# Some New Complexity Results for Composite Optimization 

Guanghui (George) Lan

Georgia Institue of Technology
Joint work with Yuyuan Ouyang (Clemson) and Yi Zhou (Georgia Tech)
Workshop on Optimization without Borders, dedicated to
Yuri Nesterov's 60th Birthday, Les Houches, France,
December 7-12, 2016

## General CP methods

Problem: $\Psi^{*}=\min _{x \in X} \Psi(x)$.

- $X$ closed and convex.
- $\psi$ is convex

Goal: to find an $\epsilon$-solution, i.e., $\bar{x} \in X$ s.t. $\Psi(\bar{x})-\Psi^{*} \leq \epsilon$.
Complexity: the number of (sub)gradient evaluations of $\psi-$

- $\psi$ is smooth: $\mathcal{O}(1 / \sqrt{\epsilon})$.
- $\psi$ is nonsmooth: $\mathcal{O}\left(1 / \epsilon^{2}\right)$.
- $\psi$ is strongly convex: $\mathcal{O}(\log (1 / \epsilon))$.


## Composite optimization problems

We consider composite problems which can be modeled as

$$
\Psi^{*}=\min _{x \in X}\{\Psi(x):=f(x)+h(x)\} .
$$

Here, $f: X \rightarrow \mathbb{R}$ is a smooth and expensive term (data fitting), $h: X \rightarrow \mathbb{R}$ is a nonsmooth regularization term (solution structures), and $X$ is a closed convex set.

## Three Challenging Cases

- $h$ or $X$ are not necessarily simple.
- $f$ given by the summation of many terms.
- $f($ or $h$ ) is possibly nonconvex.


## Existing complexity results

Problem: $\Psi^{*}:=\min _{x \in X}\{\Psi(x):=f(x)+h(x)\}$.
First-order methods: iterative methods which operate with the gradients (subgradients) of $f$ and $h$.

Complexity: number of iterations needed to find an $\epsilon$-solution, i.e., a point $\bar{x} \in X$ s.t. $\Psi(\bar{x})-\Psi^{*} \leq \epsilon$.

[^0]Tseng 08, Beck and Teboulle 09)

## Existing complexity results

Problem: $\Psi^{*}:=\min _{x \in X}\{\Psi(x):=f(x)+h(x)\}$.
First-order methods: iterative methods which operate with the gradients (subgradients) of $f$ and $h$.

Complexity: number of iterations needed to find an $\epsilon$-solution, i.e., a point $\bar{x} \in X$ s.t. $\Psi(\bar{x})-\Psi^{*} \leq \epsilon$.

## Easy case: $h$ simple, $X$ simple

$\operatorname{Pr}_{X, h}(y):=\operatorname{argmin}_{x \in X}\|y-x\|^{2}+h(x)$ is easy to compute (e.g., compressed sensing). Complexity: $\mathcal{O}(1 / \sqrt{\epsilon})$ (Nesterov 07, Tseng 08, Beck and Teboulle 09).

## More difficult cases

## $h$ general, $X$ simple

$h$ is a general nonsmooth function; $P_{X}:=\operatorname{argmin}_{x \in X}\|y-x\|^{2}$ is easy to compute. Complexity: $\mathcal{O}\left(1 / \epsilon^{2}\right)$.

## More difficult cases

## $h$ general, $X$ simple

$h$ is a general nonsmooth function; $P_{X}:=\operatorname{argmin}_{x \in X}\|y-x\|^{2}$ is easy to compute. Complexity: $\mathcal{O}\left(1 / \epsilon^{2}\right)$.

## $h$ structured, $X$ simple

$h$ is structured, e.g., $h(x)=\max _{y \in Y}\langle A x, y\rangle ; P_{X}$ is easy to compute. Complexity: $\mathcal{O}(1 / \epsilon)$.

> is easy to compute (e.g.
matrix completion).Complexity

## More difficult cases

## $h$ general, $X$ simple

$h$ is a general nonsmooth function; $P_{X}:=\operatorname{argmin}_{x \in X}\|y-x\|^{2}$ is easy to compute. Complexity: $\mathcal{O}\left(1 / \epsilon^{2}\right)$.

## $h$ structured, $X$ simple

$h$ is structured, e.g., $h(x)=\max _{y \in Y}\langle A x, y\rangle ; P_{X}$ is easy to compute. Complexity: $\mathcal{O}(1 / \epsilon)$.

## $h$ simple, $X$ complicated

$L_{X, h}(y):=\operatorname{argmin}_{x \in X}\langle y, x\rangle+h(x)$ is easy to compute (e.g., matrix completion).Complexity: $\mathcal{O}(1 / \epsilon)$.

## Motivation

$h$ simple, $X$ simple
$\mathcal{O}(1 / \sqrt{\epsilon}) \quad 100$
$h$ general, $X$ simple
$\mathcal{O}\left(1 / \epsilon^{2}\right)$
$10^{8}$
$h$ structured, $X$ simple
$10^{4}$
$h$ simple, $X$ complicated

More general $h$ or more complicated $X$

Slow convergence of first-order algorithms
v
A large number of gradient evaluations of $\nabla f$

## Motivation

| $h$ simple, $X$ simple | $\mathcal{O}(1 / \sqrt{\epsilon})$ | 100 | () |
| :---: | :---: | :---: | :---: |
| $h$ general, $X$ simple | $\mathcal{O}\left(1 / \epsilon^{2}\right)$ | $10^{8}$ | - |
| $h$ structured, $X$ simple | $\mathcal{O}(1 / \epsilon)$ | $10^{4}$ | (2) |
| $h$ simple, $X$ complicated | $\mathcal{O}(1 / \epsilon)$ | $10^{4}$ |  |

More general $h$ or more complicated $X$


Slow convergence of first-order algorithms

A large number of gradient evaluations of $\nabla f$
Question: Can we skip the computation of $\nabla f$ ?

## Composite problems

$\Psi^{*}=\min _{x \in X}\{\Psi(x):=f(x)+h(x)\}$.

- $f$ is smooth, i.e., $\exists L>0$ s.t. $\forall x, y \in X$, $\|\nabla f(y)-\nabla f(x)\| \leq L\|y-x\|$.
- $h$ is nonsmooth, i.e., $\exists M>0$ s.t. $\forall x, y \in X$, $|h(x)-h(y)| \leq M\|y-x\|$.
- $P_{X}$ is simple to compute.


## Question:

How many number of gradient evaluations of $\nabla f$ and subgradient evaluations of $h^{\prime}$ are needed to find an $\epsilon$-solution?

## Existing results

Existing algorithms evaluate $\nabla f$ and $h^{\prime}$ together at each iteration:

- Mirror-prox method (Juditsky, Nemirovski and Travel, 11):

$$
\mathcal{O}\left\{\frac{L}{\epsilon}+\frac{M^{2}}{\epsilon^{2}}\right\}
$$

- Accelerated stochastic approximation (Lan, 12):

$$
\mathcal{O}\left\{\sqrt{\frac{L}{\epsilon}}+\frac{M^{2}}{\epsilon^{2}}\right\}
$$

## Issue:

Whenever the second term dominates, the number of gradient evaluations $\nabla f$ is given by $\mathcal{O}\left(1 / \epsilon^{2}\right)$.

## Bottleneck for composite problems

- The computation of $\nabla f$, however, is often the bottleneck in comparison with that of $h^{\prime}$.
- The computation of $\nabla f$ invovles a large data set, while that of $h^{\prime}$ only involves a very sparse matrix (e.g., total variation minimization).
- Can we reduce the number of gradient evaluations for $\nabla f$ from $\mathcal{O}\left(1 / \epsilon^{2}\right)$ to $\mathcal{O}(1 / \sqrt{\epsilon})$, while still maintaining the optimal $\mathcal{O}\left(1 / \epsilon^{2}\right)$ bound on subgradient evaluations for $h^{\prime}$ ?


## The gradient sliding algorithm

Algorithm 1 The gradient sliding (GS) algorithm
Input: Initial point $x_{0} \in X$ and iteration limit $N$.
Let $\beta_{k} \geq 0, \gamma_{k} \geq 0$, and $T_{k} \geq 0$ be given and set $\bar{x}_{0}=x_{0}$. for $k=1,2, \ldots, N$ do

1. Set $\underline{x}_{k}=\left(1-\gamma_{k}\right) \bar{x}_{k-1}+\gamma_{k} x_{k-1}$ and $g_{k}=\nabla f\left(\underline{x}_{k}\right)$.
2. Set $\left(x_{k}, \tilde{x}_{k}\right)=\operatorname{PS}\left(g_{k}, x_{k-1}, \beta_{k}, T_{k}\right)$.
3. Set $\bar{x}_{k}=\left(1-\gamma_{k}\right) \bar{x}_{k-1}+\gamma_{k} \tilde{x}_{k}$.
end for
Output: $\bar{x}_{N}$.

PS: the prox-sliding procedure.

## The PS procedure

## Procedure $\left(x^{+}, \tilde{x}^{+}\right)=\operatorname{PS}(g, x, \beta, T)$

Let the parameters $p_{t}>0$ and $\theta_{t} \in[0,1], t=1, \ldots$, be given.
Set $u_{0}=\tilde{u}_{0}=x$.
for $t=1,2, \ldots, T$ do
$u_{t}=\operatorname{argmin}_{u \in X}\left\langle g+h^{\prime}\left(u_{t-1}\right), u\right\rangle+\frac{\beta}{2}\|u-x\|^{2}+\frac{\beta p_{t}}{2}\left\|u-u_{t-1}\right\|^{2}$,
$\tilde{u}_{t}=\left(1-\theta_{t}\right) \tilde{u}_{t-1}+\theta_{t} u_{t}$.
end for
Set $x^{+}=u_{T}$ and $\tilde{x}^{+}=\tilde{u}_{T}$.

Note: $\|\cdot-\cdot\|^{2} / 2$ can be replaced by the more general Bregman distance $V(x, u)=\omega(u)-\omega(x)-\langle\nabla \omega(x), u-x\rangle$.

## Remarks

When supplied with $g(\cdot), x \in X, \beta$, and $T$, the PS procedure computes a pair of approximate solutions $\left(x^{+}, \tilde{x}^{+}\right) \in X \times X$ for the problem of:

$$
\operatorname{argmin}_{u \in X}\left\{\Phi(u):=\langle g, u\rangle+h(u)+\frac{\beta}{2}\|u-x\|^{2}\right\} .
$$

In each iteration, the subproblem is given by

$$
\operatorname{argmin}_{u \in X}\left\{\Phi_{k}(u):=\left\langle\nabla f\left(\underline{x}_{k}\right), u\right\rangle+h(u)+\frac{\beta_{k}}{2}\left\|u-x_{k}\right\|^{2}\right\} .
$$

## Convergence of the PS proedure

## Proposition

If $\left\{p_{t}\right\}$ and $\left\{\theta_{t}\right\}$ in the PS procedure satisfy

$$
p_{t}=\frac{t}{2} \quad \text { and } \quad \theta_{t}=\frac{2(t+1)}{t(t+3)}
$$

then for any $t \geq 1$ and $u \in X$,

$$
\Phi\left(\tilde{u}_{t}\right)-\Phi(u)+\frac{\beta(t+1)(t+2)}{2 t(t+3)}\left\|u_{t}-u\right\|^{2} \leq \frac{M^{2}}{\beta(t+3)}+\frac{\beta\left\|u_{0}-u\right\|^{2}}{t(t+3)} .
$$

## Convergence of the GS algorithm

## Theorem

Suppose that the previous conditions on $\left\{p_{t}\right\}$ and $\left\{\theta_{t}\right\}$ hold, and that $N$ is given a priori. If

$$
\beta_{k}=\frac{2 L}{k}, \quad \gamma_{k}=\frac{2}{k+1}, \quad \text { and } T_{k}=\left\lceil\frac{M^{2} N k^{2}}{\tilde{D} L^{2}}\right\rceil
$$

for some $\tilde{D}>0$, then

$$
\Psi\left(\bar{x}_{N}\right)-\Psi\left(x^{*}\right) \leq \frac{L}{N(N+1)}\left(\frac{3\left\|x_{0}-x^{*}\right\|^{2}}{2}+2 \tilde{D}\right) .
$$

Remark: Do NOT need $N$ given a priori if $X$ is bounded.

## Complexity of the GS algorithm

Denote $D_{X}:=\max _{x_{1}, x_{2} \in X}\left\|x_{1}-x_{2}\right\|$ and set $\tilde{D}=3 D_{X}^{2} / 4$.
The number of gradient evaluations of $\nabla f$ is bounded by

$$
\sqrt{\frac{3 L D_{X}^{2}}{\epsilon}}
$$

and the number of subgradient evaluations of $h^{\prime}$ is given by
$\sum_{k=1}^{N} T_{k}$, which is bounded by

$$
\frac{4 M^{2} D_{X}^{2}}{\epsilon^{2}}+\sqrt{\frac{3 L D_{X}^{2}}{\epsilon}} .
$$

## Consequence

Significantly reduce the number of gradient evaluations of $\nabla f$ from $\mathcal{O}\left(1 / \epsilon^{2}\right)$ to $\mathcal{O}(1 / \sqrt{\epsilon})$, even though the whole objective function $\psi$ is nonsmooth in general.

## Extensions

- Gradient sliding for $\min _{x \in X} f(x)+h(x)$ :

|  | total iter. | $\nabla f$ |
| :--- | :--- | :--- |
| $h$ general nonsmooth | $\mathcal{O}\left(1 / \epsilon^{2}\right)$ | $\mathcal{O}(1 / \sqrt{\epsilon})$ |
| $h$ structured nonsmooth | $\mathcal{O}(1 / \epsilon)$ | $\mathcal{O}(1 / \sqrt{\epsilon})$ |
| $f$ strongly convex | $\mathcal{O}(1 / \epsilon)$ | $\mathcal{O}(\log (1 / \epsilon))$ |

- Conditional gradient sliding methods for problems with more complicated feasible set.

|  | total iter. (LO oracle) | $\nabla f$ |
| :--- | :--- | :--- |
| $f$ convex | $\mathcal{O}(1 / \epsilon)$ | $\mathcal{O}(1 / \sqrt{\epsilon})$ |
| $f$ strongly convex | $\mathcal{O}(1 / \epsilon)$ | $\mathcal{O}(\log (1 / \epsilon))$ |

## The problem of interest

Problem: $\psi^{*}:=\min _{x \in X}\left\{\Psi(x):=\sum_{i=1}^{m} f_{i}(x)+h(x)+\mu \omega(x)\right\}$.

- $X$ closed and convex.
- $f_{i}$ smooth convex: $\left\|\nabla f_{i}\left(x_{1}\right)-\nabla f_{i}\left(x_{2}\right)\right\|_{*} \leq L_{i}\left\|x_{1}-x_{2}\right\|$.
- $h$ simple, e.g., $I_{1}$ norm.
- $\omega$ strongly convex with modulus 1 w.r.t. an arbitrary norm.
- $\mu \geq 0$.
- Subproblem argmin ${ }_{x \in X}\langle g, x\rangle+h(x)+\mu \omega(x)$ is easy.
- Denote $f(x) \equiv \sum_{i=1}^{m} f_{i}(x)$ and $L \equiv \sum_{i=1}^{m} L_{i}$. $f$ is smooth with Lipschitz constant $\leq$ L.


## Stochastic subgradient descent for nonsmooth problems

- General stochastic programming (SP): $\min _{x \in X} \mathbb{E}_{\xi}[F(x, \xi)]$.
- Reformulation of the finite sum problem as SP:

$$
\begin{aligned}
& \text { - } \xi \in\{1, \ldots, m\}, \operatorname{Prob}\{\xi=i\}=\nu_{i} \text {, and } \\
& \quad F(x, i)=\nu_{i}^{-1} f_{i}(x)+h(x)+\mu \omega(x), i=1, \ldots, m .
\end{aligned}
$$

- Iteration complexity: $\mathcal{O}\left(1 / \epsilon^{2}\right)$ or $\mathcal{O}(1 / \epsilon)(\mu>0)$.
- Iteration cost: $m$ times cheaper than deterministic first-order methods.
- Save up to a factor of $\mathcal{O}(m)$ subgradient computations.
- For details, see Nemirovski et. al. (09).


## Required $\nabla$ f's in the smooth case

For simplicity, focus on the strongly convex case ( $\mu>0$ ). Goal: find a solution $\bar{x} \in X$ s.t. $\left\|\bar{x}-x^{*}\right\| \leq \epsilon\left\|x^{0}-x^{*}\right\|$.

- Nesterov's optimal method (Nesterov 83):

$$
\mathcal{O}\left\{m \sqrt{\frac{L_{f}}{\mu}} \log \frac{1}{\epsilon}\right\}
$$

- Accelerated stochastic approximation (Lan 12, Ghadimi and Lan 13):

$$
\mathcal{O}\left\{\sqrt{\frac{L_{f}}{\mu}} \log \frac{1}{\epsilon}+\frac{\sigma^{2}}{\mu \epsilon}\right\}
$$

Note: the optimality of the latter bound for general SP does not preclude more efficient algorithms for the finite-sum problem.

## Randomized incremental gradient methods

Each iteration requires a randomly selected $\nabla f_{i}(x)$.

- Stochastic average gradient (SAG) by Schmidt, Roux and Bach 13:

$$
\mathcal{O}\left((m+L / \mu) \log \frac{1}{\epsilon}\right) .
$$

- Similar results were obtained in Johnson and Zhang 13, Defazio et al. 14...
- Worse dependence on the $L / \mu$ than Nesterov's method.


## Coordinate ascent in the dual

$\min \left\{\sum_{i=1}^{m} \phi_{i}\left(a_{i}^{T} x\right)+h(x)\right\}, h$ strongly convex w.r.t. I I norm.
All these coordinate algorithms achieve $\mathcal{O}\left\{m+\sqrt{\frac{m L}{\mu}} \log \frac{1}{\epsilon}\right\}$.

- Shalev-Shwartz and Zhang 13, 15 (restarting stochastic dual ascent),
- Lin, Lu and Xiao, 14 ( Nesterov, Fercoq and P. Richtárik's), see also Zhang and Xiao 14 (Chambolle and Pock),
- Dang and Lan 14 (non-strongly convex), $\mathcal{O}(1 / \epsilon)$ or $\mathcal{O}(1 / \sqrt{\epsilon})$.


## Some issues:

- Deal with a more special class of problems.
- Require argmin $\left\{\langle g, y\rangle+\phi_{i}^{*}(y)+\|y\|_{*}^{2}\right\}$, not incremental gradient methods.


## Open problems and our research

## Problems:

- Can we accelerate the convergence of randomized incremental gradient method?
- What is the best possible performance we can expect?

Our contributions:

- A primal-dual gradient (PDG) method = a primal-dual look to Nesterov's method.
- A randomized PDG (RPDG).
- A new lower complexity bound.
- A game-theoretic interpretation for acceleration.

Catalyst: Lin, Mairal, and Harchaoui 15.

## Reformulation and game/economic interpretation

Let $J_{f}$ be the conjugate function of $f$. Consider
$\Psi^{*}:=\min _{x \in X}\left\{h(x)+\mu \omega(x)+\max _{g \in \mathcal{G}}\langle x, g\rangle-J_{f}(g)\right\}$

- The buyer purchases products from the supplier.
- The unit price is given by $g \in \mathbb{R}^{n}$.
- $X, h$ and $\omega$ are constraints and other local cost for the buyer.
- The profit of supplier: revenue $(\langle x, g\rangle)$ - local cost $J_{f}(g)$.


## How to achieve equilibrium?

Current order quantity $x^{0}$, and product price $g^{0}$.
Proximity control functions:

$$
\begin{aligned}
& P\left(x^{0}, x\right):=\omega(x)-\left[\omega\left(x^{0}\right)+\left\langle\omega^{\prime}\left(x^{0}\right), x-x^{0}\right\rangle\right] . \\
& D_{f}\left(g_{i}^{0}, y_{i}\right):=J_{f}(g)-\left[J_{f}\left(g^{0}\right)+\left\langle J_{f}^{\prime}\left(g^{0}\right), g-g^{0}\right\rangle\right] .
\end{aligned}
$$

Dual prox-mapping:
$\mathcal{M}_{\mathcal{G}}\left(-\tilde{x}, g^{0}, \tau\right):=\arg \min _{g \in \mathcal{G}}\left\{\langle-\tilde{x}, g\rangle+J_{f}(g)+\tau D_{f}\left(g^{0}, g\right)\right\}$.
$\tilde{x}$ is the given or predicted demand. Maximize the profit, but not too far away from $g^{0}$.

Primal prox-mapping:
$\mathcal{M}_{x}\left(g, x^{0}, \eta\right):=\arg \min _{x \in X}\left\{\langle g, x\rangle+h(x)+\mu \omega(x)+\eta P\left(x^{0}, x\right)\right\}$.
$g$ is the given or predicted price. Minimize the cost, but not too far way from $x^{0}$.

## The deterministic PDG

## Algorithm 2 The primal-dual gradient method

Let $x^{0}=x^{-1} \in X$, and the nonnegative parameters $\left\{\tau_{t}\right\},\left\{\eta_{t}\right\}$, and $\left\{\alpha_{t}\right\}$ be given.
Set $g^{0}=\nabla f\left(x^{0}\right)$. for $t=1, \ldots, k$ do

Update $z^{t}=\left(x^{t}, y^{t}\right)$ according to $\tilde{x}^{t}=\alpha_{t}\left(x^{t-1}-x^{t-2}\right)+x^{t-1}$. $g^{t}=\mathcal{M}_{\mathcal{G}}\left(-\tilde{x}^{t}, g^{t-1}, \tau_{t}\right)$. $x^{t}=\mathcal{M}_{x}\left(g^{t}, x^{t-1}, \eta_{t}\right)$.
end for

## A game/economic interpretation

- The supplier predicts the buyer's demand based on historical information: $\tilde{x}^{t}=\alpha_{t}\left(x^{t-1}-x^{t-2}\right)+x^{t-1}$.
- The supplier seeks to maximize predicted profit, but not too far away from $g^{t-1}: g^{t}=\mathcal{M}_{\mathcal{G}}\left(-\tilde{x}^{t}, g^{t-1}, \tau_{t}\right)$.
- The buyer tries to minimize the cost, but not too far away from $x^{t-1}: x^{t}=\mathcal{M}_{X}\left(g^{t}, x^{t-1}, \eta_{t}\right)$.


## PDG in gradient form

Algorithm 3 PDG method in gradient form
Input: Let $x^{0}=x^{-1} \in X$, and the nonnegative parameters $\left\{\tau_{t}\right\},\left\{\eta_{t}\right\}$, and $\left\{\alpha_{t}\right\}$ be given.
Set $x^{0}=x^{0}$.
for $t=1,2, \ldots, k$ do
$\tilde{x}^{t}=\alpha_{t}\left(x^{t-1}-x^{t-2}\right)+x^{t-1}$.
$\underline{x}^{t}=\left(\tilde{x}^{t}+\tau_{t} \underline{x}^{t-1}\right) /\left(1+\tau_{t}\right)$.
$g^{t}=\nabla f\left(\underline{x}^{t}\right)$.
$x^{t}=\mathcal{M}_{X}\left(g^{t}, x^{t-1}, \eta_{t}\right)$.
end for

Idea: set $J_{f}^{\prime}\left(g^{t-1}\right)=\underline{x}^{t-1}$.

## Relation to Nesterov's method

A variant of Nesterov's method:

$$
\begin{aligned}
\underline{x}^{t} & =\left(1-\theta_{t}\right) \bar{x}^{t-1}+\theta_{t} x^{t-1} \\
x^{t} & =M_{X}\left(\sum_{i=1}^{m} \nabla f_{i}\left(\underline{x}^{t}\right), x^{t-1}, \eta_{t}\right) \\
\bar{x}^{t} & =\left(1-\theta_{t}\right) \bar{x}^{t-1}+\theta_{t} x^{t}
\end{aligned}
$$

Note that

$$
\underline{x}^{t}=\left(1-\theta_{t}\right) \underline{x}^{t-1}+\left(1-\theta_{t}\right) \theta_{t-1}\left(x^{t-1}-x^{t-2}\right)+\theta_{t} x^{t-1} .
$$

Equivalent to PDG with $\tau_{t}=\left(1-\theta_{t}\right) / \theta_{t}$ and $\alpha_{t}=\theta_{t-1}\left(1-\theta_{t}\right) / \theta_{t}$.
Nesterov's acceleration: looking-ahead dual players.
Gradient descent: myopic dual players ( $\alpha_{t}=\tau_{t}=0$ in PDG).

## Convergence of PDG (or Nesterov's variant)

## Theorem

Define $\bar{x}^{k}:=\left(\sum_{t=1}^{k} \theta_{t}\right)^{-1} \sum_{t=1}^{k}\left(\theta_{t} x^{t}\right)$. Suppose that
$\tau_{t}=\sqrt{\frac{2 L_{t}}{\mu}}, \quad \eta_{t}=\sqrt{2 L_{f} \mu}, \quad \alpha_{t}=\alpha \equiv \frac{\sqrt{2 L_{t} / \mu}}{1+\sqrt{2 L_{f} / \mu}}, \quad$ and $\quad \theta_{t}=\frac{1}{\alpha^{t}}$.
Then,

$$
\begin{array}{ll}
P\left(x^{k}, x^{*}\right) & \leq \frac{\mu+L_{f}}{\mu} \alpha^{k} P\left(x^{0}, x^{*}\right) \\
\Psi\left(\bar{x}^{k}\right)-\Psi\left(x^{*}\right) & \leq \mu(1-\alpha)^{-1}\left[1+\frac{L_{f}}{\mu}\left(2+\frac{L_{f}}{\mu}\right)\right] \alpha^{k} P\left(x^{0}, x^{*}\right) .
\end{array}
$$

## Theorem

$$
\begin{aligned}
& \text { If } \tau_{t}=\frac{t-1}{2}, \eta_{t}=\frac{4 L_{t}}{t}, \alpha_{t}=\frac{t-1}{t} \text {, and } \theta_{t}=t \text {, then } \\
& \Psi\left(\bar{x}^{k}\right)-\Psi\left(x^{*}\right) \leq \frac{8 L_{t}}{k(k+1)} P\left(x^{0}, x^{*}\right) .
\end{aligned}
$$

## A multi-dual-player reformulation

- Let $J_{i}: \mathcal{Y}_{i} \rightarrow \mathbb{R}$ be the conjugate functions of $f_{i}$ and $\mathcal{Y}_{i}$, $i=1, \ldots, m$, denote the dual spaces. $\min _{x \in X}\left\{h(x)+\mu \omega(x)+\max _{y_{i} \in \mathcal{Y}_{i}}\left\langle x, \sum_{i} y_{i}\right\rangle-\sum_{i} J_{i}(y)\right\}$,
- Define their new dual prox-functions and dual prox-mappings as

$$
\begin{array}{cll}
D_{i}\left(y_{i}^{0}, y_{i}\right) & :=\quad J_{i}\left(y_{i}\right)-\left[J_{i}\left(y_{i}^{0}\right)+\left\langle J_{i}^{\prime}\left(y_{i}^{0}\right), y_{i}-y_{i}^{0}\right\rangle\right] \\
\mathcal{M}_{\mathcal{Y}_{i}}\left(-\tilde{x}, y_{i}^{0}, \tau\right) & := & \arg \min _{y_{i} \in \mathcal{Y}_{i}}\left\{\langle-\tilde{x}, y\rangle+J_{i}\left(y_{i}\right)+\tau D_{i}\left(y_{i}^{0}, y_{i}\right)\right\} .
\end{array}
$$

## The RPDG method

## Algorithm 4 The RPDG method

Let $x^{0}=x^{-1} \in X$, and $\left\{\tau_{t}\right\},\left\{\eta_{t}\right\}$, and $\left\{\alpha_{t}\right\}$ be given.
Set $y_{i}^{0}=\nabla f_{i}\left(x^{0}\right), i=1, \ldots, m$. for $t=1, \ldots, k$ do

Choose $i_{t}$ according to $\operatorname{Prob}\left\{i_{t}=i\right\}=p_{i}, i=1, \ldots, m$.

$$
\begin{aligned}
\tilde{x}^{t} & =\alpha_{t}\left(x^{t-1}-x^{t-2}\right)+x^{t-1} . \\
y_{i}^{t} & = \begin{cases}\mathcal{M}_{y_{i}}\left(-\tilde{x}^{t}, y_{i}^{t-1}, \tau_{t}\right), & i=i_{t}, \\
y_{i}^{t-1}, & i \neq i_{t} .\end{cases} \\
\tilde{y}_{i}^{t} & = \begin{cases}p_{i}^{-1}\left(y_{i}^{t}-y_{i}^{t-1}\right)+y_{i}^{t-1}, & i=i_{t}, \\
y_{i}^{t-1}, & i \neq i_{t} .\end{cases} \\
x^{t} & =\mathcal{M}_{X}\left(\sum_{i=1}^{m} \tilde{y}_{i}^{t}, x^{t-1}, \eta_{t}\right) .
\end{aligned}
$$

end for

## RPDG in gradient form

## Algorithm 5 RPDG

for $t=1, \ldots, k$ do
Choose $i_{t}$ according to $\operatorname{Prob}\left\{i_{t}=i\right\}=p_{i}, i=1, \ldots, m$.

$$
\begin{aligned}
\tilde{x}^{t} & =\alpha_{t}\left(x^{t-1}-x^{t-2}\right)+x^{t-1} . \\
\underline{x}_{i}^{t} & = \begin{cases}\left(1+\tau_{t}\right)^{-1}\left(\tilde{x}^{t}+\tau_{t} \underline{x}_{i}^{t-1}\right), & i=i_{t}, \\
\underline{x}_{i}^{t-1}, & i \neq i_{t} .\end{cases} \\
y_{i}^{t} & = \begin{cases}\nabla f_{i}\left(x_{i}^{t}\right), & i=i_{t}, \\
y_{i}^{t-1}, & i \neq i_{t} .\end{cases} \\
x^{t} & =\mathcal{M}_{x}\left(g^{t-1}+\left(p_{i t}^{-1}-1\right)\left(y_{i_{t}}^{t}-y_{i_{t}}^{t-1}\right), x^{t-1}, \eta_{t}\right) \\
g^{t} & =g^{t-1}+y_{i_{t}}^{t}-y_{i_{t}}^{t-1} .
\end{aligned}
$$

## end for

## Game-theoretic interpretation for RPDG

- The suppliers predict the buyer's demand as before.
- Only one randomly selcted supplier will change his/her price, arriving at $y^{t}$.
- The buyer would have used $y^{t}$ as the price, but the algorithm converges slowly (a worse depedence on $m$ ) (Dang and Lan 14).
- Add a dual prediction (estimation) step, i.e., $\tilde{y}^{t}$ s.t. $\mathbb{E}_{t}\left[\tilde{y}_{i}^{t}\right]=\hat{y}_{i}^{t}$, where $\hat{y}_{i}^{t}:=\mathcal{M}_{\mathcal{Y}_{i}}\left(-\tilde{x}^{t}, y_{i}^{t-1}, \tau_{i}^{t}\right)$.
- The buyer uses $\tilde{y}^{t}$ to determine the order quantity.


## Rate of Convergence

## Theorem

$$
\begin{aligned}
\text { Let } C & =\frac{8 L}{\mu} \cdot \text { and } \\
p_{i} & =\operatorname{Prob}\left\{i_{t}=i\right\}=\frac{1}{2 m}+\frac{L_{i}}{2 L}, i=1, \ldots, m \\
\tau_{t} & =\frac{\sqrt{(m-1)^{2}+4 m C}-(m-1)}{2 m} \\
\eta_{t} & =\frac{\mu \sqrt{(m-1)^{2}+4 m C}+\mu(m-1)}{2} \\
\alpha_{t} & =\alpha:=1-\frac{1}{(m+1)+\sqrt{(m-1)^{2}+4 m C}}
\end{aligned}
$$

## Then

$\mathbb{E}\left[P\left(x^{k}, x^{*}\right)\right] \leq\left(1+\frac{3 L_{f}}{\mu}\right) \alpha^{k} P\left(x^{0}, x^{*}\right)$,
$\mathbb{E}\left[\Psi\left(\bar{x}^{k}\right)\right]-\Psi^{*} \leq \alpha^{k / 2}(1-\alpha)^{-1}\left[\mu+2 L_{f}+\frac{L_{f}^{2}}{\mu}\right] P\left(x^{0}, x^{*}\right)$.

## The iteration complexity of RPGD

- To find a point $\bar{x} \in X$ s.t. $\mathbb{E}\left[P\left(\bar{x}, x^{*}\right)\right] \leq \epsilon$ :
$\mathcal{O}\left\{\left(m+\sqrt{\frac{m L}{\mu}}\right) \log \left[\frac{P\left(x^{0}, x^{*}\right)}{\epsilon}\right]\right\}$.
- To find a point $\bar{x} \in X$ s.t. $\operatorname{Prob}\left\{P\left(\bar{x}, x^{*}\right) \leq \epsilon\right\} \geq 1-\lambda$ for some $\lambda \in(0,1)$ :
$\mathcal{O}\left\{\left(m+\sqrt{\frac{m L}{\mu}}\right) \log \left[\frac{P\left(x^{0}, x^{*}\right)}{\lambda \epsilon}\right]\right\}$.
- Similar results hold for the ergodic sequence in terms of function values.
- A factor of up to $\mathcal{O}\left\{\min \left\{\sqrt{\frac{L}{\mu}}, \sqrt{m}\right\}\right\}$ savings on gradient computation (or price changes), at the price of more order transactions.


## Lower complexity bound

$\min _{x_{i} \in \mathbb{R}^{n}, i=1, \ldots, m}\left\{\Psi(x):=\sum_{i=1}^{m}\left[f_{i}\left(x_{i}\right)+\frac{\mu}{2}\left\|x_{i}\right\|_{2}^{2}\right]\right\}$.
$f_{i}\left(x_{i}\right)=\frac{\mu(\mathcal{Q}-1)}{4}\left[\frac{1}{2}\left\langle A x_{i}, x_{i}\right\rangle-\left\langle e_{1}, x_{i}\right\rangle\right] . \tilde{n} \equiv n / m$,
$A=\left(\begin{array}{cccccccc}2 & -1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ -1 & 2 & -1 & 0 & \cdots & 0 & 0 & 0 \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \\ 0 & 0 & 0 & 0 & \cdots & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & \cdots & 0 & -1 & \kappa\end{array}\right), \kappa=\frac{\sqrt{Q}+3}{\sqrt{Q}+1}$.

## Theorem

Denote $q:=(\sqrt{\mathcal{Q}}-1) /(\sqrt{\mathcal{Q}}+1)$. Then the iterates $\left\{x^{k}\right\}$ generated by a randomized incremental gradient method must satisfy $\frac{\mathbb{E}\left[\left\|x^{k}-x^{*}\right\|_{2}^{2}\right]}{\left\|x^{0}-x^{*}\right\|_{2}^{2}} \geq \frac{1}{2} \exp \left(-\frac{4 k \sqrt{\mathcal{Q}}}{m(\sqrt{\mathcal{Q}}+1)^{2}-4 \sqrt{\mathcal{Q}}}\right)$ for any $n \geq \underline{n}(m, k) \equiv\left[m \log \left[\left(1-\left(1-q^{2}\right) / m\right)^{k} / 2\right]\right] /(2 \log q)$.

## Complexity

## Corollary

The number of gradient evaluations performed by a randomized incremental gradient method for finding a solution $\bar{x} \in X$ s.t. $\mathbb{E}\left[\left\|\bar{x}-x^{*}\right\|_{2}^{2}\right] \leq \epsilon$ cannot be smaller than
$\Omega\left\{(\sqrt{m \mathcal{C}}+m) \log \frac{\left\|x^{0}-x^{*}\right\|_{2}^{2}}{\epsilon}\right\}$ if $n$ is sufficiently large.

## Other results in the paper

- Generalization to problems without strong convexity.
- Lower complexity bound for randomized coordinate descent methods.


## What's new?

- Gradient sliding algorithms for complex composite optimization.
- Saving gradient computation significantly without increasing \# of iterations.
- An optimal randomized incremental gradient for finite-sum optimization.
- Saving gradient computation at the expense of more iterations.
- New lower complexity bound and game-theoretic interpretation for first-order methods.


[^0]:    compressed sensing). Complexity:

