Online Bidding Algorithms for Return-on-Spend Constrained Advertisers

Joint work with
Zhe Feng [1] and Di Wang[1]

[1]Google Research, Mountainview

July 17, 2023

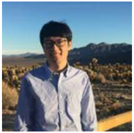 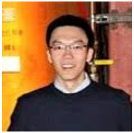

Problem Introduction

- Online advertising: a multi-billion dollar industry
- Emergence of optimization algorithms for bidding
- This talk: value maximization for the single bidder under the return-on-spend and fixed budget constraints

Problem Setup: Online Bidding under RoS Constraint

We study online bidding for a single agent. At each time step t,

Problem Setup: Online Bidding under RoS Constraint

We study online bidding for a single agent. At each time step t,

- the agent sees an ad query with its associated value v_t,

stochastically generated by Nature

Problem Setup: Online Bidding under RoS Constraint

We study online bidding for a single agent. At each time step t,

- the agent sees an ad query with its associated value v_t,
- submits their bid b_t,

Problem Setup: Online Bidding under RoS Constraint

We study online bidding for a single agent. At each time step t,

- the agent sees an ad query with its associated value v_t,
- submits their bid b_t,
- and then sees the allocation $x_t(b_t)$ and price $p_t(b_t)$.

truthful auction

$$p_t(b) = b \cdot x_t(b) - \int_{z=0}^{b} x_t(z)dz$$

Myerson (1981)

Problem Setup: Formalized

Our goal is to solve, with minimum regret and minimum constraint violation, the online bidding problem:

$$
\begin{aligned}
\underset{b_t:t=1,\cdots,T}{\text{maximize}} \quad & \sum_{t=1}^{T} v_t \cdot x_t(b_t) \\
\text{subject to} \quad & \sum_{t=1}^{T} p_t(b_t) \leq \sum_{t=1}^{T} v_t \cdot x_t(b_t), \\
& \sum_{t=1}^{T} p_t(b_t) \leq \rho T.
\end{aligned}
$$

RoS constraint

Problem Setup: Formalized

Our goal is to solve, with minimum regret and minimum constraint violation, the online bidding problem:

$$\begin{aligned} \underset{b_t: t=1,\cdots,T}{\text{maximize}} \quad & \sum_{t=1}^{T} v_t \cdot x_t(b_t) \\ \text{subject to} \quad & \sum_{t=1}^{T} p_t(b_t) \leq \sum_{t=1}^{T} v_t \cdot x_t(b_t), \\ & \sum_{t=1}^{T} p_t(b_t) \leq \rho T. \end{aligned}$$

RoS constraint

- Model first proposed by Mannor, Tsitsiklis, Yu (2009)

Problem Setup: Formalized

Our goal is to solve, with minimum regret and minimum constraint violation, the online bidding problem:

$$\begin{aligned} \underset{b_t:t=1,\cdots,T}{\text{maximize}} \quad & \sum_{t=1}^T v_t \cdot x_t(b_t) \\ \text{subject to} \quad & \underbrace{\sum_{t=1}^T p_t(b_t) \leq \sum_{t=1}^T v_t \cdot x_t(b_t)}_{\text{RoS constraint}}, \\ & \sum_{t=1}^T p_t(b_t) \leq \rho T. \end{aligned}$$

- Model first proposed by Mannor, Tsitsiklis, Yu (2009)
- The RoS constraint is **non-packing**; hence, inapplicability of:
  - bandits-with-knapsacks (e.g., Immorlica, Sankararaman, Schapire, Slivkins (2022), Castiglioni, Celli, Kroer (2022))
  - allocation-under-resource-consumption (e.g., Balseiro, Lu, Mirrokni (2020)).

Problem Setup: Formalized

Our goal is to solve, with minimum regret and minimum constraint violation, the online bidding problem:

RoS constraint

$$\begin{aligned} \underset{b_t : t=1,\cdots,T}{\text{maximize}} \quad & \sum_{t=1}^{T} v_t \cdot x_t(b_t) \\ \text{subject to} \quad & \sum_{t=1}^{T} p_t(b_t) \leq \sum_{t=1}^{T} v_t \cdot x_t(b_t), \\ & \sum_{t=1}^{T} p_t(b_t) \leq \rho T. \end{aligned}$$

We define regret relative to the best *adaptive* strategy in hindsight:

$$\text{Regret}(\text{Alg}, \overrightarrow{\gamma}) := \mathbb{E}_{\gamma \sim \mathcal{P}^T} \left[\text{Reward}(\text{Opt}, \overrightarrow{\gamma}) - \text{Reward}(\text{Alg}, \overrightarrow{\gamma}) \right].$$

$\max_{\{b_t \in \mathcal{B}\}} \sum_{t \leq T} v_t \cdot x_t(b_t)$

Our Main Result

Theorem 1: Our Main Result (Informal)

For a T-length input i.i.d. sequence of ad queries, our algorithm provably attains, under a mild technical assumption, $O(\sqrt{T}\log T)$ regret while respecting both the budget and RoS constraints.

Our Main Result

Theorem 1: Our Main Result (Informal)

For a T-length input i.i.d. sequence of ad queries, our algorithm provably attains, under a mild technical assumption, $O(\sqrt{T}\log T)$ regret while respecting both the budget and RoS constraints.

- First algorithm to attain near-optimal regret while satisfying both budget and RoS constraints in any outcome.
- In doing so, we improve upon the prior work of Balseiro, Lu, and Mirrokni (2020), which obtains similar guarantees under only budget constraints.

Related Work

- Balseiro, Lu, Mirrokni (2020): fixed-budget constraints
- Castiglioni, Celli, Marchesi, Romano, Gatti (2022): more general, with weaker guarantees for RoS
- Agrawal and Devanur (2014): requires bounded dual space
- Golrezaei, Jaillet, Liang, Mirrokni (2021): constraints hold in expectation
- Other related work:
  - AdWords: Mehta, Saberi, Vazirani, Vazirani (2007)
  - Learning to bid in repeated auctions: Borgs, Chayes, Immorlica, Jain, Etesami, Mahdian (2007), Weed, Perchet, Rigollet (2016), Feng, Podimata, Syrgkanis (2018), Badanidiyuru, Feng, Guruganesh (2021)

Our Techniques

Our Algorithmic Outline: RoS Constraints

We first solve the problem with only the RoS constraint:

$$\begin{aligned} \underset{b_t : t=1,\cdots,T}{\text{maximize}} \quad & \sum_{t=1}^{T} v_t \cdot x_t(b_t) \\ \text{subject to} \quad & \sum_{t=1}^{T} p_t(b_t) \le \sum_{t=1}^{T} v_t \cdot x_t(b_t). \end{aligned}$$

We use the primal-dual framework similar to that in Balseiro, Lu, and Mirrokni (2020), which lends our algorithm adaptivity to changing input values, with **Online Mirror Descent (OMD)** to update the dual variable (which tracks the constraint violation).

Our Algorithmic Outline: RoS Constraints

We first solve the problem with only the RoS constraint:

$$\begin{aligned} \underset{b_t:t=1,\cdots,T}{\text{maximize}} \quad & \textstyle\sum_{t=1}^{T} v_t \cdot x_t(b_t) \\ \text{subject to} \quad & \textstyle\sum_{t=1}^{T} p_t(b_t) \leq \sum_{t=1}^{T} v_t \cdot x_t(b_t). \end{aligned}$$

Our update rule solves

$$\text{maximize}_{\{b_i\}} \left\{ \sum_{i=1}^{T} v_i \cdot x_i(b_i) + \min_{\lambda \geq 0} \left[\lambda \cdot \sum_{i=1}^{T} g_i(b_i) + \frac{1}{\alpha} h(\lambda) \right] \right\}$$

where

- $g_i(b) := v_i \cdot x_i(b) - p_i(b)$ measures the constraint satisfaction, and
- h is **generalized negative entropy**, which imposes a large (exponential) penalty on constraint violation.

Our Algorithm's Updates: Approximate RoS Constraints

$$\text{maximize}_{\{b_i\}} \left\{ \sum_{i=1}^{T} v_i \cdot x_i(b_i) + \min_{\lambda \geq 0} \left[\lambda \cdot \sum_{i=1}^{T} g_i(b_i) + \frac{1}{\alpha} h(\lambda) \right] \right\}$$

Our Algorithm's Updates: Approximate RoS Constraints

$$\text{maximize}_{\{b_i\}} \left\{ \sum_{i=1}^{T} v_i \cdot x_i(b_i) + \min_{\lambda \geq 0} \left[\lambda \cdot \sum_{i=1}^{T} g_i(b_i) + \frac{1}{\alpha} h(\lambda) \right] \right\}$$

- Bid b_t maximizes the current penalty-adjusted reward:

$$b_t = \arg\max_{b \geq 0} \left[\frac{1+\lambda_t}{\lambda_t} \cdot v_t \cdot x_t(b) - p_t(b) \right] = v_t + \frac{v_t}{\lambda_t}.$$

truthful auction

Our Algorithm's Updates: Approximate RoS Constraints

$$\text{maximize}_{\{b_i\}} \left\{ \sum_{i=1}^{T} v_i \cdot x_i(b_i) + \min_{\lambda \geq 0} \left[\lambda \cdot \sum_{i=1}^{T} g_i(b_i) + \frac{1}{\alpha} h(\lambda) \right] \right\}$$

- Bid b_t maximizes the current penalty-adjusted reward:

$$b_t = \arg\max_{b \geq 0} \left[\frac{1+\lambda_t}{\lambda_t} \cdot v_t \cdot x_t(b) - p_t(b) \right] = v_t + \frac{v_t}{\lambda_t}.$$

- Dual variable λ_t penalizes (rewards) constraint violation (satisfaction):

$$\lambda_{t+1} = \arg\min_{\lambda \geq 0} \left[g_t(b_t) \cdot \lambda + \frac{1}{\alpha} V_h(\lambda, \lambda_t) \right] = \lambda_1 \cdot \exp \left[- \sum_{i \leq t} \alpha \cdot g_i(b_i) \right]$$

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\overrightarrow{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\overrightarrow{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

for $t = 1, 2, \cdots, T$

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\overrightarrow{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

for $t = 1, 2, \cdots, T$

- Observe the value v_t, and set the bid $b_t = v_t + \frac{v_t}{\lambda_t}$.

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\vec{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

for $t = 1, 2, \cdots, T$

- Observe the value v_t, and set the bid $b_t = v_t + \frac{v_t}{\lambda_t}$.
- Observe the price $p_t(b_t)$ and allocation $x_t(b_t)$

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\overrightarrow{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

for $t = 1, 2, \cdots, T$

- Observe the value v_t, and set the bid $b_t = v_t + \frac{v_t}{\lambda_t}$.
- Observe the price $p_t(b_t)$ and allocation $x_t(b_t)$
- Compute the gradient $g_t(b_t) = v_t \cdot x_t(b_t) - p_t(b_t)$.

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\overrightarrow{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

for $t = 1, 2, \cdots, T$

- Observe the value v_t, and set the bid $b_t = v_t + \frac{v_t}{\lambda_t}$.
- Observe the price $p_t(b_t)$ and allocation $x_t(b_t)$
- Compute the gradient $g_t(b_t) = v_t \cdot x_t(b_t) - p_t(b_t)$.
- Update the dual variable as $\lambda_{t+1} = \lambda_t \cdot \exp\left[-\alpha \cdot g_t(b_t)\right]$.

Our Algorithm: Approximate RoS Constraints

Input. Time horizon T and ad requests $\overrightarrow{\gamma} \sim \mathcal{P}^T$.

Initialize step size $T^{-1/2}$ and dual variable $\lambda_1 = 1$.

for $t = 1, 2, \cdots, T$

- Observe the value v_t, and set the bid $b_t = v_t + \frac{v_t}{\lambda_t}$.
- Observe the price $p_t(b_t)$ and allocation $x_t(b_t)$
- Compute the gradient $g_t(b_t) = v_t \cdot x_t(b_t) - p_t(b_t)$.
- Update the dual variable as $\lambda_{t+1} = \lambda_t \cdot \exp\left[-\alpha \cdot g_t(b_t)\right]$.

Return the sequence $\{b_t\}_{t=1}^T$ of generated bids.

Bound on Approximate RoS Violation

Theorem 2: RoS Constraint Violation Bound (Informal)

Our algorithm's RoS constraint violation at most $O(\sqrt{T}\log T)$.

Bound on Approximate RoS Violation

Theorem 2: RoS Constraint Violation Bound (Informal)

Our algorithm's RoS constraint violation at most $O(\sqrt{T}\log T)$.

The proof idea is as follows.

Bound on Approximate RoS Violation

Theorem 2: RoS Constraint Violation Bound (Informal)

Our algorithm's RoS constraint violation at most $O(\sqrt{T}\log T)$.

The proof idea is as follows.

- When the cumulative violation $-\sum_{i\le t} g_i(b_i) \gg \alpha^{-1}$, the update rule $\lambda_{t+1} = \lambda_1 \cdot \exp(-\alpha \sum_{i\le t} g_i(b_i))$ makes λ_{t+1} huge,

Bound on Approximate RoS Violation

Theorem 2: RoS Constraint Violation Bound (Informal)

Our algorithm's RoS constraint violation at most $O(\sqrt{T}\log T)$.

The proof idea is as follows.

- When the cumulative violation $-\sum_{i\le t} g_i(b_i) \gg \alpha^{-1}$, the update rule $\lambda_{t+1} = \lambda_1 \cdot \exp(-\alpha \sum_{i\le t} g_i(b_i))$ makes λ_{t+1} huge,
- ... and as a result, the update $b_{t+1} = v_{t+1} + \frac{v_{t+1}}{\lambda_{t+1}}$ prevents us from overbidding.

Bound on Regret

The primal-dual framework implies

$$\mathrm{Regret}(\mathrm{Alg}, \overrightarrow{\gamma}) \leq \mathbb{E}_{\overrightarrow{\gamma} \sim \mathcal{P}^T}\left[\sum_{t \in [T]} \lambda_t \cdot g_t(b_t)\right].$$

Bound on Regret

The primal-dual framework implies

$$\mathrm{Regret}(\mathrm{Alg}, \overrightarrow{\gamma}) \leq \mathbb{E}_{\overrightarrow{\gamma} \sim \mathcal{P}^T} \left[\sum_{t \in [T]} \lambda_t \cdot g_t(b_t) \right].$$

Bounding the R.H.S. requires novel analysis due to the existence of input instances for which λ_t can be huge.

different from fixed budget setting

Bound on Regret

The primal-dual framework implies

$$\mathrm{Regret}(\mathrm{Alg}, \overrightarrow{\gamma}) \leq \mathbb{E}_{\overrightarrow{\gamma} \sim \mathcal{P}^T} \left[\sum_{t \in [T]} \lambda_t \cdot g_t(b_t) \right] .$$

Bounding the R.H.S. requires novel analysis due to the existence of input instances for which λ_t can be huge.

Our novel insight to bound $\sum_{t \leq T} \lambda_t \cdot g_t(b_t)$ is to combine:

- structural properties of the gradient
- with a white-box OMD analysis
- and tools developed by Allen-Zhu and Orecchia (2015) for solving positive linear programs.

Towards a Regret Bound: Bounding $\sum_{t\leq T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t\leq T} g_t \cdot \lambda_t \leq O(\sqrt{T})$.

Towards a Regret Bound: Bounding $\sum_{t \le T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t \le T} g_t \cdot \lambda_t \le O(\sqrt{T})$.

Proof. We have,

Towards a Regret Bound: Bounding $\sum_{t\leq T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t\leq T} g_t \cdot \lambda_t \leq O(\sqrt{T})$.

Proof. We have,

$$\alpha g_t \lambda_t = \alpha g_t (\lambda_t - \lambda_{t+1}) + \alpha g_t \lambda_{t+1}$$

adding/subtracting

$\alpha g_t \lambda_{t+1}$

Towards a Regret Bound: Bounding $\sum_{t \leq T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t \leq T} g_t \cdot \lambda_t \leq O(\sqrt{T})$.

Proof. We have,

$$\alpha g_t \lambda_t = \alpha g_t(\lambda_t - \lambda_{t+1}) + \alpha g_t \lambda_{t+1} = \alpha g_t(\lambda_t - \lambda_{t+1}) + \lambda_{t+1} \log(\lambda_t / \lambda_{t+1})$$

dual update step

Towards a Regret Bound: Bounding $\sum_{t \le T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t \le T} g_t \cdot \lambda_t \le O(\sqrt{T})$.

Proof. We have,

$$\alpha g_t \lambda_t = \alpha g_t(\lambda_t - \lambda_{t+1}) + \alpha g_t \lambda_{t+1} = \alpha g_t(\lambda_t - \lambda_{t+1}) + \lambda_{t+1} \log(\lambda_t / \lambda_{t+1})$$
$$= (1+\alpha)(\lambda_t - \lambda_{t+1}) - \mathcal{V}^h_{\lambda_t}(\lambda_{t+1})$$

introducing Bregman divergence

Towards a Regret Bound: Bounding $\sum_{t\leq T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t\leq T} g_t \cdot \lambda_t \leq O(\sqrt{T})$.

Proof. We have,

$$\alpha g_t \lambda_t = \alpha g_t(\lambda_t - \lambda_{t+1}) + \alpha g_t \lambda_{t+1} = \alpha g_t(\lambda_t - \lambda_{t+1}) + \lambda_{t+1}\log(\lambda_t/\lambda_{t+1})$$

$$= (1+\alpha)(\lambda_t - \lambda_{t+1}) - \mathcal{V}^h_{\lambda_t}(\lambda_{t+1})$$

$$\leq (1+\alpha)(\lambda_t - \lambda_{t+1}) - \frac{(\lambda_t - \lambda_{t+1})^2}{2\max(\lambda_t, \lambda_{t+1})}$$

local strong convexity

Allen-Zhu & Orecchia

Towards a Regret Bound: Bounding $\sum_{t\leq T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t\leq T} g_t \cdot \lambda_t \leq O(\sqrt{T})$.

Proof. We have,

$$
\begin{aligned}
\alpha g_t \lambda_t = \alpha g_t(\lambda_t - \lambda_{t+1}) + \alpha g_t \lambda_{t+1} &= \alpha g_t(\lambda_t - \lambda_{t+1}) + \lambda_{t+1}\log(\lambda_t/\lambda_{t+1}) \\
&= (1+\alpha)(\lambda_t - \lambda_{t+1}) - \mathcal{V}^h_{\lambda_t}(\lambda_{t+1}) \\
&\leq (1+\alpha)(\lambda_t - \lambda_{t+1}) - \frac{(\lambda_t - \lambda_{t+1})^2}{2\max(\lambda_t, \lambda_{t+1})} \\
&\leq \lambda_t - \lambda_{t+1} + \frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})
\end{aligned}
$$

completing the square

Towards a Regret Bound: Bounding $\sum_{t\le T}\lambda_t\cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t\le T} g_t\cdot\lambda_t \le O(\sqrt{T})$.

Proof. We have,

$$
\begin{aligned}
\alpha g_t\lambda_t &= \alpha g_t(\lambda_t-\lambda_{t+1}) + \alpha g_t\lambda_{t+1} = \alpha g_t(\lambda_t-\lambda_{t+1}) + \lambda_{t+1}\log(\lambda_t/\lambda_{t+1}) \\
&= (1+\alpha)(\lambda_t-\lambda_{t+1}) - \mathcal{V}^h_{\lambda_t}(\lambda_{t+1}) \\
&\le (1+\alpha)(\lambda_t-\lambda_{t+1}) - \frac{(\lambda_t-\lambda_{t+1})^2}{2\max(\lambda_t,\lambda_{t+1})} \\
&\le \lambda_t-\lambda_{t+1} + \frac{1}{2}\alpha^2 g_t^2\max(\lambda_t,\lambda_{t+1})
\end{aligned}
$$

Towards a Regret Bound: Bounding $\sum_{t\leq T} \lambda_t \cdot g_t$

Theorem 3: Key Regret Bound Lemma (Informal)

Our algorithm's iterates satisfy $\sum_{t\leq T} g_t \cdot \lambda_t \leq O(\sqrt{T})$.

Proof. We have,

$$\begin{aligned}
\alpha g_t \lambda_t &= \alpha g_t(\lambda_t - \lambda_{t+1}) + \alpha g_t \lambda_{t+1} = \alpha g_t(\lambda_t - \lambda_{t+1}) + \lambda_{t+1}\log(\lambda_t/\lambda_{t+1}) \\
&= (1+\alpha)(\lambda_t - \lambda_{t+1}) - \mathcal{V}^h_{\lambda_t}(\lambda_{t+1}) \\
&\leq (1+\alpha)(\lambda_t - \lambda_{t+1}) - \frac{(\lambda_t - \lambda_{t+1})^2}{2\max(\lambda_t, \lambda_{t+1})} \\
&\leq \lambda_t - \lambda_{t+1} + \frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})
\end{aligned}$$

We now provide an upper bound on $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$.

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t$$

$g_t \geq 0$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t))$$

$g_t \leq 1$, $\alpha = T^{-1/2}$

$\exp(-x) \leq 1 - x/2$ for $x \in [0, 1.5]$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

dual update rule

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1}$$

$g_t \leq 0$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1} \leq \alpha^2 \cdot \frac{\lambda_{t+1}}{\lambda_t}$$

$0 \geq g_t \geq \max(-1, -1/\lambda_t)$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1} \leq \alpha^2 \cdot \frac{\lambda_{t+1}}{\lambda_t} \leq \alpha^2 \cdot \exp(-\alpha g_t)$$

dual update rule

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1} \leq \alpha^2 \cdot \frac{\lambda_{t+1}}{\lambda_t} \leq \alpha^2 \cdot \exp(-\alpha g_t) \leq 2\alpha^2$$

$g_t \geq -1, \alpha = T^{-1/2}$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1} \leq \alpha^2 \cdot \frac{\lambda_{t+1}}{\lambda_t} \leq \alpha^2 \cdot \exp(-\alpha g_t) \leq 2\alpha^2$$

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1} \leq \alpha^2 \cdot \frac{\lambda_{t+1}}{\lambda_t} \leq \alpha^2 \cdot \exp(-\alpha g_t) \leq 2\alpha^2$$

Therefore, we have $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) \leq \alpha(\lambda_t - \lambda_{t+1}) + 2\alpha^2$.

Regret Bound Continued: Bounding $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1})$

Lemma 1

Our gradients g_t satisfy the bound $\max(-1, -1/\lambda_t) \leq g_t \leq v_t x_t$

Case 1. Assume $g_t \geq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_t \leq \alpha(\lambda_t - \lambda_t \exp(-\alpha g_t)) \leq \alpha(\lambda_t - \lambda_{t+1})$$

Case 2. Assume $g_t \leq 0$. Then,

$$\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) = \frac{1}{2}\alpha^2 g_t^2 \lambda_{t+1} \leq \alpha^2 \cdot \frac{\lambda_{t+1}}{\lambda_t} \leq \alpha^2 \cdot \exp(-\alpha g_t) \leq 2\alpha^2$$

Therefore, we have $\frac{1}{2}\alpha^2 g_t^2 \max(\lambda_t, \lambda_{t+1}) \leq \alpha(\lambda_t - \lambda_{t+1}) + 2\alpha^2$.
Summing over $t \in [T]$ and dividing by α finishes the proof. ■

Putting It All Together

To get strict RoS satisfaction, we propose a simple idea:

- First, submit a sequence of bids so as to accumulate a slack on the RoS constraint.
- Next, run the existing algorithm, which suffers some bounded constraint violation, which is compensated by the slack from the first phase.

Putting It All Together

To get strict RoS satisfaction, we propose a simple idea:

- First, submit a sequence of bids so as to accumulate a slack on the RoS constraint.
- Next, run the existing algorithm, which suffers some bounded constraint violation, which is compensated by the slack from the first phase.

To satisfy both RoS and budget constraints, we combine our ideas with those of Balseiro, Lu, and Mirrokni (2020).

Thank You!