Online Submodular Maximization via Online Convex Optimization

Tareq Si Salem,[1] Gözde Özcan,[1] Iasonas Nikolaou,[2] Evimaria Terzi,[2] Stratis Ioannidis[1]

[1]**Northeastern University**, Boston, MA, USA
[2]**Boston University**, Boston, MA, USA

Submodular Maximization over Matroids

Submodular maximization is NP-hard.

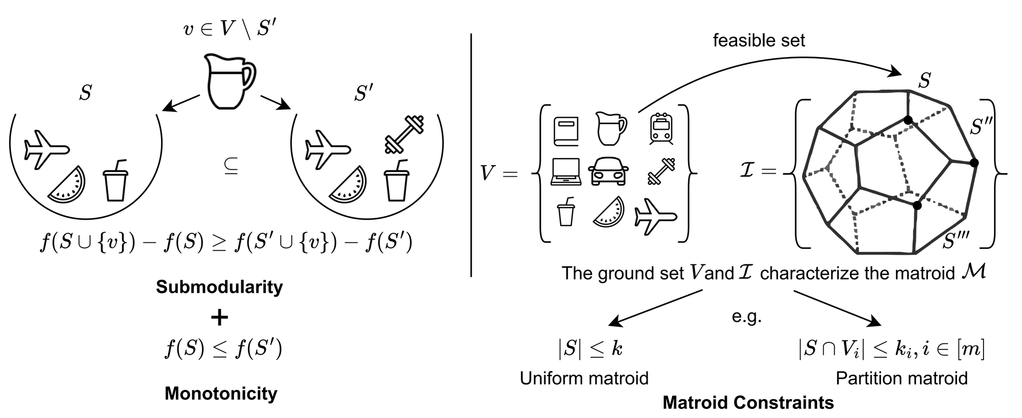

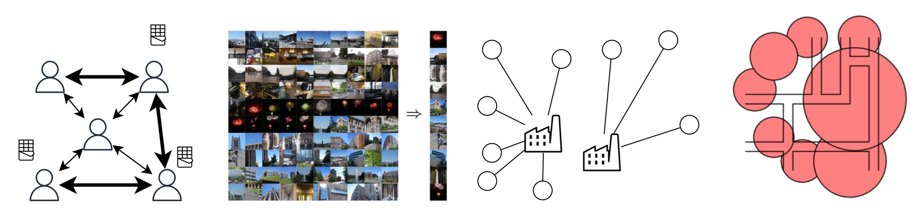

(a) Influence Maximization (b) Data Summarization (c) Facility Location (d) Sensor Placement

Online Learning and OCO

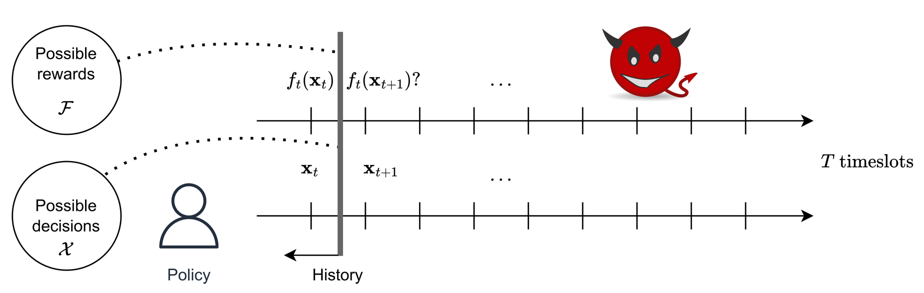

At each timeslot t:

- **Decision:** $\mathbf{x}_t \in \mathcal{X}$
- Adversary picks **reward function** $f_t : \mathcal{X} \to \mathbb{R}_{\geq 0}$ from class $\mathcal{F}$
- Decision maker accrues reward $f_t(\mathbf{x}_t)$
- Decisions based on history:
$$\mathbf{x}_t = \mathcal{P}_{\mathcal{X},t}\left((\mathbf{x}_s)_{s=1}^{t-1}, (f_s)_{s=1}^{t-1}\right)$$
- **Online policy:** $\mathcal{P}_{\mathcal{X}} = (\mathcal{P}_{\mathcal{X},t})_{t\in[T]}$

We define the **regret** at horizon T as:

$$\operatorname{regret}_T(\mathcal{P}_{\mathcal{X}}) = \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \max_{\mathbf{x}\in\mathcal{X}} \textstyle\sum_{t=1}^T f_t(\mathbf{x}) - \mathbb{E}_{\mathcal{P}_{\mathcal{X}}}\left[\sum_{t=1}^T f_t(\mathbf{x}_t)\right] \right\}.$$

In OCO, $\mathcal{X}$ is a **compact convex set** and $\mathcal{F}$ is a set of **concave functions**. Several online policies (e.g., **OGA**, **OMA**, and **FTRL**) attain $O(\sqrt{T})$ regret when functions in $\mathcal{F}$ are L-Lipschitz.

Online Submodular Maximization

In OSM, $\mathcal{X}$ is a **matroid** and $\mathcal{F}$ is a set of **monotone submodular functions**. The α-**regret** of policy $\mathcal{P}_{\mathcal{X}}$ at horizon T is:

$$\alpha\text{-regret}_T(\mathcal{P}_{\mathcal{X}}) \triangleq \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \max_{\mathbf{x}\in\mathcal{X}} \alpha \sum_{t=1}^T f_t(\mathbf{x}) - \mathop{\mathbb{E}}_{\mathcal{P}_{\mathcal{X}}}\left[\sum_{t=1}^T f_t(\mathbf{x}_t)\right] \right\}.$$

We seek policies that attain a **sublinear** (i.e., $o(T)$) regret. Intuitively, we compare the performance of the policy $\mathcal{P}_{\mathcal{X}}$ w.r.t. the best poly-time(n) α-approximation of the static offline optimum, in hindsight.

Main Contributions

- We provide a methodology for **reducing online submodular maximization (OSM) problems to online convex optimization (OCO) problems**, under a certain condition.
- We show that this condition is satisfied by a wide class of submodular functions called **weighted threshold potential (WTP) functions**.
- We establish that our reduction extends to **the dynamic regret and optimistic settings**.
- We also provide a different reduction for the **bandit setting**, for static, dynamic regret, and optimistic variants.

Case of General Set Functions

- We can **reduce** any **OSM problem** to an **OCO problem**, when the reward functions satisfy a specific property.

Assumption: Sandwich Property

There exists an $\alpha \in (0,1]$, an $L \in \mathbb{R}_{>0}$, and a **randomized rounding** $\Xi : \mathcal{Y} \to \mathcal{X}$ such that, for every $f : \mathcal{X} \to \mathbb{R}_{\geq 0} \in \mathcal{F}$ there exists a L-Lipschitz **concave function** $\tilde{f} : \mathcal{Y} \to \mathbb{R}$ s.t.

$$\tilde{f}(\mathbf{x}) \geq f(\mathbf{x}), \quad \text{for all } \mathbf{x} \in \mathcal{X}, \text{ and}$$
$$\mathbb{E}_{\Xi}[f(\Xi(\mathbf{y}))] \geq \alpha \cdot \tilde{f}(\mathbf{y}), \quad \text{for all } \mathbf{y} \in \mathcal{Y}.$$

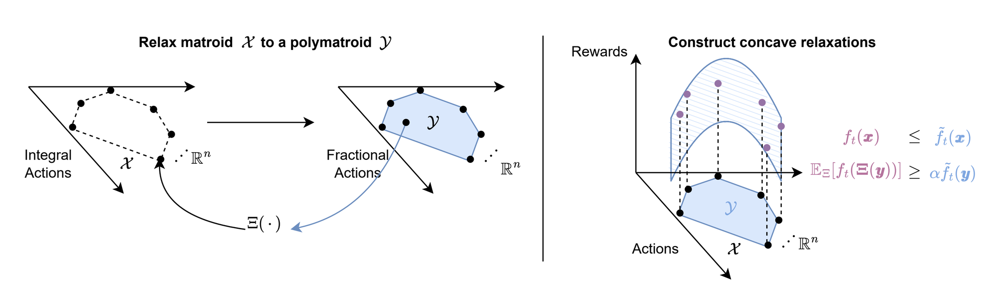

RAOCO for General Set Functions

If this property holds, we can convert any OCO policy $\mathcal{P}_{\mathcal{Y}}$ operating over $\mathcal{Y} = \operatorname{conv}(\mathcal{X})$ to a **Randomized-rounding Augmented OCO (RAOCO) policy**, denoted by $\mathcal{P}_{\mathcal{X}}$, operating over $\mathcal{X}$. The RAOCO policy amounts to (a) running $\mathcal{P}_{\mathcal{Y}}$ over the **concave relaxations** $\tilde{f}_t(\cdot)$, and (b) **rounding** the decision via Ξ. Formally, at $t \in [T]$:

$$\mathbf{y}_t = \mathcal{P}_{\mathcal{Y},t}\left((\mathbf{y}_s)_{s=1}^{t-1}, (\tilde{f}_s)_{s=1}^{t-1}\right),$$
$$\mathbf{x}_t = \Xi(\mathbf{y}_t) \in \mathcal{X}.$$

Theorem 1: *If the sandwich property holds, given an OCO policy $\mathcal{P}_{\mathcal{Y}}$, the RAOCO policy $\mathcal{P}_{\mathcal{X}}$ satisfies*

$$\alpha\text{-regret}_T(\mathcal{P}_{\mathcal{X}}) \leq \alpha \cdot \operatorname{regret}_T(\mathcal{P}_{\mathcal{Y}}).$$

Hence, OCO policies such as OGA, OMA, and FTRL combined with RAOCO yield $O(\sqrt{T})$ α-regret.

Comparison to Existing Work

Paper	Prob. Class	$(1-1/e)$-regret (Full Information)									Time
		Static			Dynamic			Optimistic			
		Uni.	Part.	Gen.	Uni.	Part.	Gen.	Uni.	Part.	Gen.	
Niazadeh et al. [2021]	GS	$r\sqrt{\log(n)T}$	×	×	×	×	×	×	×	×	$T^4 O_b$
Harvey et al. [2020]	GS	$\sqrt{r\log\left(\frac{n}{r}\right)T}$			×	×	×	×	×	×	$nr^2 + O_{\mathrm{m}} \cdot n^4/\epsilon^3 \cdot \log(n^3T/\epsilon)$
Matsuoka et al. [2021]	GS	$r\sqrt{r\log(nT)T}$	×	×	$\sqrt{r(r\log(nT)+P_T)T}$	×	×	×	×	×	nr
Streeter et al. [2009]	GS	$r^{\frac{3}{2}}\sqrt{\log(n)T}$		×	×	×	×	×	×	×	$n^2 c_{\mathrm{p}}$
Chen et al. [2018]	DR-S	$\sqrt{rnT}$			×	×	×	×	×	×	$\sqrt{T}O_{\mathrm{oco}} \cdot O_{\mathrm{m}} + nr^2$
Zhang et al. [2019]	DR-S	$T^{\frac{4}{5}}$			×	×	×	×	×	×	$T^{\frac{3}{5}}O_{\mathrm{oco}} \cdot O_{\mathrm{m}} + nr^2$
Zhang et al. [2022]	DR-S	$\sqrt{rnT}$			×	×	×	×	×	×	$O_{\mathrm{oco}} \cdot O_{\mathrm{m}} + nr^2$
Kakade et al. [2007]	LWD	$n(\alpha+2)\sqrt{T}$			×	×	×	×	×	×	TO_α
This work	WTP	$r\sqrt{\log\left(\frac{n}{r}\right)T}$			$\sqrt{r(r\log\left(\frac{n}{r}\right)+\log(n)P_T)T}$			$\sqrt{r(r\log\left(\frac{n}{r}\right)+\log(n)P_T)t} \cdot \sqrt{\sum_{t=1}^T \lVert\mathbf{g}_t - \mathbf{g}_t^\pi\rVert_\infty^2}$			nr^2

Case of Weighted Threshold Potentials

- The class of **WTP functions** satisfies the **sandwich property** under appropriate randomized rounding and concave relaxation schemes.

A **threshold potential** is defined as:

$$\Psi_{b,\mathbf{w},S}(\mathbf{x}) \triangleq \min\left\{b, \textstyle\sum_{j\in S} x_j w_j\right\}, \text{ for } \mathbf{x} \in \{0,1\}^n,$$

where $b \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ is a threshold, $S \subseteq V$, and $\mathbf{w} = (w_j)_{j\in S} \in [0,b]^{|S|}$ is a weight vector. The class of **weighted threshold potentials** (WTP) contains functions of the form:

$$f(\mathbf{x}) \triangleq \textstyle\sum_{\ell\in C} c_\ell \Psi_{b_\ell,\mathbf{w}_\ell,S_\ell}(\mathbf{x}), \quad \text{for } \mathbf{x} \in \{0,1\}^n,$$

where C is an arbitrary index set and $c_\ell \in \mathbb{R}_{\geq 0}$, for $\ell \in C$.

This class includes many important applications:

- **influence maximization** [Kempe et al., 2003]
- **facility location** [Krause and Golovin, 2014, Frieze, 1974]
- **cache networks** [Ioannidis and Yeh, 2016]
- **similarity caching** [Si Salem et al., 2022]
- **demand forecasting** [Ito and Fujimaki, 2016]
- **team formation** [Li et al., 2018]

Concave Relaxation

The relaxation of any WTP function f is itself: for $\mathbf{y} \in \mathcal{Y} \triangleq \operatorname{conv}(\mathcal{X})$:

$$\tilde{f}(\mathbf{y}) \triangleq f(\mathbf{y}) = \sum_{\ell\in C} c_\ell \min\left\{b_\ell, \sum_{j\in S_\ell} y_j w_{\ell,j}\right\}. \qquad (1)$$

The functional form of f remains unchanged while accommodating fractional inputs.

Relation to Sandwich Property

- **Concave relaxations** in (1) satisfy the **sandwich property** with $\alpha = 1 - 1/e$:

Consider the concave relaxations $\tilde{f}$ constructed from WTP functions $f \in \mathcal{F}$ via Eq. (1). Then, if these concave relaxations are uniformly Lipschitz with constant $L > 0$, the set $\mathcal{F}$ satisfies the sandwich property with $\alpha = 1 - 1/e$, for any **any negatively correlated rounding** $\Xi : \mathcal{Y} \to \mathcal{X}$.

Negatively-Correlated Rounding

- A **negatively-correlated randomized rounding** can always be constructed **if $\mathcal{X}$ is a matroid**:

(Chekuri et al. [2010, Theorem 1.1.]) Given a matroid $\mathcal{X} \subset \{0,1\}^n$, let $\mathbf{y} \in \operatorname{conv}(\mathcal{X})$ and Ξ be either **swap rounding** or **randomized pipage rounding**. Then, Ξ is **negatively correlated**.

RAOCO for WTPs

Theorem 2: *Let $\mathcal{X} \subseteq \{0,1\}^n$ be a matroid, and $\mathcal{F}$ be a subset of the WTP class for which the concave relaxations are uniformly Lipschitz continuous. Then, the RAOCO policy $\mathcal{P}_{\mathcal{X}}$ defined using swap rounding or randomized pipage rounding as Ξ, a sublinear OCO policy $\mathcal{P}_{\mathcal{Y}}$, and concave relaxations in (1) has sublinear α-regret, with $\alpha = 1 - 1/e$.*

Extensions

Dynamic Setting

In the **dynamic** setting, the decision maker compares its performance to the best sequence of decisions $(\mathbf{y}_t^\star)_{t\in[T]}$ with a path length regularity condition. We show that **RAOCO** equipped with a **modified OMA** policy has sublinear dynamic α-regret.

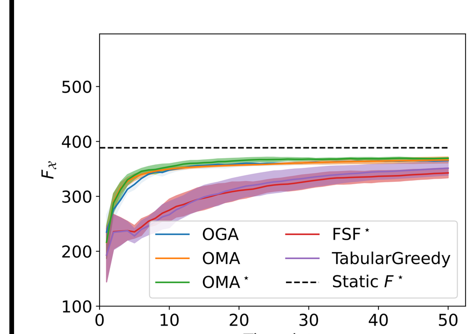

(a) **Stationary:** All policies have similar performance in the stationary setting and are able to match the performance of the static optimum.

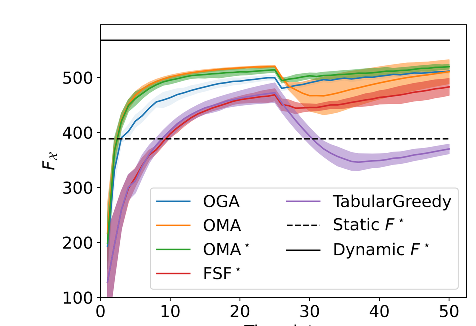

(b) **Non-Stationary:** In this setting only `OGA`, `OMA`, and `FSF*` show robustness

Optimistic Setting

In the **optimistic** setting, the decision maker has access to a predictor function $\pi_{t+1} : [0,1]^n \to \mathbb{R}_{\geq 0}$, before committing to a decision $\mathbf{x}_{t+1}$ at timeslot $t \in [T]$. We show that **RAOCO** equipped with a **modified Optimistic OMA** policy has sublinear α-regret.

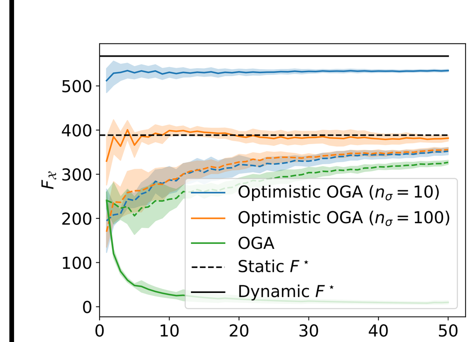

(a) **Optimistic Learning:** When the predictions are accurate `Optimistic OGA` is able to exploit these predictions to match the performance of the dynamic optimum.

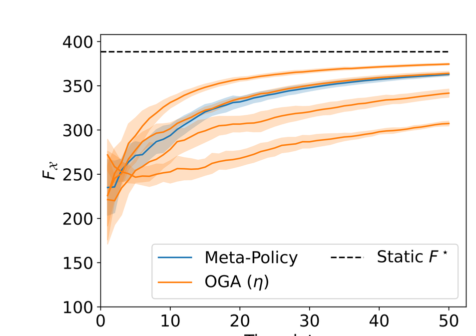

(b) **Meta Policies:** The meta-policy can learn the best configuration of `OGA` without tuning the learning rate.

Bandit Setting

In the **bandit** setting, only the reward value $f_t(\mathbf{x}_t)$ (but not $f_t(\cdot)$) is revealed. We construct a reduction to OCO based on Wan et al. [2023], for **general monotone submodular** rewards but restricted to **partition matroid constraints**, obtaining again a sublinear α-regret.

Paper	Prob. Class	$(1-1/e)$-regret (Bandit)					Time
		Static			Dynamic	Optimistic	
		Uni.	Part.	Gen.	Part.	Part.	
Kakade et al. [2007]	LWD	$n^{\frac{5}{3}}T^{\frac{2}{3}}$			×	×	TO_α
Niazadeh et al. [2021]	GS	$rn^{\frac{2}{3}}\log^{\frac{1}{3}}(n)T^{\frac{2}{3}}$	×	×	×	×	T^4O_b
Streeter et al. [2009]	GS	$\log^{\frac{1}{3}}(n)(nc_{\mathrm{p}}r)^{\frac{2}{3}}T^{\frac{2}{3}}$		×	×	×	n^2c_{p}
Wan et al. [2023]	GS	$r^{\frac{4}{3}}n^{\frac{4}{3}}T^{\frac{2}{3}}$		×	×	×	nr
This work	GS	$\alpha r^{4/3}n^{4/3}T^{2/3}$		×	$\sqrt{P_T}T^{4/5}$	$W\sqrt{P_T\sum_{l=1}^{T/W}\lVert\tilde{\mathbf{g}}_l - \mathbf{g}_l^\pi\rVert_2^2} + T(\delta + 1/W)$	nr

Numerical Results

Datasets & Constr.		$F^\star$	t	RAOCO-OGA $\bar{F}_{\mathcal{X}}/F^\star$	RAOCO-OGA std. dev.	RAOCO-OGA η	RAOCO-OMA $\bar{F}_{\mathcal{X}}/F^\star$	RAOCO-OMA std. dev.	RAOCO-OMA η	RAOCO-OMA γ	FSF* $\bar{F}_{\mathcal{X}}/F^\star$	FSF* std. dev.	FSF* η	FSF* γ	TabularGreedy $\bar{F}_{\mathcal{X}}/F^\star$	TabularGreedy std. dev.	TabularGreedy η	TabularGreedy c_{p}	Random $\bar{F}_{\mathcal{X}}/F^\star$	Random std. dev.
ZKC	uniform	0.234	33	0.902	1.85×10^{-2}	2.5	**0.965**	6.02×10^{-3}	10	0.05	0.839	5.02×10^{-3}	75	0.0	0.833	8.12×10^{-3}	160	1	0.642	3.03×10^{-2}
			66	0.924	1.60×10^{-2}		**0.967**	5.65×10^{-3}			0.896	3.83×10^{-3}			0.894	2.55×10^{-3}			0.624	2.80×10^{-2}
			99	0.945	7.70×10^{-3}		**0.982**	5.28×10^{-3}			0.933	4.17×10^{-3}			0.931	1.47×10^{-3}			0.622	1.91×10^{-2}
	partition	0.83	33	0.994	2.52×10^{-3}	8	**0.997**	6.95×10^{-4}	10	0.1	×				0.985	5.65×10^{-3}	10	1	0.953	4.88×10^{-3}
			66	0.99	8.75×10^{-4}		**0.994**	3.48×10^{-4}							0.987	2.93×10^{-3}			0.950	2.87×10^{-3}
			99	0.993	9.66×10^{-4}		**0.997**	3.37×10^{-4}							0.995	2.70×10^{-3}			0.953	1.77×10^{-3}
Epinions	uniform	0.171	50	0.845	1.61×10^{-2}	4	**0.853**	2.46×10^{-2}	10	0.01	0.703	6.03×10^{-2}	75	0.0	0.694	3.44×10^{-2}	160	1	0.632	2.99×10^{-2}
			100	0.865	1.13×10^{-2}		**0.906**	1.25×10^{-2}			0.776	1.28×10^{-2}			0.768	2.87×10^{-2}			0.615	2.32×10^{-2}
			149	0.88	7.97×10^{-3}		**0.925**	9.09×10^{-3}			0.807	1.52×10^{-2}			0.805	2.48×10^{-2}			0.629	2.36×10^{-2}
	partition	0.171	50	0.826	1.35×10^{-2}	3	**0.861**	1.82×10^{-2}	10	0.001	×				0.720	2.28×10^{-2}	160	1	0.620	6.08×10^{-3}
			100	0.854	6.73×10^{-3}		**0.908**	9.42×10^{-3}							0.786	1.27×10^{-2}			0.619	4.70×10^{-3}
			149	0.88	2.40×10^{-3}		**0.927**	6.34×10^{-3}							0.818	1.3×10^{-2}			0.625	1.10×10^{-2}
MovieLens	uniform	0.407	98	0.749	3.21×10^{-2}	0.5	**0.792**	2.60×10^{-2}	1.0	0.05	0.681	8.41×10^{-2}	1.0	0.001	0.69	1.11×10^{-2}	160	1	0.748	5.26×10^{-2}
			196	**0.786**	3.82×10^{-2}		0.781	1.36×10^{-2}			0.713	7.44×10^{-2}			0.676	1.03×10^{-2}			0.7	4.76×10^{-2}
			293	0.846	3.98×10^{-2}		**0.866**	1.02×10^{-2}			0.756	6.42×10^{-2}			0.769	1.65×10^{-2}			0.711	3.11×10^{-2}
	partition	0.419	98	0.889	5.20×10^{-2}	0.5	**0.948**	9.70×10^{-3}	10	0.001	×				0.908	2.39×10^{-2}	160	8	0.829	5.54×10^{-2}
			196	0.872	2.12×10^{-2}		**0.908**	5.14×10^{-3}							0.902	1.98×10^{-2}			0.814	1.51×10^{-2}
			293	0.926	2.34×10^{-2}		0.948	3.13×10^{-3}							**0.964**	2.01×10^{-2}			0.874	2.84×10^{-3}
SynthTF	uniform	200	33	0.984	7.05×10^{-3}	4	**0.987**	4.03×10^{-3}	0.05	0.1	0.845	4.73×10^{-2}	1	0	0.844	2.21×10^{-2}	1	2	0.605	3.13×10^{-2}
			66	0.994	3.87×10^{-3}		**0.994**	2.83×10^{-3}			0.868	2.14×10^{-2}			0.886	1.88×10^{-2}			0.603	2.01×10^{-2}
			99	0.995	2.01×10^{-3}		**0.998**	2.55×10^{-3}			0.869	2.95×10^{-2}			0.902	2.28×10^{-2}			0.612	1.18×10^{-2}
	partition	400	33	0.98	3.48×10^{-3}	3	**0.983**	2.33×10^{-3}	0.1	0.001	×				0.834	4.01×10^{-3}	1	1	0.611	9.15×10^{-3}
			66	0.990	1.75×10^{-3}		**0.991**	1.18×10^{-3}							0.852	2.04×10^{-3}			0.607	5.44×10^{-3}
			99	0.993	1.11×10^{-3}		**0.994**	7.75×10^{-4}							0.857	1.36×10^{-3}			0.601	7.71×10^{-3}

Acknowledgements

We gratefully acknowledge support from the NSF grants 1750539, 2112471, 1813406, and 1908510.

{sisalem, gozcan, ioannidis}@ece.neu.edu, {nikolaou, evimaria}@bu.edu