Supplementary Material for "Efficient Regret Minimization in Non-Convex Games"

Elad Hazan [1] **Karan Singh** [1] **Cyril Zhang** [1]

A. Proof of Theorem 4.4

Since each f_t is β-smooth, it follows that each F_t is β-smooth. Define $\widehat{\nabla f_t} = \frac{x_t - x_{t+1}}{\eta}$. Note that since the iterates $(x_t : t \in [T])$ depend on the gradient estimates, the iterates are stochastic variables, as are $\widehat{\nabla f_t}$. By β-smoothness of F_t, we have

$$
\begin{aligned}
&F_{t,w}(x_{t+1}) - F_{t,w}(x_t) \\
\leq\ & \langle \nabla F_{t,w}(x_t), x_{t+1} - x_t \rangle + \frac{\beta}{2}\|x_{t+1} - x_t\|^2 \\
=\ & -\eta \left\langle \nabla F_{t,w}(x_t), \widehat{\nabla f_t} \right\rangle + \eta^2 \frac{\beta}{2} \|\widehat{\nabla f_t}\|^2 \\
=\ & -\eta \|\nabla F_{t,w}(x_t)\|^2 - \eta \left\langle \nabla F_{t,w}(x_t), \widehat{\nabla f_t} - \nabla F_{t,w}(x_t) \right\rangle \\
&+ \eta^2 \frac{\beta}{2} \left(\|\nabla F_{t,w}(x_t)\|^2\right) \\
&+ \eta^2 \frac{\beta}{2} \left(2 \left\langle \nabla F_{t,w}(x_t), \widehat{\nabla f_t} - \nabla F_{t,w}(x_t) \right\rangle\right) \\
&+ \eta^2 \frac{\beta}{2} \left(\|\widehat{\nabla f_t} - \nabla F_{t,w}(x_t)\|^2\right) \\
=\ & -\left(\eta - \frac{\beta}{2}\eta^2\right) \|\nabla F_{t,w}(x_t)\|^2 \\
&- (\eta - \beta\eta^2) \left\langle \nabla F_{t,w}(x_t), \widehat{\nabla f_t} - \nabla F_{t,w}(x_t) \right\rangle \\
&+ \eta^2 \frac{\beta}{2} \|\widehat{\nabla f_t} - \nabla f(x_t)\|^2.
\end{aligned}
$$

Additionally, we each observe that $\widehat{\nabla f_t}$ is an average of w independently sampled unbiased gradient estimates of variance σ^2 each. It follows as a consequence that

$$
\mathbb{E}\big[\widehat{\nabla f_t} \big| x_t\big] = \nabla F_{t,w}(x_t)
$$
$$
\mathbb{E}\big[\|\widehat{\nabla f_t} - \nabla F_{t,w}(x_t)\|^2 \big| x_t\big] \leq \frac{\sigma^2}{w}
$$

Now, applying $\mathbb{E}\left[\cdot | x_t\right]$ on both sides, it follows that

$$
\begin{aligned}
&\left(\eta - \frac{\beta}{2}\eta^2\right) \cdot \mathbb{E}\|\nabla F_{t,w}(x_t)\|^2 \\
&\qquad \leq \mathbb{E}\left[F_{t,w}(x_t) - F_{t,w}(x_{t+1})\right] + \eta^2 \frac{\beta}{2} \frac{\sigma^2}{w}.
\end{aligned}
$$

Also, we note that

$$
\begin{aligned}
&F_{t+1,w}(x_{t+1}) - F_{t,w}(x_{t+1}) \\
=&\frac{1}{w}\sum_{i=0}^{w-1} f_{t+1-i}(x_{t+1}) - \frac{1}{w}\sum_{i=0}^{w-1} f_{t-i}(x_{t+1}) \\
=&\frac{1}{w}\sum_{i=-1}^{w-2} f_{t-i}(x_{t+1}) - \frac{1}{w}\sum_{i=0}^{w-1} f_{t-i}(x_{t+1}) \\
=&\frac{f_{t+1}(x_{t+1}) - f_{t-w+1}(x_{t+1})}{w} \leq \frac{2M}{w}
\end{aligned}
$$

Adding the last two inequalities, we proceed to sum the above inequality over all time steps:

$$
\mathbb{E}\left[\sum_{t=1}^{T} \|\nabla F_{t,w}(x_t)\|^2\right] \leq \frac{2M + \frac{2MT}{w} + \frac{T\beta\eta^2}{2w}\sigma^2}{\eta - \frac{\beta\eta^2}{2}}.
$$

Setting $\eta = 1/\beta$ yields the claim from the theorem.

Finally, note that for each round the number of stochastic gradient oracle calls required is w. Therefore, across all T rounds, the number of noisy oracle calls is Tw. □

B. Proof of Theorem 5.1 (ii)

Following the technique from Theorem 3.1, for $2 \leq t \leq T$, let τ_t be the number of iterations of the inner loop during the execution of Algorithm 3 during round $t-1$ (in order to generate the iterate x_t). Then, we have the following lemma:

Lemma B.1. *For any* $2 \leq t \leq T$,

$$
F_{t-1}(x_t) - F_{t-1}(x_{t-1}) \leq -\tau_t \cdot \frac{\delta^3}{2\beta w^3}.
$$

Proof. This follows by summing the inequality Lemma 5.3 for across all pairs of consecutive iterates of the inner loop

[1]Computer Science, Princeton University. Correspondence to: Elad Hazan <ehazan@princeton.edu>, Karan Singh <karans@princeton.edu>, Cyril Zhang <cyril.zhang@princeton.edu>.

Proceedings of the 34th International Conference on Machine Learning, Sydney, Australia, PMLR 70, 2017.

within the same epoch, and noting that each term $\Phi(z)$ is at least $\frac{\delta^3}{w^3}$ before the inner loop has terminated. $\square$

Finally, we write (understanding $F_0(x_0) := 0$):

$$\begin{aligned}
F_T(x_T) &= \sum_{t=1}^{T} F_t(x_t) - F_{t-1}(x_{t-1}) \\
&= \sum_{t=1}^{T} F_{t-1}(x_t) - F_{t-1}(x_{t-1}) + f_t(x_t) - f_{t-w}(x_t) \\
&\leq \sum_{t=2}^{T} \left[F_{t-1}(x_t) - F_{t-1}(x_{t-1})\right] + \frac{2MT}{w}.
\end{aligned}$$

Using Lemma B.1, we have

$$F_T(x_T) \leq \frac{2MT}{w} - \frac{\delta^3}{2\beta w^3} \cdot \sum_{t=1}^{T} \tau_t,$$

whence

$$\begin{aligned}
\tau = \sum_{t=1}^{T} \tau_t &\leq \frac{2\beta w^3}{\delta^3} \cdot \left(\frac{2MT}{w} - F_T(x_T)\right) \\
&\leq \frac{2\beta M}{\delta^3} \cdot \left(2Tw^2 + w^3\right) \\
&\leq \frac{6M}{\beta^2} \cdot Tw^2,
\end{aligned}$$

as claimed (recalling that we chose $\delta = \beta$ for this analysis). $\square$

C. Proof of Theorem 6.2

Summing up the definitions of w-regret bounds achieved by each $\mathcal{A}$, and truncating the first $w - 1$ terms, we get

$$\sum_{i=1}^{k} \sum_{t=w}^{T} \|\nabla_{\mathcal{K},\eta} F_t^i(x_t^i)\|^2 \leq \sum_{i=1}^{k} \mathfrak{R}_{w,\mathcal{A}_i}(T).$$

Thus, for some t between w and T inclusive, it holds that

$$\begin{aligned}
\sum_{i=1}^{k} \left\| \nabla_{\mathcal{K},\eta} \left[\frac{\sum_{j=0}^{w-1} \tilde{f}_{i,t-j}}{w} \right] (x_t^i) \right\|^2 &= \sum_{i=1}^{k} \|\nabla_{\mathcal{K},\eta} F_t^i(x_t^i)\|^2 \\
&\leq \sum_{i=1}^{k} \frac{\mathfrak{R}_{w,\mathcal{A}_i}(T)}{T - w}.
\end{aligned}$$

Thus, for the same t we have

$$\max_{i \in [k]} \left\| \nabla_{\mathcal{K},\eta} \left[\frac{\sum_{j=0}^{w-1} \tilde{f}_{i,t-j}}{w} \right] (x_t^i) \right\| \leq \sqrt{\sum_{i=1}^{k} \frac{\mathfrak{R}_{w,\mathcal{A}_i}(T)}{T - w}},$$

as claimed. $\square$