We prove it here for completeness.

Lemma 10. *Let $\alpha_t = \nabla \hat{f}_t(x_t) - \hat{g}_t$. Then, for $t < t'$ it holds that $\mathbb{E}[\alpha_t^\top \alpha_{t'}] = 0$, where the expectation is taken over the randomization of the algorithm used to build the estimates of the gradient $\{\hat{g}_t\}_{t=1}^T$.*

Proof.

$$
\begin{aligned}
\mathbb{E}[\alpha_t^\top \alpha_{t'}] &= \mathbb{E}[(\nabla \hat{f}_t(x_t) - \hat{g}_t)^\top(\nabla \hat{f}_{t'}(x_{t'}) - \hat{g}_{t'})]\\
&= \mathbb{E}[\nabla \hat{f}_t(x_t)^\top \nabla \hat{f}_{t'}(x_{t'})] - \mathbb{E}[\nabla \hat{f}_t(x_t)^\top \hat{g}_{t'}] - \mathbb{E}[\nabla \hat{f}_{t'}(x_{t'})^\top \hat{g}_t] + \mathbb{E}[\hat{g}_t^\top \hat{g}_{t'}]\\
&= \nabla \hat{f}_t(x_t)^\top \nabla \hat{f}_{t'}(x_{t'}) - \nabla \hat{f}_t(x_t)^\top \nabla \hat{f}_{t'}(x_{t'}) - \nabla \hat{f}_{t'}(x_{t'})^\top \nabla \hat{f}_t(x_t) + \mathbb{E}[\hat{g}_t^\top \hat{g}_{t'}]
\end{aligned}
$$

We now show that $\mathbb{E}[\hat{g}_t^\top \hat{g}_{t'}] = \nabla \hat{f}_{t'}(x_{t'})^\top \nabla \hat{f}_t(x_t)$.

$$
\begin{aligned}
\mathbb{E}[\hat{g}_t^\top \hat{g}_{t'}] &= \mathbb{E}_{1,\ldots t'-1}[\mathbb{E}_{t'}[\hat{g}_t^\top \hat{g}_{t'} | t = 1, \ldots t'-1]]\\
&= \mathbb{E}_{1,\ldots t'-1}[\hat{g}_t^\top \mathbb{E}_{t'}[\hat{g}_{t'} | t = 1, \ldots t'-1]]\\
&= \mathbb{E}_{1,\ldots t'-1}[\hat{g}_t^\top \nabla \hat{f}_{t'}(x_{t'})]\\
&= \nabla \hat{f}_t^\top(x_t) \nabla \hat{f}_{t'}(x_{t'})
\end{aligned}
$$

$\square$

Lemma 7. *Let $\{\hat{g}_t\}_{t=1}^T$ and $\{x_t\}_{t=1}^T$ be the sequences generated by* BanditSubmodPRFTL*$(\{f_i\}_{i=1}^T, M, H, L, [n], \epsilon, \gamma)$. Then,*

$$
\mathbb{E}[\sum_{t=1}^{T}\hat{g}_t^\top x_t] = \mathbb{E}[\sum_{t=1}^{T}\nabla \hat{f}_t^\top x_t],
$$

where the expectation is taken with respect to all the randomness of the algorithm.

Proof.

$$\begin{aligned}
&\mathbb{E}[\sum_{t=1}^{T} \hat{g}_t^\top x_t] \\
&= \mathbb{E}[\sum_{t=1}^{T-1} \hat{g}_t^\top x_t] + \mathbb{E}[\hat{g}_T^\top x_T] \\
&= \mathbb{E}[\sum_{t=1}^{T-1} \hat{g}_t^\top x_t] + \mathbb{E}[\mathbb{E}[\hat{g}_T^\top x_T | \tau = 1, ..., T-1]] \\
&= \mathbb{E}[\sum_{t=1}^{T-1} \hat{g}_t^\top x_t] + \mathbb{E}[x_T^\top \mathbb{E}[\hat{g}_T | \tau = 1, ..., T-1]] \\
&= \mathbb{E}[\sum_{t=1}^{T-1} \hat{g}_t^\top x_t] + \mathbb{E}[x_T^\top \mathbb{E}[\hat{g}_T | \tau = 1, ..., T-1]] \\
&= \mathbb{E}[\sum_{t=1}^{T-1} \hat{g}_t^\top x_t] + \mathbb{E}[x_T^\top \nabla \hat{f}_T] \quad \text{by Lemma 3.}
\end{aligned}$$

Repeating the argument $T-1$ more times yields the result. □

Lemma 8. *Let $\{x_t\}_{t=1}^T$ be the sequence generated by* BanditSubmodPRFTL*($\{f_i\}_{i=1}^T, M, H, L, [n], \epsilon, \gamma$). Let $\hat{x}_t$ be the non private iterate of the algorithm, that is $\hat{x}_{t+1} = \sum_{\tau=1}^t \hat{g}_\tau^\top x + \frac{H}{2}||x||^2$. Then,*

$$\mathbb{E}[\sum_{t=1}^{T} \hat{g}_t^\top x_t] \leq \mathbb{E}[\sum_{t=1}^{T} \hat{g}_t^\top \hat{x}_t] + \frac{64 n^3 M^2 T \ln^{1.5}(T)}{\epsilon \gamma H},$$

where the expectation is taken with respect to the randomness of the algorithm.

Proof. We follow the proof of Lemma 8 in [26].

Let $J_t = v_t^\top x + \frac{H}{2}||x||^2$. Let $\xi_t = v_t - \sum_{\tau=1}^t \hat{g}_t$ be the noise added by TBAP to $\sum_{\tau=1}^t \hat{g}_t$. Notice that $x_{t+1} = \arg\min_{x\in\mathcal{K}} J_t(x) + \xi_t^\top x$ and $\hat{x}_t = \arg\min_{x\in\mathcal{K}} J_t(x)$. Since J_t is H-strongly convex we have that

$$||\hat{x}_{t+1} - x_{t+1}|| \leq \frac{2||\xi_t||}{H}.$$

Since each ξ_t is formed in TBAP by adding at most $\lceil \ln(T)+1 \rceil$ vectors with norms drawn from a Gamma distribution with scale n and shape $\frac{(\lceil \ln(T)+1 \rceil) 4Mn}{\sqrt{\gamma}\epsilon}$ we can upper bound $\mathbb{E}[||\xi_t||]$ by $\frac{16 n \ln^{1.5}(T) M n}{\epsilon\sqrt{\gamma}}$.

Since $\hat{g}_t^\top$ is $\frac{4Mn}{\sqrt{\gamma}}$-Lipschitz continuous (by Lemma 4, concavity of $\sqrt{\cdot}$, and Jensen's inequality) we have that,

$$\begin{aligned}
\mathbb{E}[\sum_{t=1}^{T} \hat{g}_t^\top x_t] &\leq \mathbb{E}[\sum_{t=1}^{T} \hat{g}_t^\top \hat{x}_t] + \mathbb{E}[\frac{4Mn}{\sqrt{\gamma}} \sum_{t=1}^{T} \frac{2||\xi_t||}{H}] \\
&\leq \mathbb{E}[\sum_{t=1}^{T} \hat{g}_t^\top \hat{x}_t] + \frac{64 n^3 M^2 T \ln^{1.5}(T)}{\epsilon \gamma H}.
\end{aligned}$$

□

Theorem 9 (Regret guarantee)**.** BanditSubmodPRFTL($\{f_i\}_{i=1}^T, M, H, L, [n], \epsilon, \gamma$) *run with $H = MT^{2/3}$, $L = \frac{4Mn}{\sqrt{\gamma}}$, and $\gamma = \frac{n^{3/2}}{T^{1/3}}$ for any sequence of submodular functions $f_1, \ldots, f_T : 2^{[n]} \to [-M, M]$ for any $M, n, T > 0$ guarantees,*

$$\mathbb{E}[\text{Regret}(T)] \leq \tilde{O}\left(\frac{MnT^{2/3}}{\epsilon}\right),$$

where the expectation is taken with respect to all the internal randomness of the algorithm.

Proof of Theorem 9.

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{T} f_t(S_t) - \min_{S\subseteq[n]} \sum_{t=1}^{T} f_t(S)\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T} f_t(S_t) - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{f}_t(x)\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T} \hat{f}_t(x_t) - \sum_{t=1}^{T} \hat{f}_t(x)\right] + 2\gamma MT \quad \text{(for any } x\in\mathcal{K} \text{ by Lemma 5)} \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T} \nabla \hat{f}_t^\top (x_t - x)\right] + 2\gamma MT \quad \text{(since } \hat{f}_t \text{ is convex)} \\
&\leq \mathbb{E}\left[\sum_{t=1}^{T} \nabla \hat{f}_t^\top x_t\right] - \mathbb{E}\left[\min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{g}_t^\top x\right] + 2\gamma MT + \frac{8Mn\sqrt{T}}{\sqrt{\gamma}} \quad \text{(by Lemma 6)} \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \hat{g}_t^\top x_t\right] - \mathbb{E}\left[\min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{g}_t^\top x\right] + 2\gamma MT + \frac{8Mn\sqrt{T}}{\sqrt{\gamma}} \quad \text{(by Lemma 7)} \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \hat{g}_t^\top \hat{x}_t\right] - \mathbb{E}\left[\min_{x\in\mathcal{K}} \sum_{t=1}^{T} \hat{g}_t^\top x\right] + 2\gamma MT + \frac{8Mn\sqrt{T}}{\sqrt{\gamma}} + \frac{64n^3M^2T\ln^{1.5}(T)}{\epsilon\gamma H} \quad \text{(by Lemma 8)} \\
&\leq \mathbb{E}\left[\frac{2}{H}\sum_{t=1}^{T} ||\hat{g}_t||_2^2 + \frac{H}{2}[||x||_2^2 - ||x_1||_2^2]\right] + 2\gamma MT + \frac{8Mn\sqrt{T}}{\sqrt{\gamma}} + \frac{64n^3M^2T\ln^{1.5}(T)}{\epsilon\gamma H} \quad \text{(for any } x\in\mathcal{K} \text{ by Theorem 4)} \\
&\leq \frac{32M^2n^2T}{H\gamma} + nH + 2\gamma MT + \frac{8Mn\sqrt{T}}{\sqrt{\gamma}} + \frac{64n^3M^2T\ln^{1.5}(T)}{\epsilon\gamma H} \quad \text{(by Lemma 4)}
\end{aligned}
$$

Choosing $\gamma = \frac{n^{3/2}}{T^{1/3}}$, $H = MT^{2/3}$ yields the result. □