
On Convergence of Model Parallel Proximal Gradient Algorithm for Stale Synchronous Parallel System

Yi Zhou*

Yaoliang Yu[†]
*Syracuse University

Wei Dai[†]

Yingbin Liang*
[†]Carnegie Mellon University

Eric P. Xing[†]

Abstract

With ever growing data volume and model size, an error-tolerant, communication efficient, yet versatile parallel algorithm has become a vital part for the success of many large-scale applications. In this work we propose **msPG**, an extension of the flexible proximal gradient algorithm to the model parallel and stale synchronous setting. The worker machines of **msPG** operate asynchronously as long as they are not too far apart, and they communicate efficiently through a dedicated parameter server. Theoretically, we provide a rigorous analysis of the various convergence properties of **msPG**, and a salient feature of our analysis is its seamless generality that allows both nonsmooth and nonconvex functions. Under mild conditions, we prove the *whole* iterate sequence of **msPG** converges to a critical point (which is optimal under convexity assumptions). We further provide an economical implementation of **msPG**, completely bypassing the need of keeping a local full model. We confirm our theoretical findings through numerical experiments.

1 Introduction

Many machine learning and statistics problems fit into the general composite minimization framework:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}) + g(\mathbf{x}), \quad (1)$$

where the first term is typically a smooth empirical risk over n training samples and the second term g is a nonsmooth regularizer that promotes structures.

Appearing in Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 51. Copyright 2016 by the authors.

Popular examples under this framework include the Lasso (least squares loss f_i and various sparse regularizers g), logistic regression (logistic loss f_i), boosting (exponential loss f_i), support vector machines (hinge loss f_i and RKHS norm regularizer g), matrix completion (least squares loss f and trace norm regularizer g), etc. There is also rising interest in using nonconvex losses (mainly for robustness to outliers) and nonconvex regularizers (mainly for smaller bias in feature selection and support estimation), see e.g. [15, 26, 31, 35, 36, 38–40].

Due to the apparent importance of the composite minimization framework and the rapidly growing size in both model dimension and sample volume, there is a strong need to develop a practical *parallel* system that can solve the problem in (1) efficiently and in a scale that is impossible for a single machine [2, 8, 14, 16, 19, 21, 25, 37]. Existing systems can roughly be divided into three categories: bulk synchronous [14, 34, 37], (totally) asynchronous [8, 25], and partially asynchronous (also called stale synchronous in this work) [2, 8, 16, 19, 21, 32]. The bulk synchronous parallel mechanism (BSP) forces synchronization barriers so that the worker machines can stay on the same page to ensure correctness. However, in a real deployed parallel system BSP usually suffers from the straggler problem, that is, the performance of the whole system is bottlenecked by the *slowest* worker machine. On the other hand, asynchronous systems achieve much greater throughputs, although at the expense of potentially losing the correctness of the algorithm. The stale synchronous parallel (SSP) mechanism is a compromise between the previous two mechanisms: it allows the worker machines to operate asynchronously, as long as they are not too far apart. SSP is particularly suitable for machine learning applications, where iterative algorithms robust to small errors are usually used to find an appropriate model. This view is also practiced by many recent works building on the SSP mechanism [2, 16, 19, 21, 24, 27].

Existing parallel systems can also be divided into data parallel and model parallel. In the former case,

one usually distributes the computation involving each component function f_i in (1) into different worker machines. This is suitable when $n \gg d$, i.e. large data volume but moderate model size. A popular algorithm for this case is the stochastic gradient algorithm and its proximal versions [2, 16, 19, 21], under the SSP mechanism. In contrast, model parallel refers to the regime where $d \gg n$, i.e. large model size but moderate data volume. This is the case for many computational biology and health care problems, where collecting many samples can be very expensive but for each sample we can relatively cheaply take a large number of measurements (features). As a result we need to partition the model \mathbf{x} into different (disjoint) blocks and distribute them among many worker machines. The proximal gradient [11, 18] or its accelerated version [6] is again a natural candidate algorithm due to its nice ability of handling nonsmooth regularizers. However, such proximal gradient algorithm has not been investigated under SSP mechanism for model parallelism, although some other types of asynchronous algorithms are studied before (see Section 7). The main goal of this work is to fill in this important gap.

More specifically, we make the following contributions: 1). We propose **msPG**, an extension of the proximal gradient algorithm to the new model parallel and stale synchronous setting. 2). We provide a rigorous analysis of the convergence properties of **msPG**. Under a very general condition that allows both *nonsmooth* and *nonconvex* functions we prove in Theorem 1 that any limit point of the sequence generated by **msPG** is a critical point. Then, inspired by the recent Kurdyka-Lojasiewicz (KL) inequality [3, 5, 9, 11, 20], we further prove in Theorem 2 that the whole sequence of **msPG** in fact converges to a critical point, under mild technical assumptions that we verify for many familiar examples. Lastly, relating **msPG** to recent works on inexact proximal gradient (on a single machine), we provide, under the new model parallel and SSP setting, a simple proof of the usual sublinear $O(1/t)$ rate of convergence (assuming convexity). We remark our technical contributions with comparison to related work after each main results. 3). Building on the recent parameter server framework [19, 21], we give an economical implementation of **msPG** that completely avoids storing local full models in each worker machine. The resulting implementation only requires storing the partitioned data (with size $O(nd_i)$ for d_i assigned parameters) and communicating a vector of length n in each iteration. 4). We corroborate our theoretical findings with controlled numerical experiments.

This paper proceeds as follows: We first recall some definitions in §2, followed by the proposed algorithm **msPG** in §3. Theoretical findings are reported in §4

(with all proofs deferred to the appendix). An economical implementation of **msPG** is detailed in §5 and experimentally verified in §6. Finally, we discuss some related works in §7 and conclude in §8.

2 Preliminaries

We collect here some useful definitions that will be needed in our later analysis.

Since we consider a proper and closed¹ function $h : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ that may *not* be smooth or convex, we need a generalized notion of “derivative”.

Definition 1 (Subdifferential and critical point, [29]). *The Frechét subdifferential $\hat{\partial}h$ of h at $\mathbf{x} \in \text{dom } h$ is the set of \mathbf{u} such that*

$$\liminf_{\mathbf{z} \neq \mathbf{x}, \mathbf{z} \rightarrow \mathbf{x}} \frac{h(\mathbf{z}) - h(\mathbf{x}) - \mathbf{u}^\top (\mathbf{z} - \mathbf{x})}{\|\mathbf{z} - \mathbf{x}\|} \geq 0, \quad (2)$$

while the (limiting) subdifferential ∂h at $\mathbf{x} \in \text{dom } h$ is the graphical closure of $\hat{\partial}h$:

$$\{\mathbf{u} : \exists \mathbf{x}^k \rightarrow \mathbf{x}, h(\mathbf{x}^k) \rightarrow h(\mathbf{x}), \mathbf{u}^k \in \hat{\partial}h(\mathbf{x}^k) \rightarrow \mathbf{u}\}. \quad (3)$$

The critical points of h are $\text{crit } h := \{\mathbf{x} : \mathbf{0} \in \partial h(\mathbf{x})\}$.

Pleasantly, when h is continuously differentiable or convex, the subdifferential ∂h and critical points $\text{crit } h$ coincide with the usual notions.

Definition 2 (Distance and projection). *The distance function to a closed set $\Omega \subseteq \mathbb{R}^d$ is defined as:*

$$\text{dist}_\Omega(\mathbf{x}) := \min_{\mathbf{y} \in \Omega} \|\mathbf{y} - \mathbf{x}\|, \quad (4)$$

and the metric projection onto Ω is:

$$\text{proj}_\Omega(\mathbf{x}) := \text{argmin}_{\mathbf{y} \in \Omega} \|\mathbf{y} - \mathbf{x}\|. \quad (5)$$

Note that proj_Ω is always a singleton iff Ω is convex.

Definition 3 (Proximal map, e.g. [29]). *A natural generalization of the metric projection using a closed and proper function h is (with parameter $\eta > 0$):*

$$\text{prox}_h^\eta(\mathbf{x}) := \text{argmin}_{\mathbf{z}} h(\mathbf{z}) + \frac{1}{2\eta} \|\mathbf{z} - \mathbf{x}\|^2, \quad (6)$$

where $\|\cdot\|$ is the usual Euclidean norm.

If h decreases slower than a quadratic function (in particular, when h is bounded below), its proximal map is well-defined for all (small) η . For convex h , the proximal map is always a singleton while for nonconvex h , the proximal map can be set-valued. In the latter case we also abuse the notation $\text{prox}_h^\eta(\mathbf{x})$ for an arbitrary element from that set. The proximal map is the key component of the popular proximal gradient algorithms [6, 11, 18].

¹An extended real-valued function h is proper if its domain $\text{dom } h := \{\mathbf{x} : h(\mathbf{x}) < \infty\}$ is nonempty; it is closed iff its sublevel sets $\{\mathbf{x} : h(\mathbf{x}) \leq \alpha\}$ is closed for all $\alpha \in \mathbb{R}$.

Definition 4 (KL function, [10, 20]). A function h is called KL if for all $\bar{\mathbf{x}} \in \text{dom } \partial h$ there exist $\lambda > 0$ and a neighborhood X of $\bar{\mathbf{x}}$ such that for all $x \in X \cap \{\mathbf{x} : h(\bar{\mathbf{x}}) < h(\mathbf{x}) < h(\bar{\mathbf{x}}) + \lambda\}$ the following inequality holds

$$\varphi'(h(\mathbf{x}) - h(\bar{\mathbf{x}})) \cdot \text{dist}_{\partial h(\mathbf{x})}(\mathbf{0}) \geq 1, \quad (7)$$

where the function $\varphi : [0, \lambda) \rightarrow \mathbb{R}_+, 0 \mapsto 0$, is continuous, concave, and has continuous and positive derivative φ' on $(0, \lambda)$.

The KL inequality (7) is an important tool to bound the trajectory length of a dynamical system (see [10, 20] and the references therein for some historic developments). It has recently been used to analyze discrete-time algorithms in [1] and proximal algorithms in [3, 4, 11]. As we shall see, the function φ will serve as a Lyapunov function. Quite conveniently, most practical functions, in particular, “definable” functions and convex functions under certain growth condition, are KL. For a more detailed discussion of KL functions, including many familiar examples, see [11, Section 5] and [4, Section 4].

3 Problem Formulation

We consider the composite minimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}), \quad \text{where } F(\mathbf{x}) = f(\mathbf{x}) + g(\mathbf{x}). \quad (\text{P})$$

Usually f is a smooth loss function and g is a regularizer that promotes structure. We consider the **model parallel** scenario, that is, we decompose the d model parameters into p disjoint groups, and designate one worker machine for each group. Formally, consider the decomposition $\mathbb{R}^d = \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \times \dots \times \mathbb{R}^{d_p}$, $\mathbf{x} = (x_1, x_2, \dots, x_p)$, and let $\nabla_i f : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ be the partial gradient of f on the i -th factor space (machine). Clearly, $x_i, \nabla_i f(\mathbf{x}) \in \mathbb{R}^{d_i}$ and $\sum_{i=1}^p d_i = d$. The i -th machine is responsible for the i -th factor $x_i \in \mathbb{R}^{d_i}$, however, we also allow machine i to keep a local copy $\mathbf{x}^i \in \mathbb{R}^d$ of the *full* model parameter. This is for the convenience of evaluating the partial gradient $\nabla_i f : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$, and we will discuss in Section 5 how to implement this in an economical way. Note that unlike the **data parallel** setting, we do *not* consider explicitly distributing the computation of the gradient $\nabla_i f$.

We extend the proximal gradient algorithm [11, 18] to solve the composite problem (P) under the new model parallel setting, and we require the following standard assumptions for our convergence analysis.

Assumption 1. Regarding the functions f, g in (P):

1. They are bounded from below;

2. The function f is differentiable and the gradients $\nabla f, \nabla_i f$ are Lipschitz continuous with constant L_f and L_i , respectively. Set $L = \sum_{i=1}^p L_i$;
3. The function g is closed, and separable, i.e., $g(\mathbf{x}) = \sum_{i=1}^p g_i(x_i)$.

The first two assumptions are needed to analyze the proximal gradient algorithm even in the convex and non-distributed setting, and the third assumption is what makes model parallelism interesting (and feasible). We remark that the differentiability assumption on f can be easily relaxed by smoothing, and the separability assumption on g can also be relaxed using the proximal average idea in [36]. For brevity we do not pursue these extensions here. One salient feature of our analysis is that we do *not* assume convexity on either f or g (although our conclusions are considerably stronger under convexity).

The separability assumption above on g implies that

$$\text{prox}_g^\eta(\mathbf{x}) = (\text{prox}_{g_1}^\eta(x_1), \dots, \text{prox}_{g_p}^\eta(x_p)). \quad (8)$$

Let us introduce the update operator (on machine i):

$$U_i(\mathbf{x}^i) = U_i(\mathbf{x}^i, x_i) := \text{prox}_{g_i}^\eta(x_i - \eta \nabla_i f(\mathbf{x}^i)) - x_i, \quad (9)$$

i.e. machine i computes the i -th part of the gradient using its local model \mathbf{x}^i , updates its parameter x_i in charge using step size η , and finally applies the proximal map of the component function g_i . In a real large scale parallel system, the communication delay among machines and the unexpected shut down of machines are practical issues that bottlenecks the performance of the system, and hence a more relaxed synchronization protocol than full synchronization is needed. Consider a global clock shared by all machines and denote T_i the set of active clocks when machine i takes an update, and $\mathbb{I}_{\{t \in T_i\}}$ as the indicator function of the event $t \in T_i$. Formally, the t -th iteration on machine i can be written as:

$$\text{msPG} \begin{cases} \forall i, x_i(t+1) = x_i(t) + \mathbb{I}_{\{t \in T_i\}} U_i(\mathbf{x}^i(t)), \\ (\text{local}) \quad \mathbf{x}^i(t) = (x_1(\tau_1^i(t)), \dots, x_p(\tau_p^i(t))), \\ (\text{global}) \quad \mathbf{x}(t) = (x_1(t), \dots, x_p(t)), \end{cases}$$

where $0 \leq \tau_j^i(t) \leq t$ models the delay among machines: when machine i conducts its t -th update it only has access to $x_j(\tau_j^i(t))$, a delayed version of the factor $x_j(t)$ on the j -th machine. In the above, we collect the fresh parameters of all machines to form a global model $\mathbf{x}(t)$, which brings convenience for our analysis. We will refer to the above updates as msPG (for **m**odel parallel, **s**tale synchronous, **P**roximal **G**radient). Figure 1 illustrates the main idea of msPG.

Obviously, to establish convergence for msPG we need some control over the delay $\tau_j^i(t)$ and the active




Figure 1: The algorithm msPG under model parallelism and stale synchronism. Machine i keeps a local model $\mathbf{x}^i(t)$ that contains stale parameters of other machines (due to communication delay and network latency). These local models are used to compute the partial gradient $\nabla_i f(\mathbf{x}^i)$ which is then used to update the parameters $x_i(t)$ in each machine. See Section 5 for an economical implementation of msPG.

clocks T_i , for otherwise some machines may not make progress at all. We first impose the following assumptions on the delay and active clocks, and follow by some explanations.

Assumption 2. *The delay and skip frequency satisfy:*

1. $\forall i, \forall j, \forall t, 0 \leq t - \tau_j^i(t) \leq s$;
2. $\forall i, \forall t, \tau_i^i(t) = t$;
3. $\forall i, \forall t, T_i \cap \{t, t+1, \dots, t+s\} \neq \emptyset$.

The first assumption basically says that if machine i conducts its t -th update, then the information it gathered from other machines cannot be too obsolete (bounded by s iterations). The second assumption simply says that machine i always has the latest information on itself. The third assumption requires each machine to update at least once in every s iterations. These assumptions are very natural and have been widely adopted in previous works [2, 8, 16, 19, 22, 24, 32]. They are also in some sense unavoidable: one can construct instances such that msPG do not converge if these assumptions are violated. Clearly, when $s = 0$ (no delay) our framework reduces to the bulk synchronous proximal gradient algorithm.

4 Convergence Analysis

In this section, we conduct detailed analysis of the model parallel stale synchronous proximal gradient algorithm msPG. Our first result is as follows:

Theorem 1 (Asymptotic consistency). *Let Assumption 1 and 2 hold, and apply msPG to problem (P). If the step size $\eta < (L_f + 2Ls)^{-1}$, then the global model and local models satisfy:*

1. $\sum_{t=0}^{\infty} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 < \infty$;
2. $\lim_{t \rightarrow \infty} \|\mathbf{x}(t+1) - \mathbf{x}(t)\| = 0, \lim_{t \rightarrow \infty} \|\mathbf{x}(t) - \mathbf{x}^i(t)\| = 0$;
3. *The limit points $\omega(\{\mathbf{x}(t)\}) = \omega(\{\mathbf{x}^i(t)\}) \subseteq \text{crit } F$.*

The proof is a bit involved and can be found in Appendix A. The first assertion says that the global sequence $\mathbf{x}(t)$ defined in msPG has square summable successive differences, which we will significantly strengthen below. The second assertion implies that the successive difference of the global sequence diminishes, and the inconsistency between the local sequences and the global sequence also vanishes. These two conclusions provide some stability guarantee about our algorithm msPG. The third assertion further justifies msPG by showing that any limit point it produces is necessarily a critical point. Of course, when F is convex, any critical point is globally optimal.

The closest result to Theorem 1 we are aware of is [8, Proposition 7.5.3], where essentially the same conclusion was reached but under the much more simplified assumption that g is an indicator function of a convex set. Thus, our Theorem 1 is new even when g is a convex function such as the popular ℓ_1 norm. Furthermore, we allow g to be nonconvex, which is particularly interesting due to the rising interest in nonconvex penalties in machine learning and statistics (see e.g. [15, 26, 31, 35, 36, 38–40]). We also note that the proof of Theorem 1 (for nonconvex g) involves significantly new ideas beyond those of [8].

Interesting as it is, Theorem 1 has one significant deficiency, though: it does not tell us when there exists a limit point, and it does not guarantee the whole sequence to converge to the limit point. In fact, in the model parallel setting with delays and skips, it is even a nontrivial task to argue that the objective values $\{F(\mathbf{x}(t))\}$ do not diverge to infinity. This is in sharp contrast with the bulk synchronous setting where it is trivial to guarantee the objective values to decrease (by using a sufficiently small step size). This is where we need some further assumptions.

Assumption 3 (Sufficient Decrease). *There exists $\alpha > 0$ such that the global model $\mathbf{x}(t)$ generated by*

msPG (for problem (P)) satisfies: for all large t ,

$$F(\mathbf{x}(t+1)) \leq F(\mathbf{x}(t)) - \alpha \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2. \quad (10)$$

We will provide below a sufficient condition in Lemma 1 and many interesting examples in Example 1 to ensure Assumption 3 to hold.

Assumption 4. F is a KL function.

As we mentioned in the end of Section 2, most practical functions (including all functions in this work) are KL. Hence, this is a very mild assumption.

We are now ready to state the much more strengthened convergence guarantee for *msPG*:

Theorem 2 (Finite Length). *Let Assumption 1, 2, 3 and 4 hold, and apply *msPG* to problem (P). If the step size $\eta < (L_f + 2L_s)^{-1}$ and $\{\mathbf{x}(t)\}$ is bounded, then*

$$\sum_{t=0}^{\infty} \|\mathbf{x}(t+1) - \mathbf{x}(t)\| < \infty, \quad (11)$$

$$\forall i = 1, \dots, p, \sum_{t=0}^{\infty} \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\| < \infty. \quad (12)$$

Furthermore, $\{\mathbf{x}(t)\}$ and $\{\mathbf{x}^i(t)\}, i = 1, \dots, p$, converge to the same critical point of F .

The proof, in Appendix B, borrowed some ideas from the recent works [3, 4, 11], which, however, did not consider the parallel and asynchronous setting. Compared with the first assertion in Theorem 1, we now have the successive differences to be absolutely summable (instead of square summable). The former property is usually called finite length in dynamical systems. It is a significantly stronger property as it immediately implies that the whole sequence is Cauchy hence convergent, whereas we cannot get the same conclusion from the square summable property in Theorem 1.² We note that local maxima are excluded from being the limit in Theorem 2, thanks to Assumption 3. Also, the boundedness assumption on $\{\mathbf{x}(t)\}$ is easy to satisfy, for instance, when F has bounded sublevel sets. We refer to [4, Remark 3.3] for more conditions that guarantee the boundedness. Needless to say, if F is convex, then the whole sequence in Theorem 2 converges to a global minimizer.

The closest result to Theorem 2, to our best knowledge, is [32], which proved an s -step linear rate when f satisfies a certain error bound condition and g is the indicator function of a convex set. In contrast, we allow any convex or nonconvex function g (as long as it is KL). Our proof, inspired by [3, 4, 11], differs substantially from [32].

²A simple example would be the sequence $x(t) = \sum_{k=1}^t \frac{1}{k}$, whose successive difference is square summable but clearly $x(t)$ does not converge. Consequently, $x(t)$ is not absolutely summable.

We now provide some justifications on Assumption 3 by considering the simplified case where $\forall t, i, t \in T_i$, i.e., machines do not skip updates. The general case can also be dealt with but the analysis is much more complicated. The following condition turns out to be a good justification for Assumption 3.

Assumption 5 (Eventual Lipschitz). *The update operators $U_i, i = 1, \dots, p$ are eventually Lipschitz continuous, i.e., for all large t and small learning rate $\eta > 0$:*

$$\|U_i(\mathbf{x}^i(t+1)) - U_i(\mathbf{x}^i(t))\| \leq C_i \eta \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\|, \quad (13)$$

where $C_i \geq L_i, i = 1, \dots, p$, are positive constants.

Note that when $g \equiv 0$, $U_i = -\eta \nabla_i f$ is Lipschitz continuous (due to Assumption 1.2), thus Assumption 5 is a natural generalization to an arbitrary regularizer g . Equipped with this assumption, we can now justify Assumption 3. (Proof is in Appendix C.) In the sequel, we denote $C = \sum_{i=1}^p C_i \geq L$.

Lemma 1. *Assume $\forall t, i, t \in T_i$. Let the step size $\eta < \frac{\rho-1}{4C\rho} \frac{\sqrt{\rho}-1}{\sqrt{\rho^2+1}-1}$ for any $\rho > 1$ and all $U_i, i = 1, \dots, p$ be eventually Lipschitz continuous, then the sequences $\{\mathbf{x}(t)\}$ and $\{\mathbf{x}^i(t)\}, i = 1, \dots, p$, have finite length.*

Hence it is sufficient to further characterize Assumption 5, which turns out to be a mild condition. Instead of giving a very technical justification, we give here some popular examples where Assumption 5 holds (proof in Appendix D). Some of these will also be tested in our experiments.

Example 1. *Assume $\forall t, i, t \in T_i$, then Assumption 5 holds for the following cases (modulo a technical condition on the 1-norm):*

- $g \equiv 0$ (no regularization), $g = \|\cdot\|_0$ (nonconvex 0-norm), $g = \|\cdot\|_1$ (1-norm), $g = \|\cdot\|^2$ (squared 2-norm);
- elastic net $g = \|\cdot\|_1 + \|\cdot\|^2$ and its nonconvex variation $g = \|\cdot\|_0 + \|\cdot\|^2$;
- non-overlapping group norms $g = \|\cdot\|_{0,2}$ and $g = \|\cdot\|_{0,2} + \|\cdot\|^2$.

Further for this non-skip case ($\forall t, i, t \in T_i$), *msPG* can be cast as an inexact proximal gradient algorithm (IPGA), which, together with the finite length property in Theorem 2, provide new insights on the nature of staleness in real parallel systems. It also allows us to easily obtain the usual $O(1/t)$ rate of convergence of the objective value.

Let us introduce an error term $\mathbf{e}(t) = (e_1(t), \dots, e_p(t))$, with which we can rewrite the global sequence of *msPG* as:

$$\forall i, x_i(t+1) = \text{prox}_{g_i}^{\eta}(x_i(t) - \eta \nabla_i f(\mathbf{x}(t)) + e_i(t)), \quad (14)$$

where $e_i(t) = \eta[\nabla_i f(\mathbf{x}(t)) - \nabla_i f(\mathbf{x}^i(t))]$. This alternative representation of msPG falls under the IPGA in [30], where $\mathbf{e}(t)$ is the (gradient) error at iteration t . Note, however, that the error $\mathbf{e}(t)$ is not caused by computation but by communication delay and network latency, which only presents itself in a real stale synchronous parallel system. For convex functions F , [30] showed that the convergence rate of the objective value of IPGA can be controlled by the summability of the error magnitude $\|\mathbf{e}(t)\|$. Interestingly, our next result proves that the finite length property in Theorem 2 immediately implies the summability of the errors $\|\mathbf{e}\|$, even for nonconvex functions f and g . Moreover, for convex F , this also leads to the usual $O(1/t)$ rate of convergence in terms of the objective value.

Theorem 3 (Global rate of convergence). *If the finite length property in Theorem 2 holds, then*

1. $\sum_{t=0}^{\infty} \|\mathbf{e}(t)\| < \infty$;
2. $F(\frac{1}{t} \sum_{k=1}^t \mathbf{x}(k)) - \inf F \leq O(t^{-1})$.

The proof is in Appendix E. Intuitively, if the error $\mathbf{e}(t)$ decreases (slightly) faster than $O(1/t)$, then the rate of convergence of msPG is not affected even under the model parallel and stale synchronous setting (provided F is convex). To the best of our knowledge, this is the first *deterministic* rate of convergence result in the model parallel and stale synchronous setting.

5 Economical Implementation

In this section, we show how to economically implement msPG for the widely used linear models:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(A\mathbf{x}) + g(\mathbf{x}), \quad (15)$$

where $A \in \mathbb{R}^{n \times d}$. Typically $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is the likelihood function and $g: \mathbb{R}^d \rightarrow \mathbb{R}$ is the regularizer (we absorb the regularization constant into g). Each row of A corresponds to a sample point and we have suppressed the labels in classification or responses in regression. Support vector machines, Lasso, logistic regression, boosting, etc. all fit under this framework. Our interest here is when the model dimension d is much higher than the number of samples n (d can be up to hundreds of millions and n can be up to millions). This is the usual setup in many computational biology and health care problems.

A naive implementation of msPG might be inefficient in terms of both network communication and parameter storage. First, Each machine needs to communicate with every other machine, to exchange the latest block of parameters. If using a peer-to-peer network topology, the resulting connections will be too dense and crowded when the system holds hundreds of machines. We resolve this issue by adopting the parame-

ter server system advocated in previous works [19, 21], that is, we dedicate a specific server (which can span a set of machines if needed) to store the key parameters (will be specified later) and let each worker machine to communicate only with the server. There is a second advantage for this master-slave network topology, as we shall see momentarily.

Second, each machine needs to keep a local copy of the full model (i.e. $\mathbf{x}^i(t)$ in msPG), which can incur a very expensive storage cost when the dimension is high. This is where the linear model structure in (15) comes into help. Note that the local models $\mathbf{x}^i(t)$ are kept solely for the convenience of evaluating the partial gradient $\nabla_i f: \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$. For some problems such as the Lasso, a seemingly workaround is to pre-compute the Hessian $H = A^\top A$ and distribute the corresponding row blocks of H to each worker machine. This scheme, however, is problematic in the high dimensional setting: the pre-computation of the Hessian can be very costly, and each row block of H has a very large size ($d_i \times d$).

Instead, we use the column partition scheme [e.g. 12, 28], namely, we partition the matrix A into p column blocks $A = [A_1, \dots, A_p]$ and distribute the block $A_i \in \mathbb{R}^{n \times d_i}$ to machine i . Now the local update computed by machine i at the t -th iteration can be rewritten as

$$U_i(\mathbf{x}^i(t)) = \text{prox}_{g_i}^\eta(x_i(t) - \eta A_i^\top f'(A\mathbf{x}^i(t))) - x_i(t) \quad (16)$$

Since machine i is in charge of updating the i -th block $x_i(t)$ of the global model, to compute the local update (16) it is sufficient to have the matrix-vector product $A\mathbf{x}^i(t)$. For simplicity we initialize $\forall i, \mathbf{x}^i(0) \equiv \mathbf{0}$, then we have the following cumulative form:

$$A\mathbf{x}^i(t) = \sum_{j=1}^p A_j [\mathbf{x}^i(t)]_j = \sum_{j=1}^p \sum_{k=0}^{\tau_j^i(t)} \underbrace{A_j \mathbb{I}_{\{k \in T_j\}} U_j(\mathbf{x}^j(k))}_{\Delta_j(k)},$$

where recall that when machine i conducts its t -th iteration it only has access to a delayed copy $x_j(\tau_j^i(t))$ of the parameters in machine j . Since this matrix-vector product is needed by every machine to conduct their local updates in (16), we aggregate $\Delta_j(t) \in \mathbb{R}^n$ on the parameter server whenever it is generated and sent by the worker machines. In details, the worker machines first pull this aggregated matrix-vector product (denoted as \blacktriangle) from the server to conduct the local computation (16) in an economical way (by replacing $A\mathbf{x}^i(t)$ in (16) with \blacktriangle). Then machine i performs the simple update:

$$x_i(t+1) = x_i(t) + U_i(\mathbf{x}^i(t)). \quad (17)$$

Note that machine i does not maintain or update other blocks of parameters $x_j(t), j \neq i$. Lastly, machine i

Algorithm 1 Economic Implementation of msPG

```

1: For the server:
2:   while receives update  $\Delta_i$  from machine  $i$  do
3:      $\mathbf{\blacktriangle} \leftarrow \mathbf{\blacktriangle} + \Delta_i$ 
4:   end while
5:   while machine  $i$  sends a pull request do
6:     send  $\mathbf{\blacktriangle}$  to machine  $i$ 
7:   end while
8: For machine  $i$  at active clock  $t \in T_i$ :
9:   pull  $\mathbf{\blacktriangle}$  from the server
10:   $U_i \leftarrow \text{prox}_{g_i}^\eta(x_i - \eta A_i^\top f'(\mathbf{\blacktriangle})) - x_i$ 
11:  send  $\Delta_i = A_i U_i$  to the server
12:  update  $x_i \leftarrow x_i + U_i$ 
    
```

computes and sends the vector $\Delta_i(t) = A_i U_i(\mathbf{x}^i(t)) \in \mathbb{R}^n$ to the server, and the server immediately performs the aggregation:

$$\mathbf{\blacktriangle} \leftarrow \mathbf{\blacktriangle} + \Delta_i(t). \quad (18)$$

We summarize the above economical implementation in Algorithm 1, where $\mathbf{\blacktriangle}$ denotes the aggregation of individual matrix-vector products Δ on the server. The storage cost for each worker machine is $O(nd_i)$ (for storing A_i). Each iteration requires two matrix-vector products that cost $O(nd_i)$ in the dense case, and the communication of a length n vector between the server and the worker machines.

6 Experiments

We first test the convergence properties of msPG via a non-convex Lasso problem with the group regularizer $\|\cdot\|_{0,2}$, which takes the form

$$\min_{\mathbf{x}} \frac{1}{2} \|A\mathbf{x} - \mathbf{b}\|^2 + \lambda \sum_{i=1}^{20} \mathbb{I}(\|x_i\| = 0), \quad (19)$$

where $A \in \mathbb{R}^{1000 \times 2000}$ and we refer to Appendix F for the specifications of data generation. We use 4 machines (cores) with each handling five groups of coordinates, and consider staleness $s = 0, 10, 20, 30$, respectively. To better demonstrate the effect of staleness, we let machines only communicate when exceed the maximum staleness. This can be viewed as the worst case communication scheme and a larger s brings more staleness into the system. We set the learning rate to have the form $\eta(\alpha s) = 1/(L_f + 2L\alpha s)$, $\alpha > 0$, that is, a linear dependency on staleness s as suggested by Theorem 1. Then we run Algorithm 1 with different staleness and use $\eta(0), \eta(10), \eta^*(\alpha s)$, respectively, where $\eta^*(\alpha s)$ is the largest step size we tuned for each s that achieves a stable convergence. We track the global model $\mathbf{x}(t)$ and plot the results in Figure 2. Note that with the large step size $\eta(0)$ all instances (with nonzero

staleness) diverge hence are not presented. With $\eta(10)$ (Figure 2, left), the staleness does not substantially affect the convergence in terms of the objective value. We note that the objective curves converge to slightly different minimal values due to the non-convexity of problem (19). With $\eta^*(\alpha s)$ (Figure 2, middle), it can be observed that adding a slight penalty αs on the learning rate suffices to achieve a stable convergence, and the penalty grows as s increases, which is intuitive since a larger staleness requires a smaller step size to cancel the inconsistency. In particular, for $s = 10$ the best convergence is comparable to the bulk synchronized case $s = 0$. (Figure 2, right) further shows the asymptotic convergence behavior of the global model $\mathbf{x}(t)$ under the step size $\eta^*(\alpha s)$. It is clear that a linear convergence is eventually attained, which confirms the finite length property in Theorem 2.

Next, we verify the time and communication efficiency of msPG via an l_1 norm Lasso problem with very high dimensions, taking the form

$$\min_{\mathbf{x}} \frac{1}{2} \|A\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_1. \quad (20)$$

We generate $A \in \mathbb{R}^{n \times d}$ of size $n = 1\text{Million}$ and $d = 100\text{Millions}$. See Appendix F for the specifications of data generation. We implement Algorithm 1 on Petuum [13, 19] — a stale synchronous parallel system which eagerly updates the local parameter caches via stale synchronous communications. The system contains 100 computing nodes and each is equipped with 16 AMD Opteron processors and 16GB RAM linked by 1Gbps ethernet. We fix the learning rate $\eta = 10^{-3}$ and consider maximum staleness $s = 0, 1, 3, 5, 7$, respectively. (Figure 3, left) shows that per-iteration progress is virtually indistinguishable among various staleness settings, which is consistent with our previous experiment. (Figure 3, middle) shows that system throughput is significantly higher when we introduce staleness. This is due to lower synchronization overheads, which offsets any potential loss due to staleness in progress per iteration. We also track the distributions of staleness during the experiments, where we record in $\mathbf{\blacktriangle}$ the clocks of the freshest updates that accumulate from all the machines. Then whenever a machine pulls $\mathbf{\blacktriangle}$ from the server, it compares its local clock with these clocks and records the clock differences. (Figure 3, right) shows the distributions of staleness under different maximal staleness settings. Observe that bulk synchronous ($s = 0$) peaks at staleness 0 by design, and the distribution concentrates in small staleness area due to the eager communication mechanism of Petuum. It can be seen that a small amount of staleness is sufficient to relax the communication bottlenecks without affecting the iterative convergence rate much.

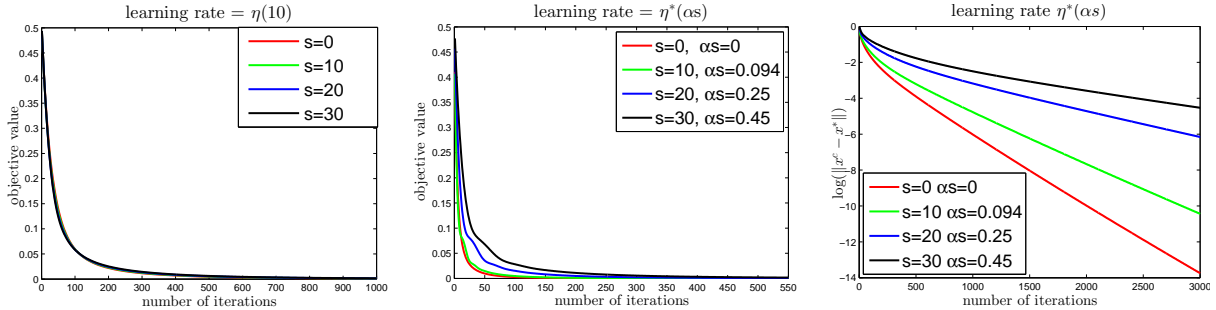
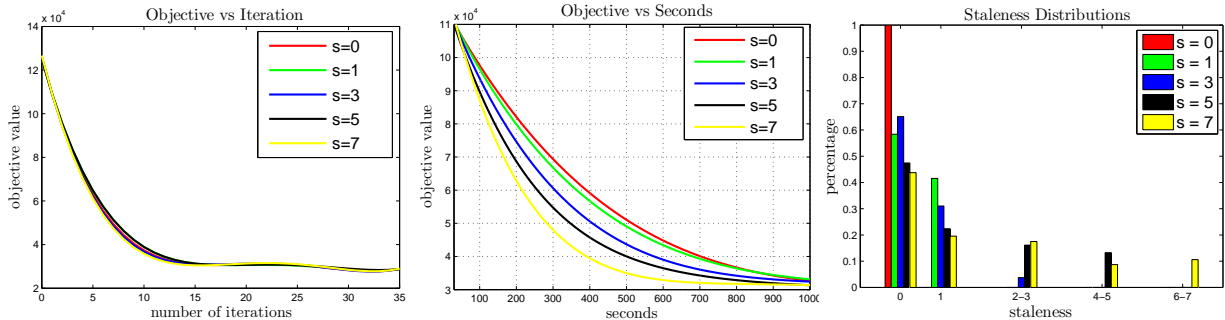

 Figure 2: Convergence curves of msPG under different staleness parameter s and step size η .


Figure 3: Efficiency of msPG on a large scale Lasso problem.

7 Related Work

The stale synchronous parallel system dates back to [7, 8, 32, 33], where it is also referred to as partially asynchronous system. These works consider using stale synchronous systems to solve different kinds of optimization problems with allowing machines to skip updates during the process. Asymptotic convergence of partially asynchronous gradient descent algorithm for solving unconstrained smooth optimizations is established in [8], with its stochastic version being analyzed in [33]. Asymptotic convergence of partially asynchronous gradient projection algorithm for solving smooth optimizations with convex constraint is established in [8], and a “B-step” linear convergence is further established in [32] with an error bound condition. Linear convergence of partially asynchronous algorithm for finding the fixed point of maximum norm contraction mappings is established in [7, 17].

Another series of work focus on SSP systems where machines are not allowed to skip updates [21–23]. In their settings, the system imposes an upper bound on the maximum clock difference between machines. Asymptotic convergence is established for proximal gradient algorithm for data parallelism [23] and for block coordinate descent [22] with a smooth objective and convex regularizer. Other works consider stochastic algorithms on stale synchronous system. [19] proposes an SSP system for stochastic gradient descent, and establishes $O(1/\sqrt{k})$ regret bound under bounded diameter and bounded sub-gradient assump-

tion. [16, 27] consider a delayed stochastic gradient descent algorithm. Linear convergence to a neighborhood of optimum is established with strong convexity assumption in [16] and with additional bounded gradient assumption in [27]. [2] proposes a distributed delayed dual averaging and mirror descent algorithm, and establishes $O(1/\sqrt{k})$ regret bound under standard stochastic assumptions.

8 Conclusion

We have proposed msPG as an extension of the proximal gradient algorithm to the model parallel and stale synchronous setting. msPG allows worker machines to operate asynchronously as long as they are not too far apart, hence greatly improves the system throughput. Theoretically, we provide a rigorous analysis of msPG that simultaneously covers nonsmooth and nonconvex functions. In particular, under mild conditions, the whole iterate sequence generated by msPG converges to a critical point. We implement msPG using the parameter server platform, and completely bypass the need of keeping a local full model. Preliminary numerical experiments confirm the effectiveness of msPG on solving very high dimensional problems.

Acknowledgement

We thank the reviewers for their valuable comments. This work is supported by NIH R01GM114311, DARPA FA87501220324 and NSF IIS1447676.

References

- [1] P.-A. Absil, R. Mahony, and B. Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM Journal on Optimization*, 16(2):531–547, 2005.
- [2] A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 24*, pages 873–881. 2011.
- [3] H. Attouch and J. Bolte. On the convergence of the proximal algorithm for nonsmooth functions involving analytic features. *Mathematical Programming*, 116(1-2):5–16, 2009.
- [4] H. Attouch, J. Bolte, P. Redont, and A. Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Lojasiewicz inequality. *Mathematics of Operations Research*, 35(2):438–457, 2010.
- [5] H. Attouch, J. Bolte, and B. Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods. *Mathematical Programming*, 137(1-2):91–129, 2013.
- [6] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Img. Sci.*, 2(1):183–202, 2009.
- [7] D. P. Bertsekas and J. N. Tsitsiklis. Convergence rate and termination of asynchronous iterative algorithms. In *Proceedings of the 3rd International Conference on Supercomputing*, pages 461–470, 1989.
- [8] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1989.
- [9] J. Bolte, A. Daniilidis, and A. Lewis. The Lojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM Journal on Optimization*, 17:1205–1223, 2007.
- [10] J. Bolte, A. Daniilidis, O. Ley, and L. Mazet. Characterizations of Lojasiewicz inequalities and applications: Subgradient flows, talweg, convexity. *Transactions of the American Mathematical Society*, 362(6):3319–3363, 2010.
- [11] J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014.
- [12] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.
- [13] W. Dai, A. Kumar, J. Wei, Q. Ho, G. Gibson, and E. P. Xing. High-performance distributed ml at scale through parameterserver consistency models. In *AAAI*, 2014.
- [14] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. *Communications of ACM*, 51(1):107–113, 2008.
- [15] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- [16] H. Feyzmahdavian, A. Aytakin, and M. Johansson. A delayed proximal gradient method with linear convergence rate. In *2014 IEEE International Workshop on Machine Learning for Signal Processing*.
- [17] H. Feyzmahdavian and M. Johansson. On the convergence rates of asynchronous iterations. In *2014 IEEE 53rd Annual Conference on Decision and Control*, pages 153–159, 2014.
- [18] M. Fukushima and H. Mine. A generalized proximal point algorithm for certain non-convex minimization problems. *International Journal of Systems Science*, 12(8):989–1000, 1981.
- [19] Q. Ho, J. Cipar, H. Cui, S. Lee, J. K. Kim, P. B. Gibbons, G. A. Gibson, G. Ganger, and E. P. Xing. More effective distributed ml via a stale synchronous parallel parameter server. In *Advances in Neural Information Processing Systems 26*, pages 1223–1231. 2013.
- [20] K. Kurdyka. On gradients of functions definable in o-minimal structures. *Annales de l’institut Fourier*, 48(3):769–783, 1998.
- [21] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su. Scaling distributed machine learning with the parameter server. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pages 583–598, 2014.
- [22] M. Li, D. G. Andersen, and A. Smola. Distributed delayed proximal gradient methods. *Big Learning NIPS Workshop*, 2013.
- [23] M. Li, L. Zhou, Z. Yang, A. Li, F. Xia, D. G. Andersen, and A. Smola. Parameter server for distributed machine learning. *Big Learning NIPS Workshop*, 2013.
- [24] J. Liu and S. J. Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.
- [25] Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein. Distributed graphlab: A framework for machine learning and data mining in the cloud. *Proc. VLDB Endow.*, 5(8):716–727, 2012.
- [26] R. Mazumder, J. H. Friedman, and T. Hastie. Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association*, 106(495):1125–1138, 2011.
- [27] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems 24*, pages 693–701. 2011.
- [28] P. Richtárik and M. Takáč. Distributed Coordinate Descent Method for Learning with Big Data. *arXiv*, 2013.
- [29] R. Rockafellar and R. Wets. *Variational Analysis*. Springer, 1997.
- [30] M. Schmidt, N. L. Roux, and F. R. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In *Advances in Neural Information Processing Systems 24*, pages 1458–1466. 2011.
- [31] Y. She. Thresholding-based iterative selection procedures for model selection and shrinkage. *Electronic Journal of Statistics*, 3:384–415, 2009.

- [32] P. Tseng. On the rate of convergence of a partially asynchronous gradient projection algorithm. *SIAM Journal on Optimization*, 1(4):603–619, 1991.
- [33] J. Tsitsiklis, D. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.
- [34] L. G. Valiant. A bridging model for parallel computation. *Communications of ACM*, 33(8):103–111, 1990.
- [35] Z. Wang, H. Liu, and T. Zhang. Optimal computational and statistical rates of convergence for sparse nonconvex learning problems. *The Annals of Statistics*, 42(6):2164–2201, 2014.
- [36] Y. Yu, X. Zheng, M. Marchetti-Bowick, and E. P. Xing. Minimizing nonconvex non-separable functions. In *The 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [37] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: Cluster computing with working sets. pages 10–10, 2010.
- [38] C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *Annals of Statistics*, 38(2):894–942, 2010.
- [39] C.-H. Zhang and T. Zhang. A general theory of concave regularization for high-dimensional sparse estimation problems. *Statistical Science*, 27(4):576–593, 2012.
- [40] Y. Zhu, X. Shen, and W. Pan. Simultaneous grouping pursuit and feature selection over an undirected graph. *Journal of the American Statistical Association*, 108(502):713–725, 2013.

Supplementary Materials

A Proof of Theorem 1

Theorem 1 (Asymptotic consistency). *Let Assumption 1 and 2 hold, and apply msPG to problem (P). If the step size $\eta < (L_f + 2Ls)^{-1}$, then the global model and local models satisfy:*

1. $\sum_{t=0}^{\infty} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 < \infty$;
2. $\lim_{t \rightarrow \infty} \|\mathbf{x}(t+1) - \mathbf{x}(t)\| = 0$, $\lim_{t \rightarrow \infty} \|\mathbf{x}(t) - \mathbf{x}^i(t)\| = 0$;
3. The limit points $\omega(\{\mathbf{x}(t)\}) = \omega(\{\mathbf{x}^i(t)\}) \subseteq \text{crit } F$.

Proof. We start from bounding the difference between the global model \mathbf{x} and the local model \mathbf{x}^i (on any machine i). Indeed, at iteration t , by the definition of the global and local models in msPG:

$$\|\mathbf{x}(t) - \mathbf{x}^i(t)\| = \sqrt{\sum_{j=1}^p \|x_j(t) - x_j(\tau_j^i(t))\|^2} \quad (21)$$

$$\text{(triangle inequality)} \leq \sqrt{\sum_{j=1}^p \left(\sum_{k=\tau_j^i(t)}^{t-1} \|x_j(k+1) - x_j(k)\| \right)^2} \quad (22)$$

$$\text{(Assumption 2.1)} \leq \sqrt{\sum_{j=1}^p \left(\sum_{k=(t-s)_+}^{t-1} \|x_j(k+1) - x_j(k)\| \right)^2} \quad (23)$$

$$= \left\| \left(\sum_{k=(t-s)_+}^{t-1} \|x_1(k+1) - x_1(k)\|, \dots, \sum_{k=(t-s)_+}^{t-1} \|x_p(k+1) - x_p(k)\| \right) \right\| \quad (24)$$

$$= \left\| \sum_{k=(t-s)_+}^{t-1} (\|x_1(k+1) - x_1(k)\|, \dots, \|x_p(k+1) - x_p(k)\|) \right\| \quad (25)$$

$$\text{(triangle inequality)} \leq \sum_{k=(t-s)_+}^{t-1} \left\| (\|x_1(k+1) - x_1(k)\|, \dots, \|x_p(k+1) - x_p(k)\|) \right\| \quad (26)$$

$$= \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\|, \quad (27)$$

where in the last equality we used the following property of the Euclidean norm:

$$\|\mathbf{x}\| = \|(x_1, \dots, x_p)\| = \|(\|x_1\|, \dots, \|x_p\|)\|. \quad (28)$$

Equation (27) bounds the inconsistency between the global model and the local models. We will repeatedly use it in the following, as it provides a bridge to jump from the global model and the local models back and forth.

Next we bound the progress of the global model $\mathbf{x}(t)$. If $t \in T_i$ (i.e., machine i updates at iteration t), then using the definition of the update operator $U_i(\mathbf{x}^i(t))$ in Equation (9) we can rewrite $x_i(t+1)$ as

$$x_i(t+1) = \text{prox}_{g_i}^{\eta}(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))), \quad (29)$$

where we recall the proximal map $\text{prox}_{g_i}^{\eta}$ from Definition 3. Thus, for all $z \in \mathbb{R}^{d_i}$:

$$g_i(x_i(t+1)) + \frac{1}{2\eta} \|x_i(t+1) - x_i(t) + \eta \nabla_i f(\mathbf{x}^i(t))\|^2 \leq g_i(z) + \frac{1}{2\eta} \|z - x_i(t) + \eta \nabla_i f(\mathbf{x}^i(t))\|^2. \quad (30)$$

Substituting with $z = x_i(t)$ and simplifying yields

$$g_i(x_i(t+1)) - g_i(x_i(t)) \leq -\frac{1}{2\eta} \|x_i(t+1) - x_i(t)\|^2 - \langle \nabla_i f(\mathbf{x}^i(t)), x_i(t+1) - x_i(t) \rangle. \quad (31)$$

(If g_i is convex, we can replace $\frac{1}{2\eta}$ with $\frac{1}{\eta}$.) Note that if $t \notin T_i$, then $x_i(t+1) = x_i(t)$ hence Equation (31) still trivially holds. On the other hand, Assumption 1.2 implies

$$f(\mathbf{x}(t+1)) - f(\mathbf{x}(t)) \leq \langle \mathbf{x}(t+1) - \mathbf{x}(t), \nabla f(\mathbf{x}(t)) \rangle + \frac{L_f}{2} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2. \quad (32)$$

Adding Equation (32) and Equation (31) (for all i) and recalling $F(\mathbf{x}) = f(\mathbf{x}) + \sum_i g_i(x_i)$, we have

$$F(\mathbf{x}(t+1)) - F(\mathbf{x}(t)) \leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \sum_{i=1}^p \langle x_i(t+1) - x_i(t), \nabla_i f(\mathbf{x}(t)) - \nabla_i f(\mathbf{x}^i(t)) \rangle \quad (33)$$

$$\text{(Cauchy-Schwarz)} \leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \sum_{i=1}^p \|x_i(t+1) - x_i(t)\| \cdot \|\nabla_i f(\mathbf{x}(t)) - \nabla_i f(\mathbf{x}^i(t))\| \quad (34)$$

$$\text{(Assumption 1.2)} \leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \sum_{i=1}^p \|x_i(t+1) - x_i(t)\| \cdot L_i \|\mathbf{x}(t) - \mathbf{x}^i(t)\| \quad (35)$$

$$\text{(Equation (27))} \leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \sum_{i=1}^p L_i \|x_i(t+1) - x_i(t)\| \cdot \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \quad (36)$$

$$\text{(Assumption 1.2)} \leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + L \|\mathbf{x}(t+1) - \mathbf{x}(t)\| \cdot \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \quad (37)$$

$$\text{(} ab \leq \frac{a^2+b^2}{2} \text{)} \leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \frac{L}{2} \sum_{k=(t-s)_+}^{t-1} \left[\|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2 + \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 \right] \quad (38)$$

$$\leq \frac{1}{2}(L_f + Ls - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \frac{L}{2} \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2. \quad (39)$$

Summing the above inequality from m to $n-1$ we have

$$F(\mathbf{x}(n)) - F(\mathbf{x}(m)) \leq \frac{1}{2}(L_f + Ls - 1/\eta) \sum_{t=m}^{n-1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + \frac{L}{2} \sum_{t=m}^{n-1} \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2 \quad (40)$$

$$\leq \frac{1}{2}(L_f + 2Ls - 1/\eta) \sum_{t=m}^{n-1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2. \quad (41)$$

Therefore, as long as $\eta < 1/(L_f + 2Ls)$, letting $m = 0$ we deduce

$$\sum_{t=0}^{n-1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 \leq \frac{2}{1/\eta - L_f - 2Ls} [F(\mathbf{x}(0)) - F(\mathbf{x}(n))] \leq \frac{2}{1/\eta - L_f - 2Ls} [F(\mathbf{x}(0)) - \inf_{\mathbf{z}} F(\mathbf{z})]. \quad (42)$$

By Assumption 1.1, F is bounded from below hence the right-hand side is finite and independent of n . Letting n goes to infinity completes the proof of item 1.

Item 2 follows immediately from item 1 and (27), whence it is clear that for all i the limit points satisfy $\omega(\{\mathbf{x}(k)\}) = \omega(\{\mathbf{x}^i(k)\})$.

To prove item 3, let \mathbf{x}^* be a limit point of $\{\mathbf{x}(t)\}_t$, i.e., there exists a subsequence $\mathbf{x}(t_m) \rightarrow \mathbf{x}^*$. Since the objective function F is closed we know $\mathbf{x}^* \in \text{dom } F$. To show $\mathbf{x}^* \in \text{crit } F$ we need to exhibit a sequence say $\mathbf{x}(k_m + 1)$ such that³

$$\mathbf{x}(k_m + 1) \rightarrow \mathbf{x}^*, F(\mathbf{x}(k_m + 1)) \rightarrow F(\mathbf{x}^*), \mathbf{0} \leftarrow \mathbf{u}(k_m + 1) \in \partial F(\mathbf{x}(k_m + 1)). \quad (43)$$

³Technically, from Definition 1 we should have the Frechét subdifferential $\hat{\partial}F$ in Equation (43), however, a usual diagonal argument allows us to use the more convenient (limiting) subdifferential.

This is the most difficult part of the proof, and the argument differs substantially from previous work (e.g. [8]).

We first prove the subdifferential goes to zero. Observe from Assumption 2.3 that the iterations $\{t, t \in T_i\}$ when machine i updates is infinite. Let $\hat{t} \in T_i$ and by the optimality condition of $x_i(\hat{t} + 1)$ in Equation (29):

$$-\frac{1}{\eta} [x_i(\hat{t} + 1) - x_i(\hat{t}) + \eta \nabla_i f(\mathbf{x}^i(\hat{t}))] \in \partial g_i(x_i(\hat{t} + 1)), \quad (44)$$

i.e. there exists $u_i(\hat{t} + 1) \in \partial g_i(x_i(\hat{t} + 1))$ such that

$$\|u_i(\hat{t} + 1) + \nabla_i f(\mathbf{x}(\hat{t} + 1))\| \leq \|u_i(\hat{t} + 1) + \nabla_i f(\mathbf{x}(\hat{t}))\| + \|\nabla_i f(\mathbf{x}(\hat{t} + 1)) - \nabla_i f(\mathbf{x}(\hat{t}))\| \quad (45)$$

$$\begin{aligned} & \text{(Equation (44), Assumption 1.2)} \leq \left\| \frac{1}{\eta} (x_i(\hat{t} + 1) - x_i(\hat{t}) + \nabla_i f(\mathbf{x}^i(\hat{t})) - \nabla_i f(\mathbf{x}(\hat{t}))) \right\| + L_i \|\mathbf{x}(\hat{t} + 1) - \mathbf{x}(\hat{t})\| \\ & \hspace{15em} (46) \end{aligned}$$

$$\text{(triangle inequality, Assumption 1.2)} \leq \frac{1}{\eta} \|x_i(\hat{t} + 1) - x_i(\hat{t})\| + L_i \|\mathbf{x}^i(\hat{t}) - \mathbf{x}(\hat{t})\| + L_i \|\mathbf{x}(\hat{t} + 1) - \mathbf{x}(\hat{t})\| \quad (47)$$

$$\begin{aligned} & \text{(Equation (27))} \leq \frac{1}{\eta} \|x_i(\hat{t} + 1) - x_i(\hat{t})\| + L_i \sum_{k=(\hat{t}-s)_+}^{\hat{t}} \|\mathbf{x}(k+1) - \mathbf{x}(k)\|. \\ & \hspace{15em} (48) \end{aligned}$$

We now use a chaining argument to remove the condition $\hat{t} \in T_i$ above. For each $t \notin T_i$ let \hat{t}_i be the *largest* element in $\{k \leq t : k \in T_i\}$. Thanks to Assumption 2.3 \hat{t}_i always exists and $t - \hat{t}_i \leq s$. Therefore, for any $t \notin T_i$, since $x_i(t+1) = x_i(\hat{t}_i + 1)$ we can certainly choose $u_i(t+1) \in \partial g_i(x_i(t+1))$ to coincide with $u_i(\hat{t}_i + 1) \in \partial g_i(x_i(\hat{t}_i + 1))$. Then:

$$\|u_i(t+1) + \nabla_i f(\mathbf{x}(t+1)) - u_i(\hat{t}_i + 1) - \nabla_i f(\mathbf{x}(\hat{t}_i + 1))\| = \|\nabla_i f(\mathbf{x}(t+1)) - \nabla_i f(\mathbf{x}(\hat{t}_i + 1))\| \quad (49)$$

$$\begin{aligned} & \text{(triangle inequality)} \leq \sum_{k=\hat{t}_i+1}^t \|\nabla_i f(\mathbf{x}(k+1)) - \nabla_i f(\mathbf{x}(k))\| \\ & \hspace{15em} (50) \end{aligned}$$

$$\begin{aligned} & \text{(Assumption 2.3)} \leq \sum_{k=(t-s+1)_+}^t \|\nabla_i f(\mathbf{x}(k+1)) - \nabla_i f(\mathbf{x}(k))\| \\ & \hspace{15em} (51) \end{aligned}$$

$$\begin{aligned} & \text{(Assumption 1.2)} \leq L_i \sum_{k=(t-s+1)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\|. \\ & \hspace{15em} (52) \end{aligned}$$

Combining the two separate cases in Equation (48) and Equation (52) above we have for all t :

$$\|\mathbf{u}(t+1) + \nabla f(\mathbf{x}(t+1))\| \leq (\sqrt{p}/\eta + 2L) \sum_{k=(t-2s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\|, \quad (53)$$

where of course $\mathbf{u}(t+1) = (u_1(t+1), \dots, u_p(t+1)) \in \partial g(\mathbf{x}(t+1))$ and we artificially introduce \sqrt{p} for convenience of subsequent proof. Therefore, from item 2 we deduce

$$\lim_{t \rightarrow \infty} \text{dist}_{\partial F(\mathbf{x}(t+1))}(\mathbf{0}) \rightarrow 0. \quad (54)$$

Next we deal with the function value convergence in Equation (43). For any $\hat{t}_i \in T_i$, using Equation (30) with $z = x_i^*$ we have

$$g_i(x_i(\hat{t}_i + 1)) + \frac{1}{2\eta} \|x_i(\hat{t}_i + 1) - x_i(\hat{t}_i) + \eta \nabla_i f(\mathbf{x}^i(\hat{t}_i))\|^2 \leq g_i(x_i^*) + \frac{1}{2\eta} \|x_i^* - x_i(\hat{t}_i) + \eta \nabla_i f(\mathbf{x}^i(\hat{t}_i))\|^2, \quad (55)$$

which, after rearrangement, yields

$$g_i(x_i(\hat{t}_i + 1)) \leq g_i(x_i^*) + \frac{1}{2\eta} \|x_i^* - x_i(\hat{t}_i)\|^2 - \frac{1}{2\eta} \|x_i(\hat{t}_i + 1) - x_i(\hat{t}_i)\|^2 + \langle x_i^* - x_i(\hat{t}_i + 1), \nabla_i f(\mathbf{x}^i(\hat{t}_i)) \rangle \quad (56)$$

$$\begin{aligned} & = g_i(x_i^*) + \frac{1}{2\eta} \|x_i^* - x_i(\hat{t}_i)\|^2 - \frac{1}{2\eta} \|x_i(\hat{t}_i + 1) - x_i(\hat{t}_i)\|^2 + \langle x_i^* - x_i(\hat{t}_i + 1), \nabla_i f(\mathbf{x}^*) \rangle \\ & \quad + \langle x_i^* - x_i(\hat{t}_i + 1), \nabla_i f(\mathbf{x}^i(\hat{t}_i)) - \nabla_i f(\mathbf{x}^*) \rangle. \\ & \hspace{15em} (57) \end{aligned}$$

We wish to deduce from the above inequality that $g_i(x_i(\hat{t}_i + 1)) \rightarrow g(x_i^*)$, but we need a uniformization device to remove the dependence on i (hence removing the condition $\hat{t}_i \in T_i$). Observing from item 2 that

$$\lim_{m \rightarrow \infty} \max_{t \in [t_m - s, t_m + s]} \|\mathbf{x}(t + 1) - \mathbf{x}^*\| \rightarrow 0, \quad \lim_{m \rightarrow \infty} \max_{t \in [t_m - s, t_m + s]} \|\mathbf{x}^i(t + 1) - \mathbf{x}^*\| \rightarrow 0. \quad (58)$$

By Assumption 2.3, $[t_m - s, t_m + s] \cap T_i \neq \emptyset$ for all i , using item 2 again and the Lipschitz continuity of ∇f , we deduce from Equation (57) that

$$\limsup_{m \rightarrow \infty} \max_{t \in [t_m - s, t_m + s] \cap T_i} g_i(x_i(t + 1)) \leq g_i(x_i^*). \quad (59)$$

Since each machine must update at least once on the intervals $[t_m - s, t_m]$ and $[t_m, t_m + s]$, let \hat{t}_m^i be the largest element of $[t_m - s, t_m] \cap T_i$. Then from the previous inequality we have

$$\limsup_{m \rightarrow \infty} \max_{t \in [\hat{t}_m^i, t_m + s] \cap T_i} g_i(x_i(t + 1)) \leq g_i(x_i^*). \quad (60)$$

Since $g_i(x_i(t + 1)) = g_i(x_i(t))$ if $t \notin T_i$ and $\hat{t}_m^i \in T_i$, it follows that

$$\max_{t \in [t_m, t_m + s]} g_i(x_i(t + 1)) \leq \max_{t \in [\hat{t}_m^i, t_m + s] \cap T_i} g_i(x_i(t + 1)), \quad (61)$$

hence

$$\limsup_{m \rightarrow \infty} \max_{t \in [t_m, t_m + s]} g_i(x_i(t + 1)) \leq g_i(x_i^*). \quad (62)$$

Choose any sequence k_m such that $k_m \in [t_m, t_m + s]$. Since $\mathbf{x}(t_m) \rightarrow \mathbf{x}^*$, from item 2 it is clear that

$$\mathbf{x}(k_m + 1) \rightarrow \mathbf{x}^*. \quad (63)$$

From Equation (62) we know for all i , $\limsup_{m \rightarrow \infty} g_i(x_i(k_m + 1)) \leq g_i(x_i^*)$ while using closedness of the function g_i we have $\liminf_{m \rightarrow \infty} g_i(x_i(k_m + 1)) \geq g_i(x_i^*)$, thus in fact $\lim_{m \rightarrow \infty} g_i(x_i(k_m + 1)) = g_i(x_i^*)$. Since f is continuous, we know

$$\lim_{m \rightarrow \infty} F(\mathbf{x}(k_m + 1)) = F(\mathbf{x}^*). \quad (64)$$

Lastly, combining Equation (54), Equation (63) and Equation (64), it follows from Definition 1 that $\mathbf{x}^* \in \text{crit } F$. \square

B Proof of Theorem 2

Theorem 2 (Finite Length). *Let Assumption 1, 2, 3 and 4 hold, and apply msPG to problem (P). If the step size $\eta < (L_f + 2L_s)^{-1}$ and $\{\mathbf{x}(t)\}$ is bounded, then*

$$\sum_{t=0}^{\infty} \|\mathbf{x}(t + 1) - \mathbf{x}(t)\| < \infty, \quad (11)$$

$$\forall i = 1, \dots, p, \quad \sum_{t=0}^{\infty} \|\mathbf{x}^i(t + 1) - \mathbf{x}^i(t)\| < \infty. \quad (12)$$

Furthermore, $\{\mathbf{x}(t)\}$ and $\{\mathbf{x}^i(t)\}, i = 1, \dots, p$, converge to the same critical point of F .

Our proof requires the following simple uniformization of the KL inequality in Definition 4:

Lemma 2 (Uniformized KL inequality, [11, Lemma 6]). *Let h be a KL function and $\Omega \subset \text{dom } h$ be a compact set. If h is constant on Ω , then there exist $\varepsilon, \lambda > 0$ and a function φ as in Definition 4, such that for all $\bar{\mathbf{x}} \in \Omega$ and all $\mathbf{x} \in \{\mathbf{x} \in \mathbb{R}^d : \text{dist}_{\Omega}(\mathbf{x}) < \varepsilon\} \cap \{\mathbf{x} : h(\bar{\mathbf{x}}) < h(\mathbf{x}) < h(\bar{\mathbf{x}}) + \lambda\}$, one has*

$$\varphi'(h(\mathbf{x}) - h(\bar{\mathbf{x}})) \cdot \text{dist}_{\partial h(\mathbf{x})}(\mathbf{0}) \geq 1.$$

The proof of this lemma is the usual covering argument.

Proof of Theorem 2. We first show that if the global sequence has finite length (i.e. (11)) then the local sequences also have finite length (i.e. (12)). Indeed,

$$\|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\| \leq \|\mathbf{x}^i(t+1) - \mathbf{x}(t+1)\| + \|\mathbf{x}(t+1) - \mathbf{x}(t)\| + \|\mathbf{x}(t) - \mathbf{x}^i(t)\| \quad (65)$$

$$\text{(Equation (27))} \leq \|\mathbf{x}(t+1) - \mathbf{x}(t)\| + \sum_{k=(t+1-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| + \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\|. \quad (66)$$

Therefore, summing from $t = 0$ to $t = n$:

$$\sum_{t=0}^n \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\| \leq \sum_{t=0}^n \left[\sum_{k=(t+1-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| + \sum_{k=(t-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \right] \quad (67)$$

$$\leq (2s+1) \sum_{t=0}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\|. \quad (68)$$

Letting n tend to infinity we have (11) \implies (12).

From Theorem 1 we know the limit points of $\{\mathbf{x}(t)\}$ and $\{\mathbf{x}^i(t)\}, i = 1, \dots, p$, coincide, and they are critical points of F .

The only thing left to prove is the finite length property of the global sequence $\mathbf{x}(t)$. If for all large t we have $\mathbf{x}(t+1) = \mathbf{x}(t)$ then the conclusion is trivial. On the other hand, we can remove all iterations t with $\mathbf{x}(t+1) = \mathbf{x}(t)$, without affecting the length of the trajectory. Thus, in the following we assume for all (large) t we have $\mathbf{x}(t+1) \neq \mathbf{x}(t)$. Thanks to Assumption 3 and Assumption 1.1, it is then clear that the objective value $F(\mathbf{x}(t))$ is strictly decreasing to a limit F^* . Since $\{\mathbf{x}(t)\}$ is assumed to be bounded, the limit point set $\Omega := \omega(\{\mathbf{x}(t)\})$ is nonempty and compact. Obviously for any $\mathbf{x}^* \in \Omega$ we have $F(\mathbf{x}^*) = F^*$. Fix any $\epsilon > 0$, clearly for t sufficiently large we have⁴ $\text{dist}_{\Omega}(\mathbf{x}(t)) \leq \epsilon$. We now have all ingredients to apply the uniformized KL inequality in Lemma 2, which implies that for all sufficiently large t , there exists a continuous and concave function φ (with additional properties listed in Definition 4) such that

$$\varphi'(F(\mathbf{x}(t)) - F^*) \cdot \text{dist}_{\partial F(\mathbf{x}(t))}(\mathbf{0}) \geq 1. \quad (69)$$

Since φ is concave, we obtain

$$\Delta_{t,t+1} := \varphi(F(\mathbf{x}(t)) - F^*) - \varphi(F(\mathbf{x}(t+1)) - F^*) \geq \varphi'(F(\mathbf{x}(t)) - F^*)(F(\mathbf{x}(t)) - F(\mathbf{x}(t+1))) \quad (70)$$

$$\text{(Assumption 3 and Equation (69))} \geq \frac{\alpha \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2}{\text{dist}_{\partial F(\mathbf{x}(t))}(\mathbf{0})}. \quad (71)$$

It is clear that the function φ (composed with F) serves as a Lyapunov function. To proceed, we need to upper bound the subdifferential $\partial F(\mathbf{x}(t))$, which has been painstakingly dealt with in the proof of Theorem 1.

Using the inequality $2\sqrt{ab} \leq a + b$ for positive numbers we obtain from Equation (71): for t sufficiently large,

$$2\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \leq \frac{\delta}{\alpha} \Delta_{t,t+1} + \frac{1}{\delta} \text{dist}_{\partial F(\mathbf{x}(t))}(\mathbf{0}), \quad (72)$$

⁴This is true for any bounded sequence, and we provide a proof for completeness: Suppose not, then there exists $\epsilon > 0$ such that for all n there exists a $t \geq n$ such that $\text{dist}_{\Omega}(\mathbf{x}(t)) > \epsilon$. Thus, we can extract a subsequence $\{\mathbf{x}(t_m)\}$ such that $\text{dist}_{\Omega}(\mathbf{x}(t_m)) > \epsilon$. However, since $\{\mathbf{x}(t)\}$ is bounded, we can extract a further subsequence, say $\{\mathbf{x}(t_{m_n})\}$, that converges, i.e. $\text{dist}_{\Omega}(\mathbf{x}(t_{m_n})) \rightarrow 0$, contradiction.

where $\delta > 0$ will be fixed later. Summing the above inequality over t from m (sufficiently large) to n :

$$2 \sum_{t=m}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\| \leq \sum_{t=m}^n \frac{\delta}{\alpha} \Delta_{t,t+1} + \sum_{t=m}^n \frac{1}{\delta} \text{dist}_{\partial F(\mathbf{x}(t))}(\mathbf{0}) \quad (73)$$

$$\text{(telescoping and Equation (53))} \leq \frac{\delta}{\alpha} \varphi(F(\mathbf{x}(m)) - F^*) + \sum_{t=m}^n \frac{\sqrt{p}/\eta + 2L}{\delta} \sum_{k=(t-2s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \quad (74)$$

$$\begin{aligned} &\leq \frac{\delta}{\alpha} \varphi(F(\mathbf{x}(m)) - F^*) + \frac{(2s+1)(\sqrt{p}/\eta + 2L)}{\delta} \sum_{k=(m-2s)_+}^{m-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \\ &\quad + \frac{(2s+1)(\sqrt{p}/\eta + 2L)}{\delta} \sum_{t=m}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\|. \end{aligned} \quad (75)$$

Setting $\delta = (2s+1)(\sqrt{p}/\eta + 2L)$ and rearranging:

$$\sum_{t=m}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\| \leq \frac{(2s+1)(\sqrt{p}/\eta + 2L)}{\alpha} \varphi(F(\mathbf{x}(m)) - F^*) + \sum_{k=(m-2s)_+}^{m-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \quad (76)$$

Since the right-hand side is finite and does not depend on n , letting n tend to infinity completes our proof for Equation (11). \square

C Proof of Lemma 1

Lemma 1. *Assume $\forall t, i, t \in T_i$. Let the step size $\eta < \frac{\rho-1}{4C\rho} \frac{\sqrt{\rho}-1}{\sqrt{\rho^{s+1}}-1}$ for any $\rho > 1$ and all $U_i, i = 1, \dots, p$ be eventually Lipschitz continuous, then the sequences $\{\mathbf{x}(t)\}$ and $\{\mathbf{x}^i(t)\}, i = 1, \dots, p$, have finite length.*

Follow the same argument of equation (27), one can bound $\|\mathbf{x}^i(t) - \mathbf{x}^i(t+1)\|$ similarly as

$$\|\mathbf{x}^i(t) - \mathbf{x}^i(t+1)\| \leq \sum_{k=(t-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\|. \quad (77)$$

For simplicity we omit the details.

Since Assumption 5 holds for all $t > t_L$, we prove the lemma by considering two complementary cases.

Case 1: There exists a $\hat{t} > t_L$ such that

$$\sum_{k=(\hat{t}-s)_+}^{\hat{t}} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \leq \frac{\sqrt{\rho^{s+1}} - 1}{\sqrt{\rho} - 1} \|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\|$$

Case 2: For all $t > t_L$ case 1 fails.

We will show that case 1 leads to the sufficient decrease property in Assumption 3 for all large t , case 2 leads to the finite length of the models.

Case 1: \hat{t} exists.

We start by proving the following lemma.

Lemma 3. *With Assumption 5 and the existence of \hat{t} . Set $\eta^{-1} > \frac{4C\rho}{\rho-1} \frac{\sqrt{\rho^{s+1}}-1}{\sqrt{\rho}-1}$, then it holds for all $t > \hat{t}$ that*

$$\|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\| \leq \sqrt{\rho} \|\mathbf{x}(\hat{t}+2) - \mathbf{x}(\hat{t}+1)\|.$$

Proof. Using the inequality $\|a\|_2^2 - \|b\|_2^2 \leq 2\|a\|\|a - b\|$, we have for all $t > \hat{t} > t_L$

$$\begin{aligned}
 \|\mathbf{x}(t+1) - \mathbf{x}(t)\|_2^2 - \|\mathbf{x}(t+2) - \mathbf{x}(t+1)\|_2^2 &\leq 2\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \|(\mathbf{x}(t+1) - \mathbf{x}(t)) - (\mathbf{x}(t+2) - \mathbf{x}(t+1))\| \\
 &\text{(no skip of update)} = 2\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \left\| \sum_{i=1}^p U_i(\mathbf{x}^i(t)) - \sum_{i=1}^p U_i(\mathbf{x}^i(t+1)) \right\| \\
 &\leq 2\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \sum_{i=1}^p \|U_i(\mathbf{x}^i(t)) - U_i(\mathbf{x}^i(t+1))\| \\
 &\text{(Assumption 5)} \leq 2\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \left(\sum_{i=1}^p C_i \eta \|\mathbf{x}^i(t) - \mathbf{x}^i(t+1)\| \right) \\
 &\text{(equation (77))} \leq 2\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \left(\sum_{i=1}^p C_i \eta \left[\sum_{k=(t-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \right] \right) \\
 &= 2C\eta \|\mathbf{x}(t+1) - \mathbf{x}(t)\| \left[\sum_{k=(t-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \right]. \tag{78}
 \end{aligned}$$

Now we use an induction argument. Since there exists $\hat{t} > t_L$ such that $\sum_{k=(\hat{t}-s)_+}^{\hat{t}} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \leq \frac{\sqrt{\rho^{s+1}} - 1}{\sqrt{\rho} - 1} \|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\|$, then set $t = \hat{t}$ in the above inequality, we obtain

$$\begin{aligned}
 \|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\|_2^2 - \|\mathbf{x}(\hat{t}+2) - \mathbf{x}(\hat{t}+1)\|_2^2 &\leq 2C\eta \frac{\sqrt{\rho^{s+1}} - 1}{\sqrt{\rho} - 1} \|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\|_2^2 \\
 &\text{(choice of } \eta) \leq \left(1 - \frac{1}{\rho}\right) \|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\|_2^2. \tag{79}
 \end{aligned}$$

After rearranging terms we conclude $\|\mathbf{x}(\hat{t}+1) - \mathbf{x}(\hat{t})\| \leq \sqrt{\rho} \|\mathbf{x}(\hat{t}+2) - \mathbf{x}(\hat{t}+1)\|$. Now we assume this relationship holds up to t ($t > \hat{t}$), then (78) becomes

$$\begin{aligned}
 \|\mathbf{x}(t+1) - \mathbf{x}(t)\|_2^2 - \|\mathbf{x}(t+2) - \mathbf{x}(t+1)\|_2^2 &\leq 2C\eta \frac{\sqrt{\rho^{s+1}} - 1}{\sqrt{\rho} - 1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|_2^2 \\
 &\text{(choice of } \eta) \leq \left(1 - \frac{1}{\rho}\right) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|_2^2
 \end{aligned}$$

we obtain $\|\mathbf{x}(t+1) - \mathbf{x}(t)\| \leq \sqrt{\rho} \|\mathbf{x}(t+2) - \mathbf{x}(t+1)\|$. This completes the lemma. \square

With this bound, inequality (37) can be further bounded for $t > \hat{t}$ as

$$\begin{aligned}
 F(\mathbf{x}(t+1)) - F(\mathbf{x}(t)) &\leq \frac{1}{2}(L_f - 1/\eta) \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2 + L \|\mathbf{x}(t+1) - \mathbf{x}(t)\| \cdot \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\|. \\
 &\leq -\alpha \|\mathbf{x}(t+1) - \mathbf{x}(t)\|^2, \tag{80}
 \end{aligned}$$

where

$$\begin{aligned}
 \alpha &\geq \frac{\eta^{-1} - L_f}{2} - \frac{L\sqrt{\rho}(1 - \sqrt{\rho}^s)}{1 - \sqrt{\rho}} \\
 (C > L, \rho > 1) &\geq \frac{2L\rho}{\rho - 1} \frac{\sqrt{\rho^{s+1}} - 1}{\sqrt{\rho} - 1} - \frac{L(\sqrt{\rho^{s+1}} - 1)}{\sqrt{\rho} - 1} - \frac{L_f}{2} \\
 &\geq \frac{L(\sqrt{\rho^{s+1}} - 1)}{\sqrt{\rho} - 1} \left(\frac{2\rho}{\rho - 1} - 1 \right) - \frac{L_f}{2} \\
 (L > L_f, \rho > 1) &> 0,
 \end{aligned}$$

This proves the sufficient decrease for all $t > \hat{t}$ of the objective value. Hence, the finite length property of the models follows from Theorem 2.

Case 2: \hat{t} does not exist

In this case we have for all $t > t_L$ it holds that $\sum_{k=(t-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \geq \frac{\sqrt{\rho^{s+1}}-1}{\sqrt{\rho}-1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|$. Set $D = \frac{\sqrt{\rho^{s+1}}-1}{\sqrt{\rho}-1}$ and sum the inequality over t from t_L to n yields

$$\begin{aligned} \sum_{k=t_L}^n \|\mathbf{x}(k+1) - \mathbf{x}(k)\| &< \frac{1}{D} \sum_{t=t_L}^n \sum_{k=(t-s)_+}^t \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \\ &< \frac{s+1}{D} \sum_{t=(t_L-s)_+}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\|, \end{aligned}$$

which after rearranging terms becomes

$$\left(1 - \frac{s+1}{D}\right) \sum_{t=t_L}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\| \leq \frac{s+1}{D} \sum_{t=(t_L-s)_+}^{t_L-1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|.$$

Since $D = \frac{\sqrt{\rho^{s+1}}-1}{\sqrt{\rho}-1} > s+1$ for $\rho > 1$ and t_L is finite, the right hand side of the above inequality is finite, and the left hand side has positive coefficient. Thus the above inequality implies

$$\sum_{t=0}^n \|\mathbf{x}(t+1) - \mathbf{x}(t)\| < +\infty.$$

Enlarge $n \rightarrow \infty$ gives the finite length property of the global model. By the proof of Appendix B, we know the finite length of global model implies the finite length of all local models.

D Proof of Example 1

We proof case by case, and the scaled version $\gamma g(\mathbf{x}), \gamma > 0$ trivially follows from the same argument.

Cases $g = 0, g = \frac{1}{2} \|\cdot\|^2$:

When $g = 0$, the update operator in eq. (9) becomes $U_i(\mathbf{x}^i(t)) = -\eta \nabla_i f(\mathbf{x}^i(t))$, which is ηL_i Lipschitz due to Assumption 1.2 .

When $g = \frac{1}{2} \|\cdot\|^2$, the update operator becomes for $i = 1, \dots, p$

$$U_i(\mathbf{x}^i(t)) = \text{prox}_{\frac{\eta}{2} \|\cdot\|_2^2}^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))) - x_i(t) = -\frac{1}{1 + \eta^{-1}} (x_i(t) + \nabla_i f(\mathbf{x}^i(t))).$$

With which we have

$$\begin{aligned} \|U_i(\mathbf{x}^i(t+1)) - U_i(\mathbf{x}^i(t))\| &\leq \frac{1}{1 + \eta^{-1}} \|(x_i(t+1) - x_i(t)) + (\nabla_i f(\mathbf{x}^i(t+1)) - \nabla_i f(\mathbf{x}^i(t)))\| \\ &\leq \eta(1 + L_i) \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\| \end{aligned}$$

Cases $g = \|\cdot\|_0, \|\cdot\|_0 + \|\cdot\|^2, \|\cdot\|_{0,2}, \|\cdot\|_{0,2} + \|\cdot\|^2$: For the non-overlapping group norms, we assign each machine a subset of groups of coordinates.

Consider $g = \|\cdot\|_0$, its proximal map on i -th coordinate can be expressed as

$$\text{prox}_{g_i}^\eta(z_i) = \begin{cases} z_i, & \text{if } |z_i| > \sqrt{2\eta} \\ 0, & \text{otherwise} \end{cases}$$

The mapping contains a hard threshold, i.e., it filters out those coordinates with magnitude less than $\sqrt{2\eta}$. This implies that any change of the support set of $\text{prox}_{g_i}^\eta(z_i)$ will induce a jump of magnitude of at least $\sqrt{2\eta}$. On

the other hand, the second assertion of Theorem 1 imply that $\lim_{t \rightarrow \infty} \|x_i(t+2) - x_i(t+1)\| = 0$, which by local update can be expressed as

$$\lim_{t \rightarrow \infty} \|\text{prox}_{g_i}^\eta(x_i(t+1) - \eta \nabla_i f(\mathbf{x}^i(t+1))) - \text{prox}_{g_i}^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t)))\| = 0.$$

Hence by the above equation and the jump of proximal map, the support Ω of $\text{prox}_{g_i}^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t)))$ (i.e., $x_i(t+1)$) must remain stable for all t sufficiently large. Moreover, the proximal map reduces to identity operator on the support set Ω . Thus, for all t sufficiently large we have

$$\begin{aligned} \|U_i(\mathbf{x}^i(t+1)) - U_i(\mathbf{x}^i(t))\| &= \|\text{prox}_{g_i}^\eta(x_i(t+1) - \eta \nabla_i f(\mathbf{x}^i(t+1))) - x_i(t+1) - \text{prox}_{g_i}^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))) - x_i(t)\| \\ &\quad (\text{support on } \Omega) = \|\text{prox}_{g_i}^\eta(x_i(t+1) - \eta \nabla_i f(\mathbf{x}^i(t+1))) - x_i(t+1) - \text{prox}_{g_i}^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))) - x_i(t)\|_\Omega \\ &\quad (\text{prox}_g^\eta \text{ is identity on } \Omega) \leq \|\eta \nabla_i f(\mathbf{x}^i(t)) - \eta \nabla_i f(\mathbf{x}^i(t+1))\| \\ &\quad \leq \eta L_i \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\|. \end{aligned}$$

Hence the operator is eventually $\mathcal{O}(\eta)$ Lipschitz.

Next we consider (without loss of generality) $g = \|\cdot\|_0 + \frac{\lambda}{2} \|\cdot\|^2$ where $\lambda > 0$. The proximal map on i -th coordinate is

$$\text{prox}_{g_i}^\eta(z_i) = \begin{cases} z_i, & \text{if } |z_i| > \sqrt{2(\eta + \eta^2 \lambda)} \\ 0, & \text{otherwise} \end{cases}$$

Thus, the mapping also contains a hard threshold. Following similar argument as previous case, we conclude that the support Ω of $\text{prox}_{g_i}^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t)))$ (i.e., $x_i(t+1)$) must remain stable for all t sufficiently large, and the proximal map reduces to identity operator on Ω . Consequently, the operator $U_i(\mathbf{x}^i(t))$ is $\mathcal{O}(\eta)$ Lipschitz for all t large.

The proof of group norms $g = \|\cdot\|_{0,2}, \|\cdot\|_{0,2} + \|\cdot\|^2$ then follows by realizing that the proximal maps have hard threshold on group support.

Cases $g = \|\cdot\|_1, \|\cdot\|_1 + \|\cdot\|^2$ with eventual stable support set of $\{\mathbf{x}(t)\}$:

For these two cases we assume that the support set of $\{\mathbf{x}(t)\}$ remains unchanged for all large t .

We just need to consider $g = \|\cdot\|_1 + \frac{\lambda}{2} \|\cdot\|^2, \lambda \geq 0$. Its proximal map on vector z_i has the form

$$\text{prox}_g^\eta(z_i) = \frac{1}{1+\eta\lambda} \text{sgn}(z_i) (|z_i| - \eta)_+.$$

Since the support set Ω of $\mathbf{x}(t)$ (i.e. $\text{prox}_g^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t)))$) is assumed to be stable after some t_L , the above soft-thresholding operator ensures that $|x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))|_\Omega > \eta$ for all large t , and we obtain

$$\begin{aligned} U_i(\mathbf{x}^i(t)) &= [x_i(t+1) - x_i(t)]_\Omega = [\text{prox}_g^\eta(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))) - x_i(t)]_\Omega \\ &= (1 + \eta\lambda)^{-1} [-\eta \nabla_i f(\mathbf{x}^i(t)) - \eta \text{sgn}(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t)))]_\Omega - \frac{\eta\lambda}{1 + \eta\lambda} [x_i(t)]_\Omega \end{aligned} \quad (81)$$

On the other hand, Theorem 1.2 and the Lipschitz gradient of f implies

$$\lim_{t \rightarrow \infty} \|[x_i(t+1) - \eta \nabla_i f(\mathbf{x}^i(t+1))] - [x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))]\| = 0.$$

Then $[\text{sgn}(x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t)))]_\Omega$ must eventually remain constant, since otherwise the condition $|x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))|_\Omega > \eta$ will induce a change of $|x_i(t) - \eta \nabla_i f(\mathbf{x}^i(t))|$ to be at least 2η and violate the above asymptotic condition. In summary, for all large t we have

$$U_i(\mathbf{x}^i(t)) = (1 + \eta\lambda)^{-1} [-\eta \nabla_i f(\mathbf{x}^i(t)) - \text{Const}]_\Omega - \frac{\eta\lambda}{1 + \eta\lambda} [x_i(t)]_\Omega$$

which further implies that

$$\begin{aligned} \|U_i(\mathbf{x}^i(t+1)) - U_i(\mathbf{x}^i(t))\| &\leq \|(1 + \eta\lambda)^{-1} [\eta \nabla_i f(\mathbf{x}^i(t+1)) - \eta \nabla_i f(\mathbf{x}^i(t))]\|_\Omega + \frac{\eta\lambda}{1 + \eta\lambda} \|x_i(t+1) - x_i(t)\|_\Omega \\ &\leq \eta L_i \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\| + \eta\lambda \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\| \\ &\leq \eta(L_i + \lambda) \|\mathbf{x}^i(t+1) - \mathbf{x}^i(t)\|. \end{aligned}$$

E Proof of Theorem 3

Theorem 3 (Global rate of convergence). *If the finite length property in Theorem 2 holds, then*

1. $\sum_{t=0}^{\infty} \|\mathbf{e}(t)\| < \infty$;
2. $F(\frac{1}{t} \sum_{k=1}^t \mathbf{x}(k)) - \inf F \leq O(t^{-1})$.

Proof. For the first assertion, note that for any n :

$$\sum_{t=0}^n \|\mathbf{e}(t)\| = \eta \sum_{t=0}^n \left\| (\nabla_1 f(\mathbf{x}^1(t)) - \nabla_1 f(\mathbf{x}(t)), \dots, \nabla_p f(\mathbf{x}^p(t)) - \nabla_p f(\mathbf{x}(t))) \right\| \quad (82)$$

$$\text{(triangle inequality, Assumption 1.2)} \leq \eta \sum_{t=0}^n \sum_{i=1}^p L_i \|\mathbf{x}(t) - \mathbf{x}^i(t)\| \quad (83)$$

$$\begin{aligned} \text{(Equation (27))} &\leq \eta \sum_{t=0}^n \left(\sum_{i=1}^p L_i \right) \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \\ &= L\eta \sum_{t=0}^n \sum_{k=(t-s)_+}^{t-1} \|\mathbf{x}(k+1) - \mathbf{x}(k)\| \end{aligned} \quad (84)$$

$$\leq Ls\eta \sum_{t=0}^{n-1} \|\mathbf{x}(t+1) - \mathbf{x}(t)\|. \quad (85)$$

Letting n tend to infinity we obtain

$$\sum_{t=0}^{\infty} \|\mathbf{e}(t)\| \leq Ls\eta \sum_{t=0}^{\infty} \|\mathbf{x}(t+1) - \mathbf{x}(t)\| < \infty. \quad (86)$$

For the second assertion, we first recall from [30] that the inexact proximal gradient algorithm in Equation (14) has the following bound, provided that F is convex:

$$F\left(\frac{1}{t} \sum_{k=1}^t \mathbf{x}(k)\right) - F^* \leq \frac{(\|\mathbf{x}(0) - \mathbf{x}^*\| + 2A_t)^2}{2t\eta}, \quad \text{where } A_t = \sum_{k=0}^t \eta \|\mathbf{e}(k)\|. \quad (87)$$

The second assertion thus follows from the first one (assuming convexity). \square

F Experiments Specifications

Specifications for $\|\cdot\|_{0,2}$ Lasso:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \sum_{i=1}^{20} \gamma_i \mathbb{I}(\|x_i\|).$$

Here $A \in \mathbb{R}^{1000 \times 2000}$, $\mathbf{b} \in \mathbb{R}^{1000}$, and $\mathbf{x} \in \mathbb{R}^{2000}$ is divided into 20 equal groups of features. Matrix A is generated from $\mathcal{N}(0, 1)$ with normalized columns. We set $\mathbf{b} = A\tilde{\mathbf{x}} + \varepsilon$, where ε is generated from $\mathcal{N}(0, 10^{-2})$ and $\tilde{\mathbf{x}}$ is a normalized vector with 8 non-zero groups of features generated from $\mathcal{N}(0, 1)$. For the non-zero groups of $\tilde{\mathbf{x}}$, we set the corresponding $\gamma_i = 10^{-4}$, and for the remaining groups we set $\gamma_i = 10^{-2}$.

We implement msPG on four cores with each core assigned five group of features. Each core stores the corresponding column blocks of A .

Specifications for $\|\cdot\|_1$ Lasso:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \gamma \|\mathbf{x}\|_1.$$

Data Generation

We generate the data column-wise. Starting from first column, we randomly pick 10^4 samples to have non-zero in column 1 and sample each value from $\text{Uniform}(-1, 1)$. We normalize it such that the ℓ_2 -norm of the column is 1. We denote these values as $\mathbf{v}_1 \in \mathbb{R}^n$. To generate column i , with probability 0.5 we randomly pick a new set of samples to have non-zero values at column i (otherwise we use the same samples from column $i - 1$). This simulates the correlations between each column. Once the samples are chosen, we assign values from $\text{Unif}(-1, 1)$. \mathbf{v}_i is again normalized. We generate ground truth regressor $\beta \in \mathbb{R}^d$ from $\mathcal{N}(0, 1)$ with 1% non-zero entries, and obtain the regressed value from $\mathbf{b} = A\beta$ where A is the design matrix.