

Asynchronous Iterations in Optimization: New Sequence Results and Sharper Algorithmic Guarantees

Hamid Reza Feyzmahdavian

*ABB Corporate Research
Västerås, Sweden*

HAMID.FEYZMAHDAVIAN@SE.ABB.COM

Mikael Johansson

*Division of Decision and Control Systems
KTH—Royal Institute of Technology
Stockholm, Sweden*

MIKAELJ@KTH.SE

Editor: Martin Jaggi

Abstract

We introduce novel convergence results for asynchronous iterations that appear in the analysis of parallel and distributed optimization algorithms. The results are simple to apply and give explicit estimates for how the degree of asynchrony impacts the convergence rates of the iterates. Our results shorten, streamline and strengthen existing convergence proofs for several asynchronous optimization methods and allow us to establish convergence guarantees for popular algorithms that were thus far lacking a complete theoretical understanding. Specifically, we use our results to derive better iteration complexity bounds for proximal incremental aggregated gradient methods, to obtain tighter guarantees depending on the average rather than maximum delay for the asynchronous stochastic gradient descent method, to provide less conservative analyses of the speedup conditions for asynchronous block-coordinate implementations of Krasnosel'skiĭ–Mann iterations, and to quantify the convergence rates for totally asynchronous iterations under various assumptions on communication delays and update rates.

Keywords: asynchronous algorithms, parallel methods, incremental methods, coordinate descent, stochastic gradient descent

1. Introduction

With the ubiquitous digitalization of the society, decision problems are rapidly expanding in size and scope. Increasingly often, we face problems where data, computations, and decisions need to be distributed on multiple nodes. These nodes may be individual cores in a CPU, different processors in a multi-CPU platform, or servers in a geographically dispersed cluster. Representative examples include machine learning on data sets that are too large to conveniently store in a single computer, real-time decision-making based on high-velocity data streams, and control and coordination of infrastructure-scale systems.

Insisting that such multi-node systems operate synchronously limits their scalability, since the performance is then dictated by the slowest node, and the system becomes fragile to node failures. Hence, there is a strong current interest in developing asynchronous algorithms for optimal decision-making (see, e.g., Niu et al. 2011; Liu et al. 2014; Peng et al. 2016; Leblond et al. 2017; Mishchenko et al. 2020 and references therein). Well-established mod-

els for parallel computations, such as bulk synchronous parallel (Valiant, 1990) or MapReduce (Dean and Ghemawat, 2008), are now being complemented by stale-synchronous parallel models (Ho et al., 2013) and fully asynchronous processing paradigms (Han and Daudjee, 2015). In many of these frameworks, the amount of asynchrony is a design parameter: in some systems, the delay is proportional to the number of parallel workers deployed (Niu et al., 2011); while other systems use communication primitives which enforce a hard limit on the maximum information delay (Ho et al., 2013). It is therefore useful to have theoretical results which characterize the level of asynchrony that can be tolerated by a given algorithm. To this end, this paper develops several theoretical tools for studying the convergence of asynchronous iterations.

The dynamics of asynchronous iterations are much richer than their synchronous counterparts, and quantifying the impact of asynchrony on the convergence rate is mathematically challenging. Some of the first results on the convergence of asynchronous iterations were derived by Chazan and Miranker (1969) for solving linear equations. This work was later extended to nonlinear iterations involving maximum norm contractions (Baudet, 1978) and monotone mappings (Bertsekas and Baz, 1987). Powerful convergence results for broad classes of asynchronous iterations under different assumptions on communication delays and update rates were presented by Bertsekas (1983), Tsitsiklis et al. (1986), and in the celebrated book of Bertsekas and Tsitsiklis (2015). Although the framework for modeling asynchronous iterations in Bertsekas and Tsitsiklis (2015) is both powerful and elegant, the most concrete results only guarantee asymptotic convergence and do not give explicit bounds on convergence times. Execution time guarantees are essential when iterative algorithms are used to find a decision under stringent real-time constraints. In this paper, we derive a number of convergence results for asynchronous iterations which explicitly quantify the impact of asynchrony on the convergence times of the iterates.

The convergence guarantees for influential asynchronous optimization algorithms such as HOGWILD! (Niu et al., 2011), Delayed SGD (Agarwal and Duchi, 2011), ASYSCD (Liu et al., 2014), ARock (Peng et al., 2016) and ASAGA (Leblond et al., 2017) have been established on a per-algorithm basis, and are often based on intricate induction proofs. Such proofs tend to be long, and sources of conservatism are hard to isolate. A closer analysis of these proofs reveals that they rely on a few common principles. In this paper, we attempt to unify these ideas, derive general convergence results for the associated sequences, and use these results to systematically provide stronger guarantees for several popular asynchronous algorithms. In contrast to the recent analysis framework proposed in Mania et al. (2017), which models the effect of asynchrony as noise, our results attempt to capture the inherent structure in the asynchronous iterations. This allows us to derive convergence results for complex optimization algorithms in a systematic and transparent manner, without introducing unnecessary conservatism. We make the following specific contributions:

- We identify two important families of sequences, characterized by certain inequalities, that appear naturally when analyzing the convergence of asynchronous optimization algorithms. For each family, we derive convergence results that allow to quantify how the degree of asynchrony affects the convergence rate guarantees. We use these sequence results to analyze several popular asynchronous optimization algorithms.

- First, we derive stronger convergence guarantees for the proximal incremental gradient method and provide a larger range of admissible step-sizes. Specifically, for L -smooth and convex objective functions, we prove an iteration complexity of $\mathcal{O}(L\tau/\epsilon)$, which improves upon the previously known rate $\mathcal{O}(L^2\tau^3/\epsilon)$ given in Sun et al. (2019). We also show that for objective functions that satisfy a quadratic functional growth condition, the iteration complexity is $\mathcal{O}(Q\tau \log(1/\epsilon))$, where $Q = L/\mu$ is the condition number. In this case, our result allows the algorithm to use larger step-sizes than those provided in Vanli et al. (2018), leading to a tighter convergence rate guarantee.
- Second, we analyze the asynchronous stochastic gradient descent method with delay-dependent step-sizes and extend the results of Koloskova et al. (2022) from non-convex to convex and strongly convex problems. We show that our sequence results are not limited to providing step-size rules and convergence rates that depend on the maximal delay. In particular, for convex problems, we derive an iteration complexity of

$$\mathcal{O}\left(\frac{L\tau_{\text{ave}}}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right),$$

where τ_{ave} is the average delay, and σ denotes the variance of stochastic gradients. For strongly convex problems, we obtain the iteration complexity

$$\mathcal{O}\left(\left(Q\tau_{\text{ave}} + \frac{\sigma^2}{\epsilon}\right) \log\left(\frac{1}{\epsilon}\right)\right).$$

Our guarantees improve the previously best known bounds given in Arjevani et al. (2020) and Stich and Karimireddy (2020), which are based on the maximal delay that can be significantly larger than the average delay. Similar to Mishchenko et al. (2022), we also provide convergence guarantees and admissible step-sizes which depend only on the number of parallel workers deployed, rather than on the gradient delays.

- Third, we give an improved analysis of the ARock framework for asynchronous block-coordinate updates of Krasnosel’skiĭ–Mann iterations. For pseudo-contractive operators, we show that ARock achieves near-linear speedup as long as the number of parallel computing elements is $o(m)$, where m is the number of decision variables. Compared to the results presented in Hannah and Yin (2017), we improve the requirement for the linear speedup property from $\mathcal{O}(\sqrt{m})$ to $o(m)$.
- Finally, we present a uniform treatment of asynchronous iterations involving block-maximum norm contractions under partial and total asynchronism. Contrary to the results in Bertsekas and Tsitsiklis (2015) which only established asymptotic convergence, we give explicit estimates of the convergence rate for various classes of bounded and unbounded communication delays and update intervals.

This paper generalizes and streamlines our earlier work (Feyzmahdavian and Johansson, 2014; Feyzmahdavian et al., 2014; Aytekin et al., 2016). Specifically, we extend the sequence result in Feyzmahdavian et al. (2014) to a family of unbounded delays, which allows us to deal with totally asynchronous iterations. Compared to Feyzmahdavian et al. (2014), we present two new Lemmas 2 and 3 in Section 3. The analysis in Feyzmahdavian

and Johansson (2014) is limited to contraction mappings in the maximum-norm, and does not provide any sequence results for analyzing asynchronous algorithms. However, Lemma 3 in Section 3 recovers the results in Feyzmahdavian and Johansson (2014) as a special case. The sequence result presented in Aytakin et al. (2016) is only applicable for deriving linear convergence rates, and is restricted to deliver step-size rules and convergence rates that depend on the maximum delay. To overcome these limitations, we introduce a novel sequence result in Lemma 5 that guarantees both linear and sublinear rates of convergence and can provide convergence bounds that depend on the average delay. In Aytakin et al. (2016), the analysis of the proximal incremental gradient method has a drawback that the guaranteed bound grows quadratically with the maximum delay τ . In our work, we improve the dependence on τ from quadratic to linear.

1.1 Notation and Preliminaries

Here, we introduce the notation and review the key definitions that will be used throughout the paper. We let \mathbb{R} , \mathbb{N} , and \mathbb{N}_0 denote the set of real numbers, natural numbers, and the set of natural numbers including zero, respectively. For any $n \in \mathbb{N}$,

$$[n] := \{1, \dots, n\}.$$

For a real number a , we denote the largest integer less than or equal to a by $\lfloor a \rfloor$ and define

$$(a)_+ := \max\{a, 0\}.$$

We use $\|\cdot\|$ to represent the standard Euclidean norm on \mathbb{R}^d and $\langle x, y \rangle$ to denote the Euclidean (dot) inner product of two vectors $x, y \in \mathbb{R}^d$. We say that a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth if it is differentiable and

$$\|\nabla f(y) - \nabla f(x)\| \leq L\|y - x\|, \quad \forall x, y \in \mathbb{R}^d.$$

We say a convex function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex if

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\mu}{2}\|y - x\|^2, \quad \forall x, y \in \mathbb{R}^d.$$

The notation $g(t) = \mathcal{O}(h(t))$ means that there exist positive constants M and t_0 such that $g(t) \leq Mh(t)$ for all $t \geq t_0$, while $g(t) = o(h(t))$ means that $\lim_{t \rightarrow \infty} g(t)/h(t) = 0$. Following standard convention, we use tilde $\tilde{\mathcal{O}}$ -notation to hide poly-logarithmic factors in the problem parameters.

2. Lyapunov Analysis for Optimization Algorithms

Convergence proofs for optimization algorithms are usually based on induction and often presented without invoking any general theorems. For more complex algorithms, this leads to lengthy derivations where it is difficult to distinguish mathematical innovations. The need to systemize convergence proofs for optimization algorithms was recognized in Polyak's insightful textbook (Polyak, 1987). Polyak argued that most results concerning convergence and rate of convergence of optimization algorithms can be derived using Lyapunov's second

method, with typical Lyapunov functions being the objective function value, the norm of its gradient, or the squared distance between the current iterate and the optimal set. In addition, he derived and collected a number of useful sequence results which allowed to shorten, unify and clarify many convergence proofs (Polyak, 1987, Chapter 2).

To make these ideas more concrete, consider the simple gradient descent method:

$$x_{k+1} = x_k - \gamma \nabla f(x_k), \quad k \in \mathbb{N}_0.$$

Assume that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth. If x^* is the minimizer of f on \mathbb{R}^d , the standard convergence proof (e.g., Theorem 2.1.15 in Nesterov 2013) establishes that the iterates satisfy

$$\|x_{k+1} - x^*\|^2 \leq \left(1 - \frac{2\gamma\mu L}{\mu + L}\right) \|x_k - x^*\|^2 - \gamma \left(\frac{2}{\mu + L} - \gamma\right) \|\nabla f(x_k)\|^2.$$

In terms of $V_k = \|x_k - x^*\|^2$ and $W_k = \|\nabla f(x_k)\|^2$, the inequality reads

$$V_{k+1} \leq \left(1 - \frac{2\gamma\mu L}{\mu + L}\right) V_k - \gamma \left(\frac{2}{\mu + L} - \gamma\right) W_k.$$

For step-sizes $\gamma \in (0, 2/(\mu + L)]$, the second term on the right-hand side can be dropped and, hence, V_k is guaranteed to decay by at least a factor $q = 1 - 2\gamma\mu L/(\mu + L)$. That is,

$$V_{k+1} \leq qV_k, \quad k \in \mathbb{N}_0. \quad (1)$$

This implies linear convergence of the iterates, i.e.,

$$V_k \leq q^k V_0, \quad k \in \mathbb{N}_0.$$

When f is convex, but not necessarily strongly convex, and L -smooth, the iterates satisfy

$$2\gamma(f(x_k) - f(x^*)) + \|x_{k+1} - x^*\|^2 \leq \|x_k - x^*\|^2 - \gamma \left(\frac{1}{L} - \gamma\right) \|\nabla f(x_k)\|^2.$$

Let V_k and W_k be as before, while $X_k = 2\gamma(f(x_k) - f(x^*))$. Then, the inequality above can be rewritten as

$$X_k + V_{k+1} \leq V_k - \gamma \left(\frac{1}{L} - \gamma\right) W_k. \quad (2)$$

If $\gamma \in (0, 1/L]$, the second term on the right-hand side is non-positive and can be dropped. Summing both sides of (2) over k and using telescoping cancellation then gives

$$\sum_{k=0}^K X_k + V_{K+1} \leq V_0, \quad K \in \mathbb{N}_0.$$

The well-known $\mathcal{O}(1/k)$ convergence rate of $f(x_k)$ to $f(x^*)$ follows from the fact that V_k is non-negative and that X_k is non-increasing (Beck, 2017, Theorem 10.21).

In the analysis of gradient descent above, it is natural to view the iteration index k as a surrogate for the accumulated execution time. This interpretation is valid when the computations are done on a single computer and the time required to perform gradient computations and iterate updates are constant across iterations. In distributed and asynchronous systems, on the other hand, computations are performed in parallel on nodes with different computational capabilities and workloads, without any global synchronization to a common clock. In these systems, the iteration index is simply an ordering of events, typically associated with reading or writing the iterate vector from memory. Since the global memory may have been updated from the time that it was read by a node until the time that the node returns its result, x_{k+1} may not depend on x_k but rather on some earlier iterate $x_{k-\tau_k}$, where $\tau_k \in \{0, \dots, k-1\}$. More generally, many asynchronous optimization algorithms result in iterations of the form

$$x_{k+1} = \mathcal{M}(x_k, x_{k-1}, \dots, x_{k-\tau_k}).$$

We call these *asynchronous iterations*. The information delay τ_k represents the age (in terms of event count) of the oldest information that the method uses in the current update, and can be seen as a measure of the amount of asynchrony in the system. When we analyze the convergence of asynchronous iterations, $\|x_{k+1} - x^*\|^2$ will not only depend on $\|x_k - x^*\|^2$, but also on $\|x_{k-1} - x^*\|^2, \dots, \|x_{k-\tau_k} - x^*\|^2$, and sometimes also on $\|\nabla f(x_{k-1})\|^2, \dots, \|\nabla f(x_{k-\tau_k})\|^2$. Hence, it will not be enough to consider simple sequence relationships such as (1) or (2). Because of asynchrony, the right-hand side of the inequalities will involve delayed versions of V_k and W_k that perturb the convergence of the synchronous iteration.

To fix ideas, consider the gradient descent method with a constant delay of τ in the gradient computation:

$$x_{k+1} = x_k - \gamma \nabla f(x_{k-\tau}). \quad (3)$$

To analyze its convergence when f is μ -strongly convex and L -smooth, we add and subtract $\gamma \nabla f(x_k)$ to the right-hand side of (3) and study the distance of the iterates to the optimum. By applying the same analysis as for the gradient descent method with $\gamma \in (0, 2/(\mu + L)]$, we obtain

$$\|x_{k+1} - x^*\|^2 \leq \left(1 - \frac{2\gamma\mu L}{\mu + L}\right) \|x_k - x^*\|^2 + \omega_k, \quad (4)$$

where the perturbation term ω_k accounts for the impact of the delay in gradient computation and is given by

$$\omega_k = 2\gamma \langle \nabla f(x_k) - \nabla f(x_{k-\tau}), x_k - \gamma \nabla f(x_k) - x^* \rangle + \gamma^2 \|\nabla f(x_k) - \nabla f(x_{k-\tau})\|^2. \quad (5)$$

According to Lemma 27 in Appendix A, we can bound ω_k by

$$\omega_k \leq (\gamma^4 L^4 \tau^2 + 2\gamma^2 L^2 \tau) \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|^2\}. \quad (6)$$

Substituting (6) into (4), we conclude that

$$\|x_{k+1} - x^*\|^2 \leq \left(1 - \frac{2\gamma\mu L}{\mu + L}\right) \|x_k - x^*\|^2 + (\gamma^4 L^4 \tau^2 + 2\gamma^2 L^2 \tau) \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|^2\}.$$

In terms of $V_k = \|x_k - x^*\|^2$, the iterates therefore satisfy

$$V_{k+1} \leq \left(1 - \frac{2\gamma\mu L}{\mu + L}\right) V_k + (\gamma^4 L^4 \tau^2 + 2\gamma^2 L^2 \tau) \max_{(k-2\tau)_+ \leq \ell \leq k} V_\ell. \quad (7)$$

We see that the inequality (7) includes delayed versions of V_k on the right-hand side and reduces to (1) when $\tau = 0$.

When f is L -smooth and convex, but not strongly convex, we first note that

$$\|x_{k+1} - x^*\|^2 = \|x_k - x^*\|^2 - 2\gamma \langle \nabla f(x_{k-\tau}), x_k - x^* \rangle + \gamma^2 \|\nabla f(x_{k-\tau})\|^2. \quad (8)$$

By Lemma 28 in Appendix A, the inner product is lower bounded by

$$f(x_k) - f^* + \frac{1}{2L} \|\nabla f(x_{k-\tau})\|^2 - \frac{\gamma^2 L \tau}{2} \sum_{\ell=(k-\tau)_+}^{k-1} \|\nabla f(x_{\ell-\tau})\|^2 \leq \langle \nabla f(x_{k-\tau}), x_k - x^* \rangle.$$

Substituting this bound into (8) yields

$$\begin{aligned} 2\gamma(f(x_k) - f^*) + \|x_{k+1} - x^*\|^2 &\leq \|x_k - x^*\|^2 + \gamma^3 L \tau \sum_{\ell=(k-\tau)_+}^{k-1} \|\nabla f(x_{\ell-\tau})\|^2 \\ &\quad - \gamma \left(\frac{1}{L} - \gamma\right) \|\nabla f(x_{k-\tau})\|^2. \end{aligned}$$

With $X_k = 2\gamma(f(x_k) - f^*)$, $V_k = \|x_k - x^*\|^2$, and $W_k = \|\nabla f(x_{k-\tau})\|^2$, the iterates hence satisfy a relationship on the form

$$X_k + V_{k+1} \leq V_k + \gamma^3 L \tau \sum_{\ell=(k-\tau)_+}^{k-1} W_\ell - \gamma \left(\frac{1}{L} - \gamma\right) W_k. \quad (9)$$

Comparing with (2), the right-hand side involves delayed versions of W_k .

In the next section, we study the convergence of sequences which include (7) and (9) as special cases. Our first set of results considers sequence relationships on the form

$$V_{k+1} \leq qV_k + p \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0. \quad (10)$$

Here, the perturbation caused by asynchrony at iteration k is modeled as a function on the order of V_ℓ scaled by a factor p , where $\ell \in [k - \tau_k, k]$ and τ_k is the age of the outdated information. Such sequences have appeared, for example, in the analysis of incremental aggregated gradient methods (Gürbüzbalaban et al., 2017), accelerated incremental aggregated gradient methods with curvature information (Wai et al., 2020), asynchronous quasi-Newton methods (Eisen et al., 2017), and asynchronous forward-backward methods for solving monotone inclusion problems (Stathopoulos and Jones, 2019).

The second set of our results considers iterate relationships on the form

$$X_k + V_{k+1} \leq q_k V_k + p_k \sum_{\ell=(k-\tau_k)_+}^k W_\ell - r_k W_k + e_k, \quad k \in \mathbb{N}_0. \quad (11)$$

Here, the perturbation due to asynchrony does not introduce delayed V_k -terms, but manifests itself through the presence of delayed W_k -terms instead. As we will show in this paper, these relationships appear naturally in the analysis of the proximal incremental aggregated gradient method (Aytekin et al., 2016), the asynchronous stochastic gradient descent method (Agarwal and Duchi, 2011), and the asynchronous Krasnosel'skii–Mann method for pseudo-contractive operators (Peng et al., 2016).

3. Novel Sequence Results for Asynchronous Iterations

In this section, we develop specific convergence results for iterations on the form (10) and (11). Our results attempt to balance simplicity, applicability and power, and provide explicit bounds on how the amount of asynchrony affects the guaranteed convergence rates. As we will demonstrate later, the results allow for a simplified and uniform treatment of several asynchronous optimization algorithms.

3.1 Results for Iterations on the form (10)

Our first result, introduced in Feyzmahdavian et al. (2014), establishes convergence properties of iterations of the form (10) when delays are bounded.

Lemma 1 *Let $\{V_k\}$ be a non-negative sequence satisfying*

$$V_{k+1} \leq qV_k + p \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0, \quad (12)$$

for non-negative constants q and p . Suppose there is a non-negative integer τ such that

$$0 \leq \tau_k \leq \tau, \quad k \in \mathbb{N}_0.$$

If $q + p < 1$, then

$$V_k \leq \rho^k V_0, \quad k \in \mathbb{N}_0,$$

where $\rho = (q + p)^{\frac{1}{1+\tau}}$.

Proof. See Lemma 3 in Feyzmahdavian et al. (2014). ■

Consider the delay-free counterpart of (12):

$$V_{k+1} \leq (q + p)V_k, \quad k \in \mathbb{N}_0.$$

Clearly, if $q + p < 1$, the sequence $\{V_k\}$ converges linearly at a rate of $\rho = q + p$. Lemma 1 shows that the convergence rate of $\{V_k\}$ is still linear in the presence of bounded delays. Lemma 1 also gives an explicit bound on the impact that an increasing delay has on the convergence rate. As can be expected, the guaranteed convergence rate deteriorates with increasing τ . More precisely, ρ is monotonically increasing in τ , and approaches one as τ tends to infinity.

The next result extends Lemma 1 to a family of unbounded delays, which allows us to deal with *totally asynchronous iterations* (Bertsekas and Tsitsiklis, 2015, Chapter 6), and shows that the sequence $\{V_k\}$ can still be guaranteed to converge.

Lemma 2 *Let $\{V_k\}$ be a non-negative sequence such that*

$$V_{k+1} \leq qV_k + p \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0, \quad (13)$$

for some non-negative scalars q and p . Suppose that the delay sequence $\{\tau_k\}$ satisfies

$$\lim_{k \rightarrow +\infty} k - \tau_k = +\infty. \quad (14)$$

If $q + p < 1$, then $\{V_k\}$ asymptotically converges to zero:

$$\lim_{k \rightarrow +\infty} V_k = 0.$$

Proof. See Appendix B.1. ■

Lemma 2 provides a test for asymptotic convergence of asynchronous iterations with delays satisfying (14). Assumption (14) holds for bounded delays, irrespectively of whether they are constant or time-varying. Moreover, delays satisfying (14) can be unbounded, as exemplified by $\tau_k = \lfloor 0.2k \rfloor$ and $\tau_k = \lfloor \sqrt{k} \rfloor$. This constraint on delays guarantees that as the iteration count k increases, the delay τ_k grows at a slower rate than time itself, thereby allowing outdated information about process updates to be eventually purged from the computation. To see this, let us assume that the update step in the gradient descent method is based on gradients computed at stale iterates rather than the current iterate, i.e.,

$$x_{k+1} = x_k - \gamma \nabla f(x_{k-\tau_k}).$$

If τ_k satisfies (14), then given any time $K_1 \in \mathbb{N}$, there exists a time $K_2 \in \mathbb{N}$ such that

$$k - \tau_k \geq K_1, \quad \forall k \geq K_2.$$

This means that given any time K_1 , out-of-date information prior to K_1 will not be used in updates after a sufficiently long time K_2 . Therefore, (14) is satisfied in asynchronous algorithms as long as no processor ceases to update (Bertsekas and Tsitsiklis, 2015).

Although Lemma 2 establishes convergence guarantees for the sequence $\{V_k\}$ also under unbounded delays, it no longer provides any finite-time guarantee or rate of convergence. The next result demonstrates that such guarantees can be obtained when we restrict how the possibly unbounded delay sequence is allowed to evolve.

Lemma 3 *Let $\{V_k\}$ be a non-negative sequence satisfying*

$$V_{k+1} \leq qV_k + p \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0, \quad (15)$$

for some non-negative constants q and p such that $q + p < 1$. In addition, assume that there exists a function $\Lambda : \mathbb{R} \rightarrow \mathbb{R}$ such that the following conditions hold:

- (i) $\Lambda(0) = 1$.
- (ii) Λ is non-increasing.

(iii) $\lim_{k \rightarrow +\infty} \Lambda(k) = 0$ and

$$(q + p)\Lambda(k - \tau_k) \leq \Lambda(k + 1), \quad k \in \mathbb{N}_0. \quad (16)$$

Then $V_k \leq \Lambda(k)V_0$ for all $k \in \mathbb{N}_0$.

Proof. See Appendix B.2. ■

According to Lemma 3, any function Λ satisfying conditions (i)—(iii) can be used to quantify how fast the sequence $\{V_k\}$ converges to zero. For example, if $\Lambda(t) = \rho^t$ with $\rho \in (0, 1)$, then $\{V_k\}$ converges at a linear rate; and if $\Lambda(t) = t^{-\eta}$ with $\eta > 0$, then $\{V_k\}$ is upper bounded by a polynomial function of time. Given q and p , it is clear from (16) that the admissible choices for Λ and, hence, the convergence bounds that we are able to guarantee depend on the delay sequence $\{\tau_k\}$. To clarify this statement, we will analyze a special case of unbounded delays in detail. Assume that $\{\tau_k\}$ can grow unbounded at a linear rate, i.e.,

$$\tau_k \leq \alpha k + \beta, \quad k \in \mathbb{N}_0, \quad (17)$$

where $\alpha \in (0, 1)$ and $\beta \geq 0$. The associated convergence result reads as follows.

Corollary 4 *Let $\{V_k\}$ be a non-negative sequence such that*

$$V_{k+1} \leq qV_k + p \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0,$$

for some non-negative scalars q and p . Suppose the delay sequence $\{\tau_k\}$ satisfies (17). If $q + p < 1$, then

$$V_k \leq \left(\frac{\alpha k}{1 - \alpha + \beta} + 1 \right)^{-\eta} V_0, \quad k \in \mathbb{N}_0,$$

where $\eta = \ln(q + p) / \ln(1 - \alpha)$.

Proof. Conditions (i)—(iii) of Lemma 3 are satisfied by the function

$$\Lambda(t) = \left(\frac{\alpha t}{1 - \alpha + \beta} + 1 \right)^{-\eta}.$$

Corollary 4 shows that for unbounded delays satisfying (17), the convergence rate of the sequence $\{V_k\}$ is $\mathcal{O}(k^{-\eta})$. Note that α , the rate at which the unbounded delays grow large, affects η . Specifically, η is monotonically decreasing with α and approaches zero as α tends to one. Hence, the guaranteed convergence rate slows down as the growth rate of the delays increases. ■

3.2 Results for Iterations on the form (11)

We will now shift our attention to the convergence result for iterations on the form (11). This result adds a lot of flexibility in how we can model and account for different perturbations that appear in the analysis of asynchronous optimization algorithms, and will be central to the developments in Subsections 4.1, 4.2, and 4.3.

Lemma 5 *Let $\{V_k\}$, $\{W_k\}$, and $\{X_k\}$ be non-negative sequences satisfying*

$$X_k + V_{k+1} \leq q_k V_k + p_k \sum_{\ell=(k-\tau_k)_+}^k W_\ell - r_k W_k + e_k, \quad k \in \mathbb{N}_0, \quad (18)$$

where $e_k \in \mathbb{R}$, $q_k \in [0, 1]$, and $p_k, r_k \geq 0$ for all k . Suppose that there is a non-negative integer τ such that

$$0 \leq \tau_k \leq \tau, \quad k \in \mathbb{N}_0.$$

For every $K \in \mathbb{N}_0$, the following statements hold:

1. Assume that $q_k = 1$ for $k \in \mathbb{N}_0$. If

$$\sum_{\ell=0}^{\tau} p_{k+\ell} \leq r_k \quad (19)$$

is satisfied for all $k \in \mathbb{N}_0$, then

$$\begin{aligned} \sum_{k=0}^K X_k &\leq V_0 + \sum_{k=0}^K e_k, \\ V_{K+1} &\leq V_0 + \sum_{k=0}^K e_k. \end{aligned}$$

2. Assume that $p_k = p > 0$ and $r_k = r > 0$ for $k \in \mathbb{N}_0$. Assume also that there exists a constant $q \in (0, 1)$ such that $q_k \geq q$ for $k \in \mathbb{N}_0$. If

$$2\tau + 1 \leq \min \left\{ \frac{1}{1-q}, \frac{r}{p} \right\},$$

then

$$\begin{aligned} V_{K+1} &\leq Q_{K+1} \left(V_0 + \sum_{k=0}^K \frac{e_k}{Q_{k+1}} \right), \\ \sum_{k=0}^K \frac{X_k}{Q_{k+1}} &\leq V_0 + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}, \end{aligned}$$

where Q_k is defined as

$$Q_k = \prod_{\ell=0}^{k-1} q_\ell, \quad k \in \mathbb{N},$$

with $Q_0 = 1$.

Proof. See Appendix B.3. ■

Consider the non-delayed counterpart of (18) with $e_k \equiv 0$, $q_k = q \in (0, 1)$, $p_k = p > 0$ and $r_k = r > 0$:

$$X_k + V_{k+1} \leq qV_k + (p - r)W_k, \quad k \in \mathbb{N}_0.$$

Assume that $q \in (0, 1)$ and $p \leq r$, or equivalently,

$$1 \leq \frac{1}{1 - q} \quad \text{and} \quad 1 \leq \frac{r}{p}.$$

In this case, the sequence $\{V_k\}$ converges linearly to zero at a rate of q . In general, the existence of delays may impair performance, induce oscillations and even instability. However, Lemma 5 shows that for the delayed iteration (18), the convergence rate of $\{V_k\}$ is still q if the maximum delay bound τ satisfies

$$2\tau + 1 \leq \frac{1}{1 - q} \quad \text{and} \quad 2\tau + 1 \leq \frac{r}{p}.$$

This means that up to certain value of the delay, the iteration (18) and its delay-free counterpart have the same guaranteed convergence rate.

4. Applications to Asynchronous Optimization Algorithms

Data-driven optimization problems can grow large both in the number of decision variables and in the number of data points that are used to define the objective and constraints. It may therefore make sense to parallelize the associated optimization algorithms over both data and decision variables, see Figure 1. One popular framework for parallelizing algorithms in the data dimension is the parameter server (Li et al., 2013b). Here, a master node (the server) maintains the decision vector, while the data is divided between a number of worker nodes. When a worker node is queried by the server, it computes and returns the gradient of the part of the objective function defined by its own data. The master maintains an estimate of the gradient of the full objective function, and executes a (proximal) gradient update whenever it receives gradient information from one of the workers. As soon as the master completes an update, it queries idle worker nodes with the updated decision vector. If the asynchrony, measured in terms of the maximum number of iterations carried out by the master between two consecutive gradient updates from any worker, is bounded, then convergence can be guaranteed under mild assumptions on the objective function (Li et al., 2013a; Aytakin et al., 2016).

A natural way to parallelize problems with high-dimensional decision vectors is to use block-coordinate updates. In these methods, the decision vector is divided into sub-vectors, and different processing elements update the sub-vectors in parallel. In the partially and totally asynchronous models of Bertsekas and Tsitsiklis (2015), each processing element is responsible for storing and updating one sub-vector, and it does so using delayed information of the remaining decision variables retrieved from the other (remote) processors. Under weak assumptions on the communications delays and update rates of individual processors, convergence can be proven for contraction mappings with respect to the block-maximum

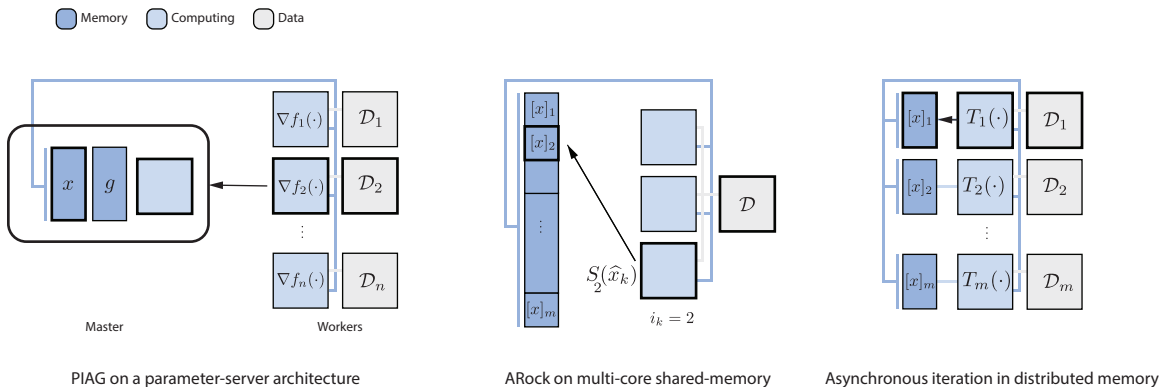


Figure 1: Three different parallel architectures studied for various algorithms in Section 4. The PIAG and SGD algorithms run on the parameter server (left) distribute data over multiple nodes that are able to evaluate the corresponding loss function gradients, while the master maintains and updates the decision vector. In the ARock framework (middle), multiple computing units access shared memory and update (randomly selected) sub-vectors of the overall decision vector in parallel. Finally, the totally asynchronous framework (right) allows us to model loosely coupled distributed architectures where computing nodes retrieve parts of the global decision vector from remote nodes, and evaluate components of an operator to update their local decisions.

norm (Bertsekas and Tsitsiklis, 2015, Section 6.3). However, only some special combinations of algorithms and optimization problems result in iterations that are contractive with respect to the block-maximum norm (Bertsekas and Tsitsiklis, 2015, Section 3.1). Another type of block-coordinate updates are used in the ARock framework (Peng et al., 2016). Here, the decision vector is stored in shared memory, and the parallel computing elements pick sub-vectors uniformly at random to update whenever they terminate their previous work. Under an assumption of bounded asynchrony, convergence of ARock can be established for a wide range of objective functions (Peng et al., 2016).

In the remaining parts of this paper, we demonstrate how the sequence results introduced in the previous section allows us to strengthen the existing convergence guarantees for the algorithms discussed above. Specifically, we improve iteration complexity bounds for the proximal incremental aggregated gradient (PIAG) method, which is suitable for implementation in the parameter server framework, with respect to both the amount of asynchrony and the problem conditioning; we derive tighter guarantees for the asynchronous stochastic gradient descent method, which depend on the average delay rather than on maximal delay; we prove the linear rate of convergence for ARock under larger step-sizes and provide better scalability properties with respect to the number of parallel computing elements; and we describe a unified Lyapunov-based approach for analysis of totally and partially asynchronous iterations involving maximum norm contractions, that allows to derive convergence rate guarantees also outside the partially asynchronous regime.

4.1 Proximal Incremental Aggregated Gradient Method

We begin by considering composite optimization problems of the form

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad P(x) := F(x) + R(x). \quad (20)$$

Here, x is the decision variable, F is the average of many component functions f_i , i.e.,

$$F(x) = \frac{1}{n} \sum_{i=1}^n f_i(x),$$

and R is a proper closed convex function that may be non-differentiable and extended real-valued. We use \mathcal{X}^* to denote the set of optimal solutions of (20) and P^* to denote the corresponding optimal value. We impose the following assumptions on Problem (20).

Assumption 1 *The optimal set \mathcal{X}^* is non-empty.*

Assumption 2 *Each function $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$, $i \in [n]$, is convex and L_i -smooth.*

Note that under Assumption 2, the average function F is L_F -smooth (Xiao and Zhang, 2014), where

$$L_F \leq L := \frac{1}{n} \sum_{i=1}^n L_i. \quad (21)$$

In the optimization problem (20), the role of the regularization term R is to favor solutions with certain structures. Common choices of R include: the ℓ_1 norm, $R(x) = \lambda \|x\|_1$ with $\lambda > 0$, used to promote sparsity in solutions; and the indicator function of a non-empty closed convex set $\mathcal{X} \subseteq \mathbb{R}^d$,

$$R(x) = \begin{cases} 0, & \text{if } x \in \mathcal{X}, \\ +\infty, & \text{otherwise,} \end{cases}$$

used to force the admissible solutions to lie in \mathcal{X} . A comprehensive catalog of regularization terms is given in Beck (2017).

Optimization problems on the form (20) are known as *regularized empirical risk minimization* problems and often arise in machine learning, signal processing, and statistical estimation (see, e.g., Hastie et al. 2009). In such problems, we are given a collection of n training samples $\{(a_1, b_1), \dots, (a_n, b_n)\}$, where each $a_i \in \mathbb{R}^d$ is a feature vector, and each $b_i \in \mathbb{R}$ is the desired response. A classical example is the least-squares regression where the component functions are given by

$$f_i(x) = \frac{1}{2}(a_i^\top x - b_i)^2, \quad i \in [n],$$

and popular choices of the regularization terms include $R(x) = \lambda_1 \|x\|_2^2$ (ridge regression), $R(x) = \lambda_2 \|x\|_1$ (Lasso), or $R(x) = \lambda_1 \|x\|_2^2 + \lambda_2 \|x\|_1$ (elastic net) for some non-negative

parameters λ_1 and λ_2 . Another example is logistic regression for binary classification problems, where each $b_i \in \{-1, 1\}$ is the desired class label and the component functions are

$$f_i(x) = \log \left(1 + \exp(-b_i a_i^\top x) \right), \quad i \in [n].$$

A standard method for solving Problem (20) is the *proximal gradient* (PG) method, which consists of a gradient step followed by a proximal mapping. More precisely, the PG method is described by Algorithm 1, where γ is a positive step-size, and the prox-operator (proximal mapping) is defined as

$$\text{prox}_{\gamma R}(x) = \underset{u \in \mathbb{R}^d}{\text{argmin}} \left\{ \frac{1}{2} \|u - x\|^2 + \gamma R(u) \right\}.$$

Under Assumptions 1 and 2, the iterates generated by the PG method with $\gamma = \frac{1}{L}$ satisfy

$$P(x_k) - P^* \leq \frac{L \|x_0 - x^*\|^2}{2k} \quad (22)$$

for all $k \in \mathbb{N}$ (Beck, 2017, Theorem 10.21). This means that Algorithm 1 achieves the convergence rate $\mathcal{O}(1/k)$ in function values to the optimal value.

Algorithm 1 Proximal Gradient (PG) Method

Input: $x_0 \in \mathbb{R}^d$, step-size $\gamma > 0$, number of iterations $K \in \mathbb{N}$

- 1: Initialize $k \leftarrow 0$
 - 2: **while** $k < K$ **do**
 - 3: Set $g_k \leftarrow \frac{1}{n} \sum_{i=1}^n \nabla f_i(x_k)$
 - 4: Set $x_{k+1} \leftarrow \text{prox}_{\gamma R}(x_k - \gamma g_k)$
 - 5: Set $k \leftarrow k + 1$
 - 6: **end while**
-

Each iteration of the PG method requires computing the gradients for all n component functions. When n is large, this per iteration cost is expensive, and hence often results in slow convergence. An effective alternative is the *proximal incremental aggregated gradient* (PIAG) method that exploits the additive structure of (20) and operates on a *single* component function at a time, rather than on the entire cost function (Tseng and Yun, 2014). The PIAG method evaluates the gradient of only one component function per iteration, but keeps a memory of the most recent gradients of all component functions to approximate the full gradient ∇F . Specifically, at iteration k , the method will have stored $\nabla f_i(x_{[i]})$ for all $i \in [n]$, where $x_{[i]}$ represents the latest iterate at which ∇f_i was evaluated. An integer $j \in [n]$ is then chosen, and the full gradient $\nabla F(x_k)$ is approximated by

$$g_k = \frac{1}{n} \left(\nabla f_j(x_k) - \nabla f_j(x_{[j]}) + \sum_{i=1}^n \nabla f_i(x_{[i]}) \right).$$

The aggregated gradient vector g_k is employed to update the current iterate x_k via

$$x_{k+1} = \text{prox}_{\gamma R}(x_k - \gamma g_k).$$

Thus, the PIAG method uses outdated gradients from previous iterations for the components f_i , $i \neq j$, and does not need to compute gradients of these components at iteration k .

A formal description of the PIAG method is presented as Algorithm 2. Let us define $s_{i,k}$ as the iteration number in which the gradient of the component function f_i is updated for the last time before the completion of the k th iteration. Then, we can rewrite the aggregated gradient vector g_k as

$$g_k = \frac{1}{n} \sum_{i=1}^n \nabla f_i(x_{s_{i,k}}).$$

Note that $s_{i,k} \in \{0, \dots, k\}$ for each $i \in [n]$ and $k \in \mathbb{N}_0$. In a traditional (serial) implementation of PIAG, the selection of j on line 7 follows a deterministic rule (for example, the cyclic rule) and the gradient of the component function f_j is computed at the current iterate, i.e., $s_{j,k} = k$. In the parameter server implementation, the index j will not be chosen, but rather assigned to the identity of the worker that returns its gradient to the master at iteration k . Since workers can exchange information with the master independently of each other, worker j may evaluate ∇f_j at an outdated iterate $x_{s_{j,k}}$, where $s_{j,k} < k$, and send the result to the server. We assume that each component function is sampled at least once in the past τ iterations of the PIAG method. In other words, there is a fixed non-negative integer τ such that the indices $s_{i,k}$ satisfy

$$(k - \tau)_+ \leq s_{i,k} \leq k, \quad i \in [n], k \in \mathbb{N}_0.$$

The value τ can be viewed as an upper bound on the delay encountered by the gradients of the component functions. For example, if the component functions are chosen one by one using a deterministic cyclic order on the index set $\{1, \dots, n\}$, then $\tau = n - 1$.

Algorithm 2 Proximal Incremental Aggregated Gradient (PIAG) Method

Input: $x_0 \in \mathbb{R}^d$, step-size $\gamma > 0$, number of iterations $K \in \mathbb{N}$

- 1: **for** $i = 1$ to n **do**
 - 2: Compute $\nabla f_i(x_0)$
 - 3: Store $\nabla f_i(x_{[i]}) \leftarrow \nabla f_i(x_0)$
 - 4: **end for**
 - 5: Initialize $k \leftarrow 0$
 - 6: **while** $k < K$ **do**
 - 7: Choose j from $\{1, \dots, n\}$
 - 8: Compute $\nabla f_j(x_k)$
 - 9: Set $g_k \leftarrow \frac{1}{n} (\nabla f_j(x_k) - \nabla f_j(x_{[j]}) + \sum_{i=1}^n \nabla f_i(x_{[i]}))$
 - 10: Store $\nabla f_j(x_{[j]}) \leftarrow \nabla f_j(x_k)$
 - 11: Set $x_{k+1} \leftarrow \text{prox}_{\gamma R}(x_k - \gamma g_k)$
 - 12: Set $k \leftarrow k + 1$
 - 13: **end while**
-

The following result shows that in the analysis of the PIAG method, we can establish iterate relationships on the form (11).

Lemma 6 *Let Assumptions 1 and 2 hold. Suppose that $\{\alpha_k\}$ is a sequence defined by*

$$\alpha_k = k + \alpha_0, \quad k \in \mathbb{N}_0,$$

where α_0 is a non-negative scalar. Let

$$V_k = 2\gamma\alpha_k(P(x_k) - P^*) + \|x_k - x^*\|^2,$$

and $W_k = \|x_{k+1} - x_k\|^2$ for $k \in \mathbb{N}_0$. Then, the iterates $\{x_k\}$ generated by Algorithm 2 satisfy

$$V_{k+1} \leq V_k + \gamma L(\alpha_k + \tau + 1) \sum_{\ell=(k-\tau)_+}^k W_\ell - (2\alpha_k + 1 - \gamma L\tau\alpha_k) W_k, \quad k \in \mathbb{N}_0.$$

Proof. See Appendix C.1. ■

Using this iterate relationship, the sequence result in Lemma 5 yields the following convergence guarantee for the PIAG method.

Theorem 7 *Let Assumptions 1 and 2 hold. Suppose that $\gamma \in (0, \gamma_{\max}]$ with*

$$\gamma_{\max} = \frac{1}{L(2\tau + 1)}.$$

Then, for every $k \in \mathbb{N}$ and any $x^ \in \mathcal{X}^*$, the iterates $\{x_k\}$ generated by Algorithm 2 satisfy*

$$P(x_k) - P^* \leq \frac{\frac{1}{2\gamma}\|x_0 - x^*\|^2 + \tau(P(x_0) - P^*)}{k + \tau}.$$

Proof. See Appendix C.2. ■

According to Theorem 7, the PIAG iterations converge at a rate of $\mathcal{O}(1/k)$ when the constant step-size γ is appropriately tuned. The convergence rate depends on the choice of γ . For example, if we pick

$$\gamma = \frac{1}{L(2\tau + 1)}, \tag{23}$$

then the corresponding $\{x_k\}$ converges in terms of function values with the rate

$$P(x_k) - P^* \leq \frac{L\|x_0 - x^*\|^2 + 2\tau(L\|x_0 - x^*\|^2 + P(x_0) - P^*)}{2(k + \tau)}, \quad k \in \mathbb{N}. \tag{24}$$

One can verify that the right-hand side is monotonically increasing in τ . Hence, the guaranteed convergence rate slows down as the delays increase in magnitude. In the case that $\tau = 0$, the bound (24) reduces to (22), which is achieved by the PG method. From (24), we can see that if

$$k \geq K_\epsilon = \frac{L\|x_0 - x^*\|^2 + 2\tau(L\|x_0 - x^*\|^2 + P(x_0) - P^*)}{2\epsilon} - \tau,$$

then PIAG with the step-size choice (23) achieves an accuracy of $P(x_k) - P^* \leq \epsilon$. This shows that the PIAG method has an iteration complexity of $\mathcal{O}(L(\tau + 1)/\epsilon)$ for convex problems. Therefore, as τ increases, the complexity bound deteriorates linearly with τ . Note that the linear dependence on the maximum delay bound τ is unavoidable and cannot further be improved (Arjevani et al., 2020).

Remark 8 *Sun et al. (2019) analyzed convergence of the PIAG method under Assumptions 1 and 2 and proved $\mathcal{O}(C/k)$ convergence rate, where C is a positive constant. While the constant C is implicit in Sun et al. (2019), Huang et al. (2021) showed that the analysis in Sun et al. (2019) guarantees an $\mathcal{O}(\tau^3 L^2/\epsilon)$ iteration complexity for the PIAG method. In comparison with this result, Theorem 7 gives a better dependence on the Lipschitz constant (L vs. L^2) and on the maximum delay bound (τ vs. τ^3) in the iteration complexity.*

Next, we restrict our attention to composite optimization problems under the following quadratic functional growth condition.

Assumption 3 *There exists a constant $\mu > 0$ such that*

$$P(x) - P^* \geq \frac{\mu}{2} \|x - \Pi_{\mathcal{X}^*}(x)\|^2, \quad \forall x \in \text{dom } R, \quad (25)$$

where $\text{dom } R$, defined as $\text{dom } R := \{x \in \mathbb{R}^d \mid R(x) < +\infty\}$, is the effective domain of the function R , and $\Pi_{\mathcal{X}^*}(\cdot)$ denotes the Euclidean-norm projection onto the set \mathcal{X}^* , i.e.,

$$\Pi_{\mathcal{X}^*}(x) = \operatorname{argmin}_{u \in \mathcal{X}^*} \|u - x\|.$$

Assumption 3 implies that the objective function grows faster than the squared distance between any feasible point and the optimal set. While every strongly convex function satisfies the quadratic functional growth condition (25), the converse is not true in general (Necoara et al., 2019). For example, if $A \in \mathbb{R}^{m \times d}$ is rank deficient and $b \in \mathbb{R}^m$, the function $F(x) = \|Ax - b\|^2$ is not strongly convex, but it satisfies Assumption 3. Other examples of objective functions which satisfy the quadratic functional growth condition can be found in Necoara et al. (2019). Let us define the condition number of the optimization problem (20) as $Q = L/\mu$. The role of Q in determining the linear convergence rate of (proximal) gradient methods is well-known (Nesterov, 2013).

We next show that under Assumption 3, an iterate relationship on the form (11) appears in the analysis of the PIAG method.

Lemma 9 *Suppose that Assumptions 1, 2 and 3 hold. Let*

$$V_k = \frac{2}{L} (P(x_k) - P^*) + \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2,$$

and $W_k = \|x_{k+1} - x_k\|^2$ for $k \in \mathbb{N}_0$. Then, the iterates $\{x_k\}$ generated by Algorithm 2 satisfy

$$V_{k+1} \leq \left(\frac{1}{1 + \gamma\mu\theta} \right) V_k + \frac{1 + \gamma L(\tau + 1)}{1 + \gamma\mu\theta} \sum_{\ell=(k-\tau)_+}^k W_\ell - \frac{\frac{2}{\gamma L} + 1 - \tau}{1 + \gamma\mu\theta} W_k,$$

where $\theta = \frac{Q}{Q+1}$.

Proof. See Appendix C.3. ■

We use Lemma 9 together with Lemma 5 to derive the convergence rate of the PIAG method for optimization problems whose objective functions satisfy the quadratic functional growth condition.

Theorem 10 *Let Assumptions 1, 2 and 3 hold. Suppose that the step-size γ is set to*

$$\gamma = \frac{h}{L(2\tau + 1)}, \quad h \in (0, 1].$$

Then, for every $k \in \mathbb{N}$, the iterates $\{x_k\}$ generated by Algorithm 2 satisfy

$$\begin{aligned} \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2 &\leq \left(1 - \frac{1}{1 + (Q + 1)(2\tau + 1)/h}\right)^k \left(\frac{2}{L}(P(x_0) - P^*) + \|x_0 - \Pi_{\mathcal{X}^*}(x_0)\|^2\right), \\ P(x_k) - P^* &\leq \left(1 - \frac{1}{1 + (Q + 1)(2\tau + 1)/h}\right)^k \left(P(x_0) - P^* + \frac{L}{2}\|x_0 - \Pi_{\mathcal{X}^*}(x_0)\|^2\right). \end{aligned}$$

Proof. See Appendix C.4. ■

Theorem 10 demonstrates that under Assumption 3, the PIAG method is linearly convergent by taking a constant step-size inversely proportional to the maximum delay τ . The best guaranteed convergence rate is obtained for the step-size

$$\gamma = \frac{1}{L(2\tau + 1)}.$$

With this choice of γ , the iterates converge linearly in terms of function values with rate

$$P(x_k) - P^* \leq \left(1 - \frac{1}{1 + (Q + 1)(2\tau + 1)}\right)^k \epsilon_0,$$

where $\epsilon_0 = P(x_0) - P^* + \frac{L}{2}\|x_0 - \Pi_{\mathcal{X}^*}(x_0)\|^2$. Taking the logarithm of both sides yields

$$\log(P(x_k) - P^*) \leq k \log\left(1 - \frac{1}{1 + (Q + 1)(2\tau + 1)}\right) + \log(\epsilon_0).$$

Since $\log(1 + x) \leq x$ for any $x > -1$, it follows that

$$\log(P(x_k) - P^*) \leq -\frac{k}{1 + (Q + 1)(2\tau + 1)} + \log(\epsilon_0).$$

Therefore, if the number of iterations satisfy

$$k \geq K_\epsilon = (1 + (Q + 1)(2\tau + 1)) \log\left(\frac{\epsilon_0}{\epsilon}\right),$$

then $P(x_k) - P^* \leq \epsilon$. We conclude that the PIAG method achieves an iteration complexity of $\mathcal{O}(Q(\tau + 1) \log(1/\epsilon))$ for optimization problems satisfying the quadratic functional growth condition. Note that when $\tau = 0$, this bound becomes $\mathcal{O}(Q \log(1/\epsilon))$, which is the iteration complexity for the PG method (Beck, 2017, Theorem 10.30).

As discussed before, if the component functions are selected in a fixed cyclic order, then $\tau = n - 1$. It follows from Theorem 7 and Theorem 10 that the iteration complexity of the PIAG method with cyclic sampling is $\mathcal{O}(nL/\epsilon)$ for convex problems and $\mathcal{O}(nQ \log(1/\epsilon))$ for problems whose objective functions satisfy the quadratic functional growth condition. Each iteration of the PIAG method requires only one gradient evaluation, compared to n gradient computations in the PG method. Therefore, in terms of the total number of component gradients evaluated to find an ϵ -optimal solution, the iteration complexity of PG and the cyclic PIAG are the same.

Remark 11 *Schmidt et al. (2017) proposed a randomized variant of PIAG, called stochastic average gradient (SAG), where the component functions are sampled uniformly at random. The iteration complexity of the SAG method, in expectation, is $\mathcal{O}(\max\{n, Q\} \log(1/\epsilon))$ for strongly convex problems and $\mathcal{O}((n + L)/\epsilon)$ for convex problems. The maximum allowable step-size for the SAG method is larger than that of the PIAG method, which can lead to improved empirical performance (Schmidt et al., 2017, Figure 1). Note, however, that in some applications, the component functions must be processed in a particular deterministic order and, hence, random sampling is not possible. For example, in source localization or distributed parameter estimation over wireless networks, sensors may only communicate with their neighbors subject to certain constraints in terms of geography and distance, which can restrict the updates to follow a specific deterministic order (Blatt et al., 2007).*

4.1.1 COMPARISON OF OUR ANALYSIS WITH PRIOR WORK

Blatt et al. (2007) proposed the incremental aggregated gradient (IAG) method for solving unconstrained optimization problems. For a special case where each component function is quadratic, they showed that the IAG method with a constant step-size achieves a linear rate of convergence. The authors obtained these results using a perturbation analysis of the eigenvalues of a periodic linear system. However, the analysis in Blatt et al. (2007) only applies to quadratic objective functions and provides neither an explicit convergence rate nor an explicit upper bound on the step-sizes that ensure linear convergence.

Tseng and Yun (2014) proved global convergence and local linear convergence for the PIAG method in a more general setting where the gradients of the component functions are Lipschitz continuous and each component function satisfies a local error bound assumption. For a L -smooth and (possibly) non-convex function F , the proof of Theorem 4.1 in Tseng and Yun (2014) shows that the iterates generated by PIAG satisfy

$$P(x_{k+1}) - P^* \leq P(x_k) - P^* + \frac{\gamma^2 L}{2} \sum_{\ell=(k-\tau)_+}^k \|d_\ell\|^2 - \gamma \left(1 - \frac{\gamma L \tau}{2}\right) \|d_k\|^2, \quad (26)$$

where d_k is the search direction at iteration k and given by

$$d_k = \operatorname{argmin}_{d \in \mathbb{R}^d} \left\{ \langle g_k, d \rangle + \frac{1}{2} \|d\|^2 + R(x_k + d) \right\}.$$

By defining $V_k = P(x_k) - P^*$, $X_k = \gamma(1 - h)\|d_k\|^2$ with $h \in (0, 1)$, and $W_k = \|d_k\|^2$, we can rewrite (26) as

$$X_k + V_{k+1} \leq V_k + \frac{\gamma^2 L}{2} \sum_{\ell=(k-\tau)_+}^k W_\ell - \gamma \left(h - \frac{\gamma L \tau}{2}\right) W_k.$$

The iterates satisfy a relationship on the form (11). Thus, according to Lemma 5, if

$$\frac{\gamma^2 L(\tau + 1)}{2} \leq \gamma \left(h - \frac{\gamma L \tau}{2}\right),$$

then $\sum_{k=0}^K X_k \leq V_0$ for every $K \in \mathbb{N}$. This, in turn, implies that

$$\sum_{k=0}^K \|d_k\|^2 \leq \frac{P(x_0) - P^*}{\gamma(1-h)}, \quad \text{for } \gamma = \frac{2h}{L(2\tau+1)},$$

which is the same result as in Tseng and Yun (2014) for non-convex objective functions. While the analysis in Tseng and Yun (2014) is more general than that in Blatt et al. (2007), as the authors did not limit the objective function to be strongly convex and quadratic, explicit rate estimates and a characterization of the step-size needed for linear convergence were still missing in Tseng and Yun (2014).

Gürbüzbalaban et al. (2017) provided the first explicit linear rate result for the IAG method. According to the proof of Theorem 3.3 in Gürbüzbalaban et al. (2017), the iterates generated by IAG satisfy a relationship of the form (10) given by

$$V_{k+1} \leq \left(1 - \frac{2\gamma\mu L}{\mu + L}\right) V_k + (9\gamma^4 L^4 \tau^2 + 6\gamma^2 L^2 \tau) \max_{(k-2\tau)_+ \leq \ell \leq k} V_\ell,$$

where $V_k = \|x_k - x^*\|^2$. As in Gürbüzbalaban et al. (2017), we can use Lemma 1 to obtain an iteration complexity of $\mathcal{O}(Q^2 \tau^2 \log(1/\epsilon))$ for the IAG method. However, this analysis has a drawback that the guaranteed bound grows quadratically with both the condition number Q and the maximum delay τ .

The quadratic dependence of the iteration complexity on Q was improved by Aytakin et al. (2016) for strongly convex composite objectives and by Zhang et al. (2021) under a quadratic functional growth condition. Specifically, the proof of Theorem 1 in Aytakin et al. (2016) shows that for the iterates generated by PIAG, we have

$$V_{k+1} \leq \left(\frac{1}{1+\gamma\mu}\right) V_k + \left(\frac{\gamma L(\tau+1)}{1+\gamma\mu}\right) \sum_{\ell=(k-\tau)_+}^k W_\ell - \left(\frac{1}{1+\gamma\mu}\right) W_k,$$

where $V_k = \|x_k - x^*\|^2$ and $W_k = \|x_{k+1} - x_k\|^2$. Since the iterates satisfy a relationship of the form (11), it follows from Lemma 5 that

$$\|x_k - x^*\|^2 \leq \left(\frac{1}{1+\gamma\mu}\right) \|x_0 - x^*\|^2, \quad \text{for } \gamma \in \left(0, \frac{1}{L(2\tau+1)(\tau+1)}\right].$$

The above convergence rate leads to an iteration complexity of $\mathcal{O}(Q\tau^2 \log(1/\epsilon))$, which matches the bound derived in Aytakin et al. (2016). However, this analysis has some limitations. First, the number of iterations required to reach an ϵ -optimal solution increases quadratically with the maximal delay τ . Second, the constant step-size that guarantees linear convergence is inversely proportional to the square of τ .

Vanli et al. (2018) improved the quadratic dependence on τ to a linear one by showing that the iteration complexity of the PIAG method is $\mathcal{O}(Q\tau \log(1/\epsilon))$. According to the proof of Theorem 3.9 in Vanli et al. (2018), the iterates generated by PIAG satisfy

$$V_{k+1} \leq \left(\frac{1}{1+\gamma\mu/16}\right) V_k + \left(\frac{3L}{4(1+\gamma\mu/16)}\right) \sum_{\ell=(k-\tau)_+}^{k-1} W_\ell - \left(\frac{1}{4\gamma(1+\gamma\mu/16)}\right) W_k,$$

where $V_k = P(x_k) - P^*$ and $W_k = \|x_{k+1} - x_k\|^2$. Therefore, by Lemma 5, we have

$$P(x_k) - P^* \leq \left(\frac{1}{1 + \gamma\mu/16} \right)^k (P(x_0) - P^*), \quad \text{for } \gamma \in \left(0, \frac{1}{3L(2\tau + 1)} \right],$$

which is the same bound as in Vanli et al. (2018). In comparison with this result, Theorem 10 allows PIAG to use larger step-sizes that leads to a tighter guaranteed convergence rate. This improvement is achieved through our choice of the sequence V_k , which includes two terms: $P(x_k) - P^*$ and $\|x_k - x^*\|^2$. In addition, the analysis in Vanli et al. (2018) only applies to strongly convex problems.

We have thus shown that our sequence results can be used to obtain the convergence guarantees established in Tseng and Yun (2014); Gürbüzbalaban et al. (2017); Aytakin et al. (2016); Vanli et al. (2018), as their analysis involves recurrences of the form (10) or (11). Upon comparing our analysis with previous work, it becomes clear that for a specific algorithm, such as PIAG, the range of step-sizes and convergence rates guaranteed by Lemmas 1- 5 depend heavily on the choice of sequences V_k , W_k , and X_k . Selecting these sequences involves considering the characteristics of the algorithm and properties of the optimization problem that the algorithm aims to solve. For example, for the PIAG method, we see that

$$V_k = P(x_k) - P^*$$

for non-convex problems,

$$V_k = 2\gamma(k + \tau)(P(x_k) - P^*) + \|x_k - x^*\|^2$$

for convex problems, and

$$V_k = \frac{2}{L}(P(x_k) - P^*) + \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2$$

for problems whose objective functions satisfy the quadratic functional growth condition result in sharper theoretical bounds and larger step-sizes than the previous state-of-the-art.

4.2 Asynchronous SGD

To demonstrate the versatility of the sequence results, and that they are not limited to deliver step-size rules and convergence rates that depend on the maximal delay, we now consider stochastic optimization problems of the form

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad F(x) := \mathbb{E}_{\xi \sim \mathcal{D}}[f(x, \xi)]. \quad (27)$$

Here, x is the decision variable, ξ is a random variable drawn from the probability distribution \mathcal{D} , and $f(\cdot, \xi) : \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable for each ξ . This objective captures, for example, supervised learning where x represents parameters of a machine learning model to be trained. In this case, \mathcal{D} is an unknown distribution of labelled examples, ξ is a data point, $f(x, \xi)$ is the loss of the model with parameters x on the data point ξ , and F is the generalization error. We use \mathcal{X}^* to denote the set of optimal solutions of Problem (27) and F^* to denote the corresponding optimal value.

In machine learning applications, the distribution \mathcal{D} is often unknown, which makes it challenging to solve Problem (27). To support these applications, we do not assume knowledge of F , only access to a stochastic oracle. Each time the oracle is queried with an $x \in \mathbb{R}^d$, it generates an independent and identically distributed (i.i.d.) sample ξ from \mathcal{D} and returns $\nabla f(x, \xi)$, which is an unbiased estimate of $\nabla F(x)$, i.e.,

$$\nabla F(x) = \mathbb{E}_{\xi \sim \mathcal{D}}[\nabla f(x, \xi)].$$

We then use the stochastic gradient $\nabla f(x, \xi)$, instead of $\nabla F(x)$, in the update rule of the optimization algorithm that attempts to minimize F .

The classical stochastic gradient descent (SGD) method is among the first and the most commonly used algorithms developed for solving Problem (27). Its popularity comes mainly from the fact that it is easy to implement and has low computational cost per iteration. The SGD method proceeds iteratively by drawing an i.i.d sample ξ_k from \mathcal{D} , computing $\nabla f(x_k, \xi_k)$, and updating the current vector x_k via

$$x_{k+1} = x_k - \gamma_k \nabla f(x_k, \xi_k),$$

where γ_k is a step-size (or learning rate). For a L -smooth and convex function F , the iteration complexity of the SGD method is

$$\mathcal{O}\left(\frac{L}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right),$$

which is $\mathcal{O}(\sigma^2/\epsilon^2)$ asymptotically in ϵ (Lan, 2012).

The SGD method is inherently serial in the sense that gradient computations take place on a single processor which has access to the whole data set and updates iterations sequentially, i.e., one after another. However, it is often infeasible for a single machine to store and process the vast amounts of data that we encounter in practical problems. In these situations, it is common to implement the SGD method in a master-worker architecture in which several worker processors compute stochastic gradients in parallel based on their portions of the data set while a master processor stores the decision vector and updates the current iterate. The master-worker implementation can be executed in two ways: synchronous and asynchronous.

In the synchronous case, the master will perform an update and broadcast the new decision vector to the workers when it has collected stochastic gradients from *all* the workers. Given M workers, the master performs the following mini-batch SGD update

$$x_{k+1} = x_k - \frac{\gamma_k}{M} \sum_{m=1}^M \nabla f(x_k, \xi_{k,m}),$$

when stochastic gradients $\nabla f(x_k, \xi_{k,1}), \dots, \nabla f(x_k, \xi_{k,M})$ are computed and communicated back by workers. In terms of the total number of stochastic gradients evaluated to find an ϵ -optimal solution, the iteration complexity of the mini-batch SGD method is

$$\mathcal{O}\left(\frac{ML}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right) \tag{28}$$

for L -smooth and convex objective functions (Dekel et al. 2012). This bound is asymptotically $\mathcal{O}(\sigma^2/\epsilon^2)$, which is exactly the asymptotic iteration complexity achieved by the SGD method. However, using M workers in parallel, the mini-batch SGD method can achieve updates at a rate roughly M times faster. This means that mini-batch SGD is expected to enjoy a near-linear speedup in the number of workers.

A main drawback of the mini-batch SGD method is that the workers need to synchronize in each round and compute stochastic gradients at the same decision vector. Due to various factors, such as differences in computational capabilities and communication bandwidth, or interference from other running jobs, some workers may evaluate stochastic gradients slower than others. This causes faster workers to be idle during each iteration and the algorithm suffers from the *straggler* problem, in which the algorithm can only move forward at the pace of the slowest worker. The asynchronous SGD method offers a solution to this issue by allowing the workers to compute gradients at different rates without synchronization and letting the master perform updates using outdated gradients. In other words, there is no need for workers to wait for the others to finish their gradient computations, and the master can update the decision vector every time it receives a stochastic gradient from some worker. For this method, the master performs the update

$$x_{k+1} = x_k - \gamma_k \nabla f(x_{k-\tau_k}, \xi_k), \quad (29)$$

where τ_k is the delay of the gradient at iteration k . The value of τ_k , which is often much greater than zero, captures the staleness of the information used to compute the stochastic gradient involved in the update of x_k . In (29), the index of stochastic noise ξ_k is equal to the iteration number k to show that previous iterates $x_{k'}$ for $k' \leq k$ do not depend on this stochastic noise (Agarwal and Duchi, 2011). Algorithm 3 describes the asynchronous SGD method executed in a master-worker setting with one master and M workers.

Algorithm 3 Asynchronous SGD

Input: $x_0 \in \mathbb{R}^d$

- 1: Server sends x_0 to all workers
 - 2: **for** $k = 0, 1, \dots$ **do**
 - 3: Workers compute stochastic gradients at the assigned points in parallel
 - 4: Gradient $\nabla f(x_{k-\tau_k}, \xi_k)$ arrives from some worker $m_k \in \{1, \dots, M\}$
 - 5: Server updates $x_{k+1} \leftarrow x_k - \gamma_k \nabla f(x_{k-\tau_k}, \xi_k)$
 - 6: Server sends x_{k+1} to worker m_k
 - 7: **end for**
-

To characterize the iteration complexity and the convergence rate of Algorithm 3, we make the following assumptions.

Assumption 4 *The objective function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth.*

Assumption 5 *The stochastic gradients have bounded variance, i.e., there exists a constant $\sigma \geq 0$ such that*

$$\mathbb{E}_{\xi \sim \mathcal{D}} [\|\nabla f(x, \xi) - \nabla F(x)\|^2] \leq \sigma^2.$$

Assumption 6 *There is a non-negative integer τ_{\max} such that*

$$0 \leq \tau_k \leq \tau_{\max}, \quad k \in \mathbb{N}_0.$$

Most existing theoretical guarantees for the asynchronous SGD method show that the number of iterations required to reach an ϵ -optimal solution grows with the maximum delay τ_{\max} (see, e.g., Agarwal and Duchi 2011; Lian et al. 2015; Feyzmahdavian et al. 2016; Mania et al. 2017; Arjevani et al. 2020; Stich and Karimireddy 2020). In particular, Arjevani et al. (2020) analyzed positive semi-definite quadratic functions and Stich and Karimireddy (2020) considered general L -smooth convex functions and proved that if the delays are always constant ($\tau_k = \tau_{\max}$ for $k \in \mathbb{N}_0$), the asynchronous SGD method attains an iteration complexity bound of the form

$$\mathcal{O}\left(\frac{L(\tau_{\max} + 1)}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right). \quad (30)$$

This bound is unable to exclude the possibility that the performance of asynchronous SGD degrades proportionally to τ_{\max} . In heterogeneous environments involving workers with different computational speeds, τ_{\max} is effectively determined by the slowest worker. This suggests that similar to mini-batch SGD, the asynchronous SGD method may face the straggler issue. Moreover, the step-size used to achieve the bound (30) is inversely proportional to τ_{\max} , and may hence be quite small, which degrades the performance of the algorithm in practice.

To mitigate the impact of stragglers on asynchronous optimization, Aviv et al. (2021) and Cohen et al. (2021) developed delay-adaptive SGD methods whose rates depend on the *average* delay τ_{ave} rather than on τ_{\max} . Performance of these methods are more robust to asynchrony since τ_{ave} can be significantly smaller than τ_{\max} , especially when some workers are much slower than others. For instance, consider the implementation of the asynchronous SGD method using two workers, where one worker is 1000 times faster than the other in computing gradients. When the master waits for the slower worker to finish computing a gradient, it can use the faster worker’s gradients to produce 1000 updates. As a result, τ_k is equal to 0 for $k = 0, \dots, 999$, and $\tau_{1000} = 1000$. This indicates that the average delay τ_{ave} would be approximately 1, while $\tau_{\max} = 1000$. Although this scenario is hypothetical and constructed to make a point, actual delays tend to have $\tau_{\max} \gg \tau_{\text{ave}}$. For example, the delay measurements for a 40 worker implementation of asynchronous SGD reported in Figure 1 in Mishchenko et al. (2022) have $\tau_{\max} = 1200$, while $\tau_{\text{ave}} = 40$.

Although the methods developed in Aviv et al. (2021) and Cohen et al. (2021) are more robust to straggling workers, there are some limitations to their convergence guarantees. Specifically, the rates in Aviv et al. (2021) were derived under the assumption that the gradients are uniformly bounded, while the results in Cohen et al. (2021) only hold with probability $\frac{1}{2}$. Koloskova et al. (2022) and Mishchenko et al. (2022) have recently addressed these limitations, and shown that asynchronous SGD is always faster than mini-batch SGD regardless of the delay patterns. In particular, Koloskova et al. (2022) proposed a delay-adaptive step-size rule, under which asynchronous SGD achieves an iteration complexity

$$\mathcal{O}\left(\frac{L(\tau_{\text{ave}} + 1)}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right)$$

for L -smooth non-convex functions. Mishchenko et al. (2022) provided convergence guarantees for L -smooth convex, strongly convex, and non-convex functions which depend only on the number of workers used to implement the algorithm.

Similar to Koloskova et al. (2022), we investigate the convergence of the asynchronous SGD method with delay-dependent step-sizes. We define the step-size sequence as

$$\gamma_k = \begin{cases} \gamma, & \tau_k \leq \tau_{\text{th}}, \\ 0, & \tau_k > \tau_{\text{th}}, \end{cases} \quad (31)$$

where the threshold parameter τ_{th} is a non-negative constant satisfying

$$\min \{2\tau_{\text{ave}}, \tau_{\text{max}}\} \leq \tau_{\text{th}}.$$

The adaptive rule (31) sets the step-size to a constant value γ when the delay of the k th iteration is at most τ_{th} . Otherwise, it sets the step-size to zero, effectively dropping the gradients with large delays. Note that (31) reduces to a constant step-size rule by setting $\tau_{\text{th}} \geq \tau_{\text{max}}$. In this case, the step-size remains fixed at γ for all iterations, regardless of the delay values. The main goal of this section is to use the sequence results from Section 3 to: (i) extend the results of Koloskova et al. (2022) to L -smooth convex and strongly convex functions; (ii) improve upon the previously best-known rate of the asynchronous SGD method with constant step-sizes for strongly convex functions given in Stich and Karimireddy (2020); and (iii) recover the convergence rates derived in Mishchenko et al. (2022).

Remark 12 *To use the delay-dependent step-size rule (31), prior knowledge of τ_{ave} is required. This is challenging since gradient delays are usually difficult to predict before implementing the algorithm. However, Koloskova et al. (2022) and Mishchenko et al. (2022) have shown that for the asynchronous SGD method running on M workers, it holds that*

$$\tau_{\text{ave}} \leq M - 1.$$

Therefore, Algorithm 3 can be implemented without prior knowledge of τ_{ave} by setting

$$\tau_{\text{th}} = 2(M - 1).$$

The following result shows that in the analysis of the asynchronous SGD method, we can establish iterate relationships on the form (11).

Lemma 13 *Suppose that Assumptions 4–6 hold. Let*

$$V_k = \mathbb{E}[\|x_k - x^*\|^2], \quad X_k = 2\gamma_k \mathbb{E}[F(x_k) - F^*], \quad \text{and} \quad W_k = \gamma_k^2 \mathbb{E}[\|\nabla F(x_{k-\tau_k})\|^2],$$

where the expectation is over all random variables ξ_0, \dots, ξ_{k-1} . For the iterates generated by Algorithm 3, the following statements hold:

1. If F is convex, then

$$X_k + V_{k+1} \leq V_k + 2\gamma_k \tau_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} W_\ell - \left(\frac{1}{\gamma_k L} - 1\right) W_k + \left(\gamma_k^2 + 2\gamma_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2\right) \sigma^2$$

for $\gamma_k > 0$, and $V_{k+1} = V_k$ for $\gamma_k = 0$.

2. If F is μ -strongly convex, then

$$V_{k+1} \leq (1 - \gamma_k \mu) V_k + 2\gamma_k \tau_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} W_\ell - \left(\frac{1}{\gamma_k L} - 1 \right) W_k + \left(\gamma_k^2 + 2\gamma_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2 \right) \sigma^2$$

for $\gamma_k > 0$, and $V_{k+1} = V_k$ for $\gamma_k = 0$.

Proof. See Appendix D.1. ■

Using this iterate relationship, Lemma 5 allows us to derive convergence guarantees for the asynchronous SGD method. Our proof is based on the observation that, even if some gradients with large delays are discarded, gradients with delays less than or equal to the threshold parameter τ_{th} are sufficient to ensure the convergence of the algorithm.

Theorem 14 *Suppose Assumptions 4–6 hold. Let τ_{th} be a non-negative constant such that*

$$\min \{2\tau_{\text{ave}}, \tau_{\text{max}}\} \leq \tau_{\text{th}},$$

where τ_{ave} is the average delay, i.e.,

$$\tau_{\text{ave}} = \frac{1}{K+1} \sum_{k=0}^K \tau_k.$$

Using the delay-dependent step-size rule (31) in Algorithm 3 yields the following results:

1. For convex F and any $x^* \in \mathcal{X}^*$, choosing

$$\gamma \in \left(0, \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)} \right],$$

guarantees that

$$\mathbb{E}[F(\bar{x}_K) - F^*] \leq \frac{\|x_0 - x^*\|^2}{\gamma(K+1)} + (1 + \sqrt{2})\gamma\sigma^2, \quad K \in \mathbb{N}_0,$$

where \bar{x}_K is the weighted average of x_0, \dots, x_K defined as

$$\bar{x}_K = \frac{1}{\sum_{k=0}^K \gamma_k} \sum_{k=0}^K \gamma_k x_k.$$

2. For strongly convex F , choosing

$$\gamma \in \left(0, \frac{1}{L(2\tau_{\text{th}} + 1)} \right],$$

ensures that

$$\mathbb{E}[\|x_K - x^*\|^2] \leq \exp\left(-\frac{\gamma\mu K}{2}\right) \|x_0 - x^*\|^2 + \frac{2\gamma\sigma^2}{\mu}, \quad K \in \mathbb{N}.$$

Proof. See Appendix D.2. ■

According to Theorem 14, the asynchronous SGD method converges to a ball around the optimum at a rate of $\mathcal{O}(1/K)$ for convex functions and at a linear rate for strongly convex functions. The choice of γ in the delay-adaptive step-size rule (31) affects both the convergence rate and the residual error: decreasing γ reduces the residual error, but it also results in a slower convergence. Next, we present a possible strategy for selecting γ to achieve an ϵ -optimal solution.

Theorem 15 *Suppose that Assumptions 4–6 hold. Given any $\epsilon > 0$, the following statements hold:*

1. *For convex F , choosing*

$$\gamma = \min \left\{ \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)}, \frac{\epsilon}{2(\sqrt{2} + 1)\sigma^2} \right\},$$

implies $\mathbb{E}[F(\bar{x}_K) - F^] \leq \epsilon$ after*

$$\mathcal{O} \left(\frac{L(\tau_{\text{th}} + 1)}{\epsilon} + \frac{\sigma^2}{\epsilon^2} \right)$$

iterations.

2. *For strongly convex F , setting γ to*

$$\gamma = \min \left\{ \frac{1}{L(2\tau_{\text{th}} + 1)}, \frac{\epsilon\mu}{4\sigma^2} \right\},$$

guarantees that $\mathbb{E}[\|x_K - x^\|^2] \leq \epsilon$ after*

$$\mathcal{O} \left(\left(\frac{L(\tau_{\text{th}} + 1)}{\mu} + \frac{\sigma^2}{\epsilon\mu^2} \right) \log \left(\frac{1}{\epsilon} \right) \right)$$

iterations.

Proof. See Appendix D.3. ■

According to Theorem 15, the value of γ needed to achieve an ϵ -optimal solution depends on two terms. The role of the first term is to decrease the effects of asynchrony on the convergence of the algorithm, while the second term aims to control the noise from stochastic gradient information. The first term is monotonically decreasing in the threshold parameter τ_{th} . At the same time, the guaranteed bounds on the iteration complexity are monotonically increasing in τ_{th} . Therefore, choosing a smaller value for τ_{th} allows the algorithm to use larger step-sizes by discarding more gradients with long delays, and results in fewer iterations required to attain an ϵ -optimal solution.

When the threshold parameter τ_{th} is set to $2\tau_{\text{ave}}$, Theorem 15 shows that the asynchronous SGD method with delay-adaptive step-sizes achieves an iteration complexity of

$$\mathcal{O} \left(\frac{L(\tau_{\text{ave}} + 1)}{\epsilon} + \frac{\sigma^2}{\epsilon^2} \right)$$

for convex objective functions. The average delay τ_{ave} only appears in the first term and its negative impact on the convergence rate is asymptotically negligible once

$$\tau_{\text{ave}} \leq \mathcal{O}\left(\frac{\sigma^2}{L\epsilon}\right). \quad (32)$$

In this case, the asynchronous SGD method has the same asymptotic iteration complexity as the serial SGD method. Our asymptotic rate $\mathcal{O}(\sigma^2/\epsilon^2)$ for asynchronous SGD is consistent with the results in Arjevani et al. (2020) and Stich and Karimireddy (2020). However, their requirement to guarantee such rate is

$$\tau_{\text{max}} \leq \mathcal{O}\left(\frac{\sigma^2}{L\epsilon}\right),$$

which is more conservative than (32) since necessarily $\tau_{\text{ave}} \leq \tau_{\text{max}}$.

According to Remark 12, we can select the threshold parameter as $\tau_{\text{th}} = 2(M - 1)$. By doing so, it follows from Theorem 15 that for convex problems, we have

$$\mathcal{O}\left(\frac{LM}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right).$$

By comparing this bound with (28), we can see that the asynchronous SGD method with delay-adaptive step-sizes attains the same iteration complexity as the mini-batch SGD method. Since each update of mini-batch SGD takes the time needed by the slowest worker, its expected per-iteration time is slower than that of asynchronous SGD. This means that asynchronous SGD outperforms mini-batch SGD regardless of the delays in the gradients, as previously proved by Koloskova et al. (2022) and Mishchenko et al. (2022).

By setting $\tau_{\text{th}} = 2\tau_{\text{ave}}$, the iteration complexity for strongly convex functions is

$$\mathcal{O}\left(\left(\frac{L(\tau_{\text{ave}} + 1)}{\mu} + \frac{\sigma^2}{\epsilon\mu^2}\right) \log\left(\frac{1}{\epsilon}\right)\right).$$

In the case that $\tau_{\text{ave}} = 0$, the preceding guaranteed bound reduces to the one obtained in Gower et al. (2019) for the serial SGD method. We can see that the average delay τ_{ave} can be as large as $\mathcal{O}(\sigma^2/(\mu L\epsilon))$ without affecting the asymptotic rate of Algorithm 3.

An alternative approach to tune the step-size in Algorithm 3 is to use γ that depends on a prior knowledge of the number of iterations to be performed (Lan, 2012; Mishchenko et al., 2022). Assume that the number of iterations is fixed in advance, say equal to \mathcal{K} . The following result provides the convergence rate of Algorithm 3 in terms of \mathcal{K} .

Theorem 16 *Let Assumptions 4–6 hold. Given $\mathcal{K} \in \mathbb{N}$, the following statements hold:*

1. *For convex F and*

$$\gamma = \min \left\{ \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)}, \frac{\|x_0 - x^*\|}{\sigma\sqrt{\sqrt{2} + 1}\sqrt{\mathcal{K} + 1}} \right\},$$

Algorithm 3 ensures

$$\mathbb{E}[F(\bar{x}_{\mathcal{K}}) - F^*] = \mathcal{O}\left(\frac{L(\tau_{\text{th}} + 1)}{\mathcal{K}} + \frac{\sigma}{\sqrt{\mathcal{K}}}\right).$$

2. For strongly convex F and

$$\gamma = \min \left\{ \frac{1}{L(2\tau_{\text{th}} + 1)}, \frac{2}{\mu\mathcal{K}} \log \left(1 + \frac{\mu^2\mathcal{K}\|x_0 - x^*\|^2}{4\sigma^2} \right) \right\},$$

Algorithm 3 guarantees that

$$\mathbb{E}[\|x_{\mathcal{K}} - x^*\|^2] = \tilde{\mathcal{O}} \left(\exp \left(-\frac{\mu\mathcal{K}}{L(\tau_{\text{th}} + 1)} \right) + \frac{\sigma^2}{\mathcal{K}} \right).$$

Proof. See Appendix D.4. ■

For strongly convex problems, the previously best known convergence rate under constant step-sizes was given in Stich and Karimireddy (2020) and is expressed as

$$\tilde{\mathcal{O}} \left(\tau_{\text{max}} \exp \left(-\frac{\mu\mathcal{K}}{L(\tau_{\text{max}} + 1)} \right) + \frac{\sigma^2}{\mathcal{K}} \right).$$

However, when we set $\tau_{\text{th}} = \tau_{\text{max}}$, Theorem 16 tells us that the asynchronous SGD method with a constant step-size converges at a rate of

$$\tilde{\mathcal{O}} \left(\exp \left(-\frac{\mu\mathcal{K}}{L(\tau_{\text{max}} + 1)} \right) + \frac{\sigma^2}{\mathcal{K}} \right).$$

Therefore, we obtain a sharper guaranteed convergence rate than the one presented in Stich and Karimireddy (2020). Note also that our theoretical guarantee matches with the bound in Arjevani et al. (2020) derived for strongly convex quadratic functions.

By selecting the threshold parameter as $\tau_{\text{th}} = 2(M - 1)$, it follows from Theorem 16 that the asynchronous SGD method achieves a convergence rate of

$$\mathbb{E}[F(\bar{x}_{\mathcal{K}}) - F^*] = \mathcal{O} \left(\frac{LM}{\mathcal{K}} + \frac{\sigma}{\sqrt{\mathcal{K}}} \right)$$

for convex functions and

$$\tilde{\mathcal{O}} \left(\exp \left(-\frac{\mu\mathcal{K}}{LM} \right) + \frac{\sigma^2}{\mathcal{K}} \right)$$

for strongly convex functions, which match the rates provided by Mishchenko et al. (2022).

4.2.1 COMPARISON OF OUR ANALYSIS WITH PRIOR WORK

Arjevani et al. (2020) analyzed the asynchronous SGD method for convex quadratic functions and provided the first tight convergence rate. However, their proof technique is based on the generating function method, which is only applicable to quadratic functions. To extend these results to general smooth functions, Stich and Karimireddy (2020) employed the perturbed iterate technique, which was originally developed by Mania et al. (2017) to analyze asynchronous optimization algorithms. In Theorem 16, we establish a convergence rate for convex problems that matches the one derived by Stich and Karimireddy (2020) while also providing a tighter guarantee for strongly convex problems. Next, we show how

to recover the convergence rate result for non-convex objective functions given in Stich and Karimireddy (2020). For L -smooth (possibly non-convex) functions F , the proof of Theorem 16 in Stich and Karimireddy (2020) shows that the iterates generated by asynchronous SGD satisfy

$$\begin{aligned} \mathbb{E}[F(\tilde{x}_{k+1}) - F^*] &\leq \mathbb{E}[F(\tilde{x}_k) - F^*] + \gamma^3 L^2 \tau_{\max} \sum_{\ell=(k-\tau_{\max})_+}^{k-1} \mathbb{E}[\|\nabla F(x_\ell)\|^2] \\ &\quad - \frac{\gamma(1-L\gamma)}{2} \mathbb{E}[\|\nabla F(x_k)\|^2] + \frac{(3\gamma^3 L^2 \tau_{\max} + \gamma^2 L)\sigma^2}{2}, \end{aligned} \quad (33)$$

where $\{\tilde{x}_k\}$ with $\tilde{x}_0 = x_0$ is a virtual sequence defined in the proof. Let

$$V_k = \mathbb{E}[F(\tilde{x}_k) - F^*], \quad W_k = \mathbb{E}[\|\nabla F(x_k)\|^2], \quad X_k = \frac{\gamma}{4} \mathbb{E}[\|\nabla F(x_k)\|^2],$$

and $e = (3\gamma^3 L^2 \tau_{\max} + \gamma^2 L)\sigma^2/2$. Then, we can rewrite (33) as

$$X_k + V_{k+1} \leq V_k + \gamma^3 L^2 \tau_{\max} \sum_{\ell=(k-\tau_{\max})_+}^{k-1} W_\ell - \frac{\gamma(1-2L\gamma)}{4} W_k + e,$$

which is on the form (11). Thus, according to Lemma 5, if

$$\gamma^3 L^2 \tau_{\max}^2 \leq \frac{\gamma(1-2L\gamma)}{4},$$

then $\sum_{k=0}^K X_k \leq V_0 + (K+1)e$ for $K \in \mathbb{N}$. This implies that

$$\sum_{k=0}^K \mathbb{E}[\|\nabla F(x_k)\|^2] \leq \frac{4(F(x_0) - F^*)}{\gamma} + 5\gamma L(K+1)\sigma^2, \quad \text{for } \gamma \in \left(0, \frac{1}{2L(\tau_{\max} + 1)}\right].$$

By setting

$$\gamma = \min \left\{ \frac{1}{2L(\tau_{\max} + 1)}, \sqrt{\frac{4(F(x_0) - F^*)}{5L(K+1)\sigma^2}} \right\},$$

which minimizes the right-hand side of the inequality above with respect to γ , we obtain

$$\min_{0 \leq k \leq K} \{\mathbb{E}[\|\nabla F(x_k)\|^2]\} = \mathcal{O} \left(\frac{L(\tau_{\max} + 1)(F(x_0) - F^*)}{K+1} + \sigma \sqrt{\frac{L(F(x_0) - F^*)}{K+1}} \right).$$

This convergence guarantee matches the bound derived in Stich and Karimireddy (2020).

Also focusing on non-convex problems, Koloskova et al. (2022) proposed a delay-adaptive step-size rule that eliminates the dependence of the convergence rate on τ_{\max} . According to the proof of Theorem 8 in Koloskova et al. (2022), the following inequality holds for the sequence generated by asynchronous SGD:

$$X_k + V_{k+1} \leq V_k + \gamma_k \tau_k L^2 \sum_{\ell=(k-\tau_k)_+}^{k-1} W_\ell - \frac{1}{4\gamma_k} W_k + \left(\gamma_k^2 L + \gamma_k L^2 \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2 \right) \sigma^2,$$

where $V_k = \mathbb{E}[F(x_k) - F^*]$, $X_k = \frac{\gamma_k}{2} \mathbb{E}[\|\nabla F(x_k)\|^2]$, and $W_k = \gamma_k^2 \mathbb{E}[\|\nabla F(x_{k-\tau_k})\|^2]$. Similar to the proof for Theorem 14, it follows from Lemma 11 that

$$\frac{1}{\sum_{k=0}^K \gamma_k} \sum_{k=0}^K \gamma_k \mathbb{E}[\|\nabla F(x_k)\|^2] \leq \frac{4(F(x_0) - F^*)}{\gamma(K+1)} + 6\gamma L\sigma^2, \quad \text{for } \gamma \in \left(0, \frac{1}{2L(\tau_{\text{th}} + 1)}\right].$$

By setting $\tau_{\text{th}} = 2\tau_{\text{ave}}$, the above convergence rate leads to an iteration complexity of

$$\mathcal{O}\left(\frac{\tau_{\text{ave}}}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right).$$

This bound is the same as the one provided in Koloskova et al. (2022). Alternatively, choosing $\tau_{\text{th}} = 2(M-1)$ results in

$$\mathcal{O}\left(\frac{M}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right),$$

which matches the iteration complexity derived in Mishchenko et al. (2022) using the perturbed iterate technique. While we recover the convergence guarantees presented in Mishchenko et al. (2022), recurrences of the form (10) and (11) did not appear in their proofs. The reason is that the analysis in Mishchenko et al. (2022) relies on upper bounds for the delayed terms that are expressed in terms of M , rather than on the delay values themselves.

4.3 Asynchronous Coordinate Update Methods

Many popular optimization algorithms, such as gradient descent, projected gradient descent, proximal-point, forward-backward splitting, Douglas-Rachford splitting, and the alternating direction method of multipliers (ADMM), can be viewed as special instances of the Krasnosel'skiĭ–Mann (KM) method for operators (Peng et al., 2016). The only way they differ is in their choice of operator. The KM method has the form

$$x_{k+1} = (1 - \gamma)x_k + \gamma T(x_k), \quad k \in \mathbb{N}_0, \quad (34)$$

where $\gamma \in (0, 1)$ is the step-size, and $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a nonexpansive operator, i.e.,

$$\|T(x) - T(y)\| \leq \|x - y\|, \quad \forall x, y \in \mathbb{R}^d.$$

Any vector $x^* \in \mathbb{R}^d$ satisfying $T(x^*) = x^*$ is called a *fixed point* of T and the KM method can be viewed as an algorithm for finding such a fixed point. Let $S = I_d - T$, where I_d is the identity operator on \mathbb{R}^d . Then, (34) can be rewritten as

$$x_{k+1} = x_k - \gamma S(x_k), \quad k \in \mathbb{N}_0. \quad (35)$$

This shows that the KM method can be interpreted as a taking a step of length γ in the opposite direction of S evaluated at the current iterate. For example, the gradient descent method for minimization of a L -smooth convex function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ can be formulated as the KM method with

$$S = \frac{2}{L} \nabla f, \quad \text{or equivalently,} \quad T = I_d - \frac{2}{L} \nabla f.$$

We represent x as $x = ([x]_1, \dots, [x]_m)$, where $[x]_i \in \mathbb{R}^{d_i}$ with d_1, \dots, d_m being positive integer numbers satisfying $d = d_1 + \dots + d_m$. We denote by $S_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ the i th block of S , so $S(x) = (S_1(x), \dots, S_m(x))$. Thus, we can rewrite (35) as

$$[x_{k+1}]_i = [x_k]_i - \gamma S_i(x_k), \quad i \in [m].$$

ARock is an algorithmic framework for parallelizing the KM method in an asynchronous fashion (Peng et al., 2016). In ARock, multiple agents (machines, processors, or cores) have access to a shared memory for storing the vector x , and are able to read and update x simultaneously without using locks. Conceptually, ARock lets each agent repeat the following steps:

- Read x from the shared memory without locks and save it in a local cache as \hat{x} ;
- Choose an index $i \in [m]$ uniformly at random and use \hat{x} to compute $S_i(\hat{x})$;
- Update component i of the shared x via

$$[x]_i \leftarrow [x]_i - \gamma S_i(\hat{x}).$$

Since the agents are being run independently without synchronization, while one agent is busy reading x and evaluating $S_i(x)$, other agents may repeatedly update the value stored in the shared memory. Therefore, the value \hat{x} read from the shared memory may differ from the value of x to which the update is made later. In other words, each agent can update the shared memory using possibly out-of-date information.

Algorithm 4 ARock

Input: $x_0 \in \mathbb{R}^d$, step-size $\gamma > 0$, number of iterations $K \in \mathbb{N}$

- 1: Initialize global counter $k \leftarrow 0$
 - 2: **while** $k < K$ **do**
 - 3: Read each position of shared memory denoted by \hat{x}_k
 - 4: Sample i_k from $\{1, \dots, m\}$ with equal probability $\frac{1}{m}$
 - 5: Set $[x_{k+1}]_{i_k} \leftarrow [x_k]_{i_k} - \gamma S_{i_k}(\hat{x}_k)$
 - 6: Set $k \leftarrow k + 1$
 - 7: **end while**
-

Algorithm 4 describes ARock. We assume that the write operation on line 5 is atomic, in the sense that the updated result will successfully appear in the shared memory by the end of the execution. In practice, this assumption can be enforced through compare-and-swap operations (Niu et al., 2011). Updating a scalar is a single atomic instruction on most modern hardware. Thus, the atomic write assumption in Algorithm 4 naturally holds when each block is a single scalar, i.e., $m = d$ and $d_i = 1$ for all $i \in [m]$.

We define one *iteration* of ARock as a modification on any block of x in the shared memory. A global counter k is introduced to track the total number of iterations so that x_k is the value of x in the shared memory after k iterations. We use i_k to denote the component that is updated at iteration k , and \hat{x}_k for the value of x that is used in the calculation of

$S_{i_k}(\cdot)$. Since every iteration changes one block of the shared memory and all writes are required to be atomic, we have

$$x_k = \hat{x}_k + \sum_{j \in J_k} (x_{j+1} - x_j), \quad (36)$$

where $J_k \subseteq \{0, \dots, k-1\}$. The set J_k contains the indices of the iterations during which a block of shared memory is updated between the time when \hat{x}_k is read and the time when x_k is written. Thus, the sum in (36) represents all updates of the shared vector that have occurred from the time that the agent i_k begins reading \hat{x}_k from memory until it finishes evaluating $S_{i_k}(\hat{x}_k)$ and writes the result to memory (see Peng et al. 2016 for a more detailed discussion). Note that individual blocks of the shared vector may be updated by multiple agents while it is in the process of being read by another agent. Therefore, the components of \hat{x}_k may have different ages and the vector \hat{x}_k may never actually exist in the shared memory during the execution of Algorithm 4. This phenomenon is known as *inconsistent read* (Liu and Wright, 2015).

To analyze the convergence of ARock, we need to make a few assumptions similar to Peng et al. (2016).

Assumption 7 *For Algorithm 4, the following properties hold:*

1. **(Pseudo-contractivity)** *The operator $T = I_d - S$ has a fixed point and is pseudo-contractive with respect to the Euclidean norm with contraction modulus c . That is, there exists $c \in (0, 1)$ such that*

$$\|T(x) - x^*\| \leq c\|x - x^*\|, \quad \forall x \in \mathbb{R}^d.$$

2. **(Bounded delay)** *There is a non-negative integer τ such that*

$$(k - \tau)_+ \leq \min\{j \mid j \in J_k\}, \quad \forall k \in \mathbb{N}_0.$$

3. **(Independence)** *Random variables i_k for $k = 0, \dots, K$ are independent of each other.*

Under Assumption 7.1, the serial KM iteration (35) with $\gamma \in (0, 1]$ converges to the fixed point x^* at a linear rate (Bertsekas and Tsitsiklis, 2015, Chapter 3). Assumption 7.2 guarantees that during any update cycle of an agent, the vector x in the shared memory is updated at most τ times by other agents. Therefore, no component of \hat{x}_k is older than τ for all $k \in \mathbb{N}_0$. The value of τ is an indicator of the degree of asynchrony in ARock. In practice, τ will depend on the number of agents involved in the computation. If all agents are working at the same rate, we would expect τ to be a multiple of the number of agents (Wright, 2015). Similar to Peng et al. (2016) and most results on asynchronous stochastic optimization for shared memory architecture (e.g., Niu et al. 2011; Liu et al. 2014; Liu and Wright 2015), we assume that the age of the components of \hat{x}_k is independent of the block i_k being updated at iteration k . However, this may not hold in practice if, for example, some blocks are more expensive to update than others (Leblond et al., 2018). This independence assumption can be relaxed using several techniques such as *before read labeling* (Mania et al., 2017), *after*

read labeling (Leblond et al., 2018), single coordinate consistent ordering (Cheung et al., 2020), and probabilistic models of asynchrony (Sun et al., 2017; Cannelli et al., 2020).

The following result shows that in the analysis of ARock, we can establish iterate relationships on the form (11).

Lemma 17 *Suppose Assumption 7 holds. Let $V_k = \mathbb{E}[\|x_k - x^*\|^2]$ and $W_k = \mathbb{E}[\|S(\hat{x}_k)\|^2]$ for $k \in \mathbb{N}_0$, where the expectation is over all choices of index i_k up to step k . Then, the iterates generated by Algorithm 4 satisfy*

$$\begin{aligned} V_{k+1} \leq & \left(1 - \frac{\gamma(1-c^2)}{m(1+\gamma(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}))}\right) V_k + \frac{2\gamma^2}{m\tau} \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right) \sum_{\ell=(k-\tau)_+}^{k-1} W_\ell \\ & - \frac{\gamma}{m} \left(1 - \gamma \left(1 + \frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right) W_k. \end{aligned}$$

Proof. See Appendix E.2. ■

Using this iterate relationship, the sequence result in Lemma 5 allows us to derive new convergence guarantees for ARock.

Theorem 18 *Suppose that Assumption 7 holds and that the step-size γ is set to*

$$\gamma = \frac{h}{1 + 5\left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)},$$

with $h \in (0, 1]$. Then, the iterates generated by Algorithm 4 satisfy

$$\mathbb{E}[\|x_k - x^*\|^2] \leq \left(1 - \frac{h(1-c^2)}{m(1+6\left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right))}\right)^k \|x_0 - x^*\|^2, \quad k \in \mathbb{N}_0. \quad (37)$$

Moreover, the algorithm reaches an accuracy of $\mathbb{E}[\|x_k - x^*\|^2] \leq \epsilon$ after

$$k \geq K_\epsilon = \frac{m(1+6\left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right))}{h(1-c^2)} \log\left(\frac{\|x_0 - x^*\|^2}{\epsilon}\right) \quad (38)$$

iterations.

Proof. See Appendix E.3. ■

Theorem 18 shows that for pseudo-contractive operators, ARock converges in expectation at a linear rate. To quantify how τ can affect the convergence of ARock, we define

$$\Gamma = \frac{\tau}{m} + \sqrt{\frac{\tau}{m}}.$$

Clearly, Γ is monotonically increasing in τ and is equal to zero for $\tau = 0$. The maximum allowable step-size and the linear rate of ARock depend on Γ . As Γ tends to infinity, the maximum allowable step-size decreases and approaches zero, while the convergence factor increases and approaches one. Therefore, ARock should take a smaller step-size for a larger

Γ , which will lead to a slower convergence rate. Note that if $\Gamma \approx 0$, or equivalently, $\tau \ll m$, then ARock with step-size $\gamma \approx 1$ has the linear rate

$$\rho = 1 - \frac{1 - c^2}{m},$$

which is precisely the rate for the serial stochastic KM method (Hannah and Yin, 2017). As discussed before, τ is related to the number of agents used in the algorithm. Therefore, the number of agents can be of the order of $o(m)$ without appreciably degrading the convergence rate of ARock.

The serial KM method (35) with constant step-size $\gamma_{\text{KM}} = h \in (0, 1]$ needs

$$k \geq \frac{1}{h(1 - c^2)} \log \left(\frac{\|x_0 - x^*\|^2}{\epsilon} \right)$$

iterations to satisfy $\|x_k - x^*\|^2 \leq \epsilon$ (Hannah and Yin, 2017). Each iteration of (35) requires computing all m blocks of S and updating the whole vector x . Thus, the overall complexity of the serial KM method, in terms of the total number of blocks updated to find an ϵ -optimal solution, is

$$K_{\text{KM}} = \frac{m}{h(1 - c^2)} \log \left(\frac{\|x_0 - x^*\|^2}{\epsilon} \right).$$

On the other hand, by Theorem 18, ARock with step-size $\gamma_{\text{ARock}} = h/(1 + 5\Gamma)$ performs

$$K_{\text{ARock}} = \frac{m(1 + 6\Gamma)}{h(1 - c^2)} \log \left(\frac{\|x_0 - x^*\|^2}{\epsilon} \right)$$

component updates to return an ϵ -optimal solution in average. In the case that $\Gamma \approx 0$, we have $\gamma_{\text{ARock}} \approx \gamma_{\text{KM}}$ and $K_{\text{ARock}} \approx K_{\text{KM}}$. Hence, as long as τ is bounded by $o(m)$, ARock can use the same step-size as the serial KM method and achieve the same iteration bound. Furthermore, since ARock runs on p agents in parallel, updates can occur roughly p times more frequently, leading to a near-linear speedup in the number of agents.

4.3.1 COMPARISON OF OUR ANALYSIS WITH PRIOR WORK

Peng et al. (2016) proposed ARock for finding fixed points of nonexpansive operators in an asynchronous parallel fashion. They proved that if T is nonexpansive and $S = I - T$ is quasi- μ -strongly monotone, then ARock with step-size

$$\gamma = \frac{1}{1 + \mathcal{O}\left(\frac{\tau^2}{\sqrt{m}}\right)}$$

converges linearly to a fixed point and achieves a linear speedup for $\tau \leq \mathcal{O}(m^{1/4})$.¹ However, we will show that using Lemma 5 in their proof allows us to improve the linear speed

1. The operator $S : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is quasi- μ -strongly monotone if it satisfies $\langle x - y, Sx \rangle \geq \mu \|x - y\|^2$ for all $x \in \mathbb{R}^d$ and $y \in \{z \in \mathbb{R}^d \mid Sz = 0\}$.

condition from $\tau \leq \mathcal{O}(m^{1/4})$ to $\tau \leq \mathcal{O}(m^{1/2})$. According to the proof of Theorem 4 in Peng et al. (2016), the iterates generated by ARock satisfy

$$V_{k+1} \leq \left(1 - \frac{\gamma\mu}{2m}\right) V_k + \frac{\gamma^2}{m^2} (\gamma\mu\tau + \sqrt{m}) \sum_{\ell=(k-\tau)_+}^{k-1} W_\ell - \frac{\gamma^2}{m} \left(\frac{1}{2\gamma} - 1 - \frac{\tau}{\sqrt{m}}\right) W_k,$$

where $V_k = \mathbb{E}[\|x_k - x^*\|^2]$ and $W_k = \mathbb{E}[\|S(\hat{x}_k)\|^2]$ for $k \in \mathbb{N}_0$. Since this inequality is on the form (11), it follows from Lemma 5 that if

$$\frac{\gamma^2}{m^2} (\gamma\mu\tau + \sqrt{m}) (2\tau) \leq \frac{\gamma^2}{m} \left(\frac{1}{2\gamma} - 1 - \frac{\tau}{\sqrt{m}}\right),$$

then V_k converges at a linear rate. This implies that

$$\mathbb{E}[\|x_k - x^*\|^2] \leq \left(1 - \frac{\gamma\mu}{2m}\right)^k \|x_0 - x^*\|^2, \quad \text{for } \gamma \in \left(0, \frac{1}{2\left(1 + \frac{\tau(\sqrt{\mu}+3)}{\sqrt{m}}\right)}\right).$$

Therefore, τ can be as large as $\mathcal{O}(m^{1/2})$ without affecting the maximum allowable step-size. Hence, using Lemma 5 in the proof of Theorem 4 in Peng et al. (2016) improves the upper bound on τ by a factor of $m^{1/4}$.

Hannah and Yin (2017) improved the results in Peng et al. (2016) by showing that for pseudo-contractive operators,

$$\gamma = \frac{1}{1 + \mathcal{O}\left(\frac{\tau}{\sqrt{m}}\right)}$$

guarantees the linear convergence of ARock and $\tau \leq \mathcal{O}(m^{1/2})$ ensures a linear speedup. Compared to the result presented in Hannah and Yin (2017), not only can Theorem 18 provide a larger value for the maximal allowable step-size, but it also improves the requirement for the linear speedup property from $\tau \leq \mathcal{O}(m^{1/2})$ to $\tau \leq o(m)$. The analysis in Hannah and Yin (2017) involves a recurrence of the form

$$V_{k+1} \leq qV_k + p \sum_{\ell=(k-\tau_k)_+}^k W_\ell - rW_k, \quad k \in \mathbb{N}_0,$$

with the same quantities for V_k and W_k as our analysis. However, our coefficients for q , p , and r are different, since we employ tighter upper bounds on the delayed terms. It is, in this case, the use of Lemmas 32 and 33 to bound $\mathbb{E}_k[\|x_k - \hat{x}_k\|^2]$ and $\mathbb{E}_k[\langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle]$, respectively, that are the keys to our improved results.

4.3.2 A SPECIAL CASE: ASYNCHRONOUS COORDINATE DESCENT METHOD

We now present a special case of ARock, namely the asynchronous coordinate descent algorithm for minimizing a class of composite objective functions. Specifically, we consider the problem

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad P(x) := F(x) + R(x), \quad (39)$$

where $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth, and $R : \mathbb{R}^d \rightarrow \mathbb{R}$ is separable in all coordinates, i.e.,

$$R(x) = \sum_{i=1}^m r_i([x]_i).$$

Here, each $r_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R} \cup \{+\infty\}$, $i \in [m]$, is a closed, convex, and extended real-valued function. The best known examples of separable regularizers include ℓ_1 norm, ℓ_2 norm square, and the indicator function of box constraints (Fercoq and Richtárik, 2015). The minimizer of Problem (39) is the unique fixed point of T_{prox} defined as

$$T_{\text{prox}}(x) = \text{prox}_{\frac{2}{\mu+L}R} \left(x - \frac{2}{\mu+L} \nabla F(x) \right).$$

The operator T_{prox} is contractive with contraction modulus

$$c = \frac{Q-1}{Q+1},$$

where $Q = L/\mu$ (Bertsekas, 2015). To solve (39), we apply ARock to $S = I_d - T_{\text{prox}}$. Then, the update rule of Algorithm 4 at the k th iteration becomes

$$[x_{k+1}]_{i_k} \leftarrow [x_k]_{i_k} - \gamma \left([\hat{x}_k]_{i_k} - \text{prox}_{\frac{2}{\mu+L}r_{i_k}} \left([\hat{x}_k]_{i_k} - \frac{2}{\mu+L} \nabla_{i_k} F(\hat{x}_k) \right) \right),$$

where $\nabla_i F(x)$ denotes the partial gradient of F with respect to $[x]_i$. According to Theorem 18, the iterates generated by ARock with step-size $\gamma = 1/(1+5\Gamma)$ satisfy

$$\mathbb{E}[\|x_k - x^*\|^2] \leq \left(1 - \frac{4Q}{m(1+6\Gamma)(Q+1)^2} \right)^k \|x_0 - x^*\|^2, \quad k \in \mathbb{N}_0.$$

If $\Gamma \approx 0$, or equivalently, $\tau \ll m$, then ARock using $\gamma \approx 1$ converges linearly at a rate of

$$\rho_{\text{ARock}} = 1 - \frac{4Q}{m(Q+1)^2}. \quad (40)$$

Note that when $m = 1$, the preceding linear rate reduces to

$$\rho_{\text{GD}} = \left(\frac{Q-1}{Q+1} \right)^2,$$

which is the best convergence rate for the gradient descent method applied to strongly convex optimization.

Remark 19 *Liu and Wright (2015) proposed an asynchronous coordinate descent algorithm for solving optimization problems of the form (39). They proved that the linear speedup is achievable if $\tau \leq \mathcal{O}(m^{1/4})$. For a special case of Problem (39) where $R(x) \equiv 0$, Liu et al. (2014) showed that the asynchronous coordinate descent method can enjoy the linear speedup if $\tau \leq \mathcal{O}(m^{1/2})$. In comparison with Liu and Wright (2015) and Liu et al. (2014), our requirement for the linear speedup property is $\tau \leq o(m)$ and, hence, allows a larger value for τ . Recently, Cheung et al. (2020) analyzed convergence of the asynchronous coordinate descent method for composite objective functions without assuming independence between i_k and \hat{x}_k . Their analysis guarantees the linear speedup for $\tau \leq \mathcal{O}(m^{1/2})$.*

Remark 20 *In Theorem 18, the linear convergence of ARock is given in terms of the expected quadratic distance from the iterates to the fixed point. Note however that the literature on coordinate descent algorithms usually establishes convergence results using coordinate-wise Lipschitz constants of the function F (see, e.g., Liu and Wright 2015; Liu et al. 2014; Nesterov 2012). This allows to provide larger step-sizes, which can lead to potentially better convergence bounds, especially in terms of the function values $P(x_k) - P^*$.*

4.4 A Lyapunov Approach to Analysis of Totally Asynchronous Iterations

Finally, we study iterations involving maximum norm pseudo-contractions under the general asynchronous model introduced by Bertsekas and Tsitsiklis (2015), which allows for heterogeneous and time-varying communication delays and update rates. Such iterations arise, for example, in algorithms for the solution of certain classes of linear equations, optimization problems and variational inequalities (Bertsekas and Tsitsiklis, 2015; Moallemi and Roy, 2010; Hale et al., 2017), distributed algorithms for averaging (Mehyar et al., 2007), power control algorithms for wireless networks (Feyzmahdavian et al., 2012), and reinforcement algorithms for solving discounted Markov decision processes (Zeng et al., 2020). We will demonstrate how the convergence results in Bertsekas and Tsitsiklis (2015) for maximum norm contractions can be derived and extended using Lemmas 1–3. This allows to unify and expand the existing results for partially and totally asynchronous iterations.

We consider iterative algorithms on the form

$$x_{k+1} = T(x_k), \quad k \in \mathbb{N}_0, \quad (41)$$

where $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a continuous mapping. This iteration aims to find a fixed point of T , that is, a vector $x^* \in \mathbb{R}^d$ satisfying $x^* = T(x^*)$. Similar to Subsection 4.3, we decompose the space \mathbb{R}^d as a Cartesian product of m subspaces:

$$\mathbb{R}^d = \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_m}, \quad d = \sum_{i=1}^m d_i.$$

Accordingly, we can partition any vector $x \in \mathbb{R}^d$ as $x = ([x]_1, \dots, [x]_m)$ with $[x]_i \in \mathbb{R}^{d_i}$, $i \in [m]$. We denote by $T_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ the i th component of T , so $T(x) = (T_1(x), \dots, T_m(x))$. Then, we rewrite (41) as

$$[x_{k+1}]_i = T_i([x_k]_1, \dots, [x_k]_m), \quad i \in [m]. \quad (42)$$

This iteration can be viewed as a network of m agents, each responsible for updating one of the m blocks of x so as to find a global fixed point of the operator T .

Let us fix some norms $\|\cdot\|_i$ for the spaces \mathbb{R}^{d_i} . The *block-maximum norm* $\|\cdot\|_{b,\infty}^w$ on \mathbb{R}^d is defined as

$$\|x\|_{b,\infty}^w = \max_{i \in [m]} w_i \| [x]_i \|_i,$$

where w_i are positive constants. Note that when $d_i = 1$ for each $i \in [m]$, the block-maximum norm reduces to the maximum norm

$$\|x\|_\infty^w = \max_{i \in [m]} w_i |[x]_i|.$$

The following definition introduces *pseudo-contractive* mappings with respect to the block-maximum norm.

Definition 21 *A mapping $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called a pseudo-contraction with respect to the block-maximum norm if it has a fixed-point x^* and the property*

$$\|T(x) - x^*\|_{b,\infty}^w \leq c \|x - x^*\|_{b,\infty}^w, \quad \forall x \in \mathbb{R}^d.$$

Here, c , called the contraction modulus of T , is a constant belonging to $(0, 1)$.

Pseudo-contractions have at most one fixed point (Bertsekas and Tsitsiklis, 2015, Proposition 3.1.2). Note that we follow Bertsekas and Tsitsiklis (2015) and include the existence of a fixed point in the definition of pseudo-contractions. Examples of iterative algorithms that involve pseudo-contractions with respect to the block-maximum norm can be found in (Bertsekas and Tsitsiklis, 2015, Chapter 3). An important property of pseudo-contractions is that the iterates produced by (42) converge at a linear rate to the fixed point (Bertsekas and Tsitsiklis, 2015, Proposition 3.1.2). More precisely, using the Lyapunov function $V_k = \|x_k - x^*\|_{b,\infty}^w$, one can show that the sequence $\{x_k\}$ generated by (42) with initial vector $x_0 \in \mathbb{R}^d$ satisfies $V_{k+1} \leq cV_k$ for all $k \in \mathbb{N}_0$. This implies that

$$\|x_k - x^*\|_{b,\infty}^w \leq c^k \|x_0 - x^*\|_{b,\infty}^w, \quad k \in \mathbb{N}_0.$$

The algorithm described by (42) is *synchronous* in the sense that all agents update their states at the same time and have instantaneous access to the states of all other agents. Synchronous execution is possible if there are no communication failures in the network and if all agents operate in sync with a global clock. In practice, these requirements are hard to satisfy: local clocks in different agents tend to drift, global synchronization mechanisms are complex to implement and carry substantial execution overhead, and the communication latency between agents can be significant and unpredictable. Insisting on synchronous operation in an inherently asynchronous environment forces agents to spend a significant time idle, waiting for the slowest agents (perhaps due to lower processing power or higher workload per iteration) to complete its work.

In an *asynchronous* implementation of the iteration (42), each agent can update its state at its own pace, using possibly outdated information from the other agents. This leads to iterations on the form

$$[x_{k+1}]_i = \begin{cases} T_i([x_{s_{i1,k}}]_1, \dots, [x_{s_{im,k}}]_m), & k \in \mathcal{K}_i, \\ [x_k]_i, & k \notin \mathcal{K}_i. \end{cases} \quad (43)$$

Here, \mathcal{K}_i is the set of times when agent i executes an update, and $s_{ij,k}$ is the time at which the most recent version of $[x]_j$ available to agent i at the time k was computed. The sets \mathcal{K}_i need not be known to all the agents. Thus, there is no requirement for a shared global clock. Since the agents use only values computed in the previous iterations, $s_{ij,k} \leq k$ for all $k \in \mathbb{N}_0$. We can view the value

$$\tau_{ij,k} := k - s_{ij,k},$$

as the communication delay from agent j to agent i at time k . Clearly, the synchronous iteration (42) is a special case of (43) where $\mathcal{K}_i = \mathbb{N}_0$ and $\tau_{ij,k} = 0$ for all $i, j \in [m]$ and all $k \in \mathbb{N}_0$.

Based on the assumptions on communication delays and update rates, Bertsekas and Tsitsiklis (2015) introduced a classification of asynchronous algorithms as either *totally asynchronous* or *partially asynchronous*. The totally asynchronous model is characterized by the following assumption.

Assumption 8 (Total Asynchronism) *For the asynchronous iteration (43), the following properties hold:*

1. *The sets \mathcal{K}_i are infinite subsets of \mathbb{N}_0 for each $i \in [m]$.*
2. *$\lim_{k \rightarrow \infty} s_{ij,k} = +\infty$ for all $i, j \in [m]$.*

Assumption 8.1 guarantees that no agent ceases to execute its update, while Assumption 8.2 guarantees that outdated information about the agent updates is eventually purged from the computation (compare the discussion after Lemma 2). Under total asynchronism, the communication delays $\tau_{ij,k}$ can become unbounded as k increases. This is the main difference with partially asynchronous iterations, where delays are bounded; in particular, the following assumption holds.

Assumption 9 (Partial Asynchronism) *For the asynchronous iteration (43), there exist non-negative integers B and D such that the following conditions hold:*

1. *At least one of the elements of the set $\{k, k + 1, \dots, k + B\}$ belongs to \mathcal{K}_i for each $i \in [m]$ and all $k \in \mathbb{N}_0$.*
2. *$k - D \leq s_{ij,k} \leq k$ for all $i, j \in [m]$ and all $k \in \mathcal{K}_i$.*
3. *$s_{ii,k} = k$ for all $i \in [m]$ and $k \in \mathcal{K}_i$.*

Assumptions 9.1 and 9.2 ensure that the time interval between updates executed by each agent and the communication delays are bounded by B and D , respectively. This means that no agent waits an arbitrarily long time to compute or to receive a message from another agent. Assumption 9.3 states that agent i always uses the latest version of its own component $[x]_i$. Note that when $B = D = 0$, the asynchronous iteration (43) under partial asynchronism reduces to the synchronous iteration (42).

We will now present a uniform analysis of the iteration (43) involving block-maximum norm pseudo-contractions under both partial and total asynchronism, and study its convergence rate under different assumptions on the communication delays and update rates. To this end, we introduce τ_k to represent the maximum age of the outdated information being used to update blocks at global time $k \in \mathbb{N}_0$. Specifically, we define τ_k as

$$\tau_k := k - \min_{i \in [m]} \min_{j \in [m]} s_{ij, t_i(k)}, \quad (44)$$

where $t_i(k)$ is the most recent update time of agent i at $k \in \mathbb{N}_0$, i.e.,

$$t_i(k) = \max\{\kappa \mid \kappa \leq k \wedge \kappa \in \mathcal{K}_i\}.$$

In this way, if $k \in \mathcal{K}_i$ then $t_i(k) = k$, and if $k \in (\kappa^-, \kappa^+)$ for two consecutive elements κ^- and κ^+ of \mathcal{K}_i , then $t_i(k) = \kappa^-$. For simplicity, we assume that $0 \in \mathcal{K}_i$ for each i , so that $t_i(k)$ is well-defined for all $k \in \mathbb{N}_0$.

The next result shows that for the asynchronous iteration (43), if $V_k = \|x_k - x^*\|_{b,\infty}^w$ is used as a candidate Lyapunov function (similarly to the convergence analysis for the synchronous case), we can establish iterate relationships on the form (10).

Lemma 22 *Suppose that T is pseudo-contractive with respect to the block-maximum norm with contraction modulus c . Let $V_k = \|x_k - x^*\|_{b,\infty}^w$. Then, the iterates $\{x_k\}$ generated by the asynchronous iteration (43) satisfy*

$$V_{k+1} \leq c \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0,$$

where τ_k is defined in (44).

Proof. See Appendix F.1. ■

We apply Lemma 2 to show that for pseudo-contractions with respect to the block-maximum norm, the asynchronous iteration (43) converges asymptotically to the fixed point under total asynchronism.

Theorem 23 *Let Assumption 8 hold. Suppose that T is pseudo-contractive with respect to the block-maximum norm. Then, the sequence $\{x_k\}$ generated by the asynchronous iteration (43) converges asymptotically to the unique fixed point of T .*

Proof. See Appendix F.2. ■

While convergent synchronous algorithms may diverge in the face of asynchronism, Theorem 23 shows that pseudo-contractive mappings in the block-maximum norm can tolerate arbitrarily large communication delays and update intervals satisfying Assumption 8. In addition, as the next result shows, these iterations admit an explicit convergence rate bound when they are executed in a partially asynchronous fashion.

Theorem 24 *Let Assumption 9 hold. Suppose that T is pseudo-contractive with respect to the block-maximum norm with contraction modulus c . Then, the sequence $\{x_k\}$ generated by the asynchronous iteration (43) satisfies*

$$\|x_k - x^*\|_{b,\infty}^w \leq c^{\frac{k}{B+D+1}} \|x_0 - x^*\|_{b,\infty}^w, \quad k \in \mathbb{N}_0.$$

Proof. See Appendix F.3. ■

According to Theorem 24, the asynchronous iteration (43) involving block-maximum norm pseudo-contractions remains linearly convergent for bounded communication delays and update rates. Note that $c^{\frac{1}{B+D+1}}$ is monotonically increasing with B and D , and approaches one as either B or D tends to infinity. Hence, the guaranteed convergence rate of (43) slows down as either the delays increase in magnitude or agents execute less frequently. The convergence rate is directly related to the number of iterations required for the algorithm to converge. According to Theorem 24, the asynchronous iteration (43) needs

$$k \geq \frac{B+D+1}{\log(1/c)} \log \left(\frac{\|x_0 - x^*\|_{b,\infty}^w}{\epsilon} \right)$$

iterations to satisfy $\|x_k - x^*\|_{b,\infty}^w \leq \epsilon$. We can see that the time to reach a fixed target accuracy deteriorates linearly with both B and D .

We now use Lemma 3 to develop a result that provides guaranteed convergence rates for the asynchronous algorithm (43) under a rather broad family of communication delays and update rates, in between the partially and totally asynchronous models.

Theorem 25 *Suppose that T is pseudo-contractive with respect to the block-maximum norm with contraction modulus c . Suppose also that there exists a function $\Lambda : \mathbb{R} \rightarrow \mathbb{R}$ such that the following conditions hold:*

- (i) $\Lambda(0) = 1$.
- (ii) Λ is non-increasing.
- (iii) $\lim_{k \rightarrow +\infty} \Lambda(k) = 0$ and

$$c\Lambda(k - \tau_k) \leq \Lambda(k + 1), \quad k \in \mathbb{N}_0,$$

where τ_k is defined in (44).

Then, the iterates $\{x_k\}$ generated by the asynchronous iteration (43) satisfy

$$\|x_k - x^*\|_{b,\infty}^w \leq \Lambda(k) \|x_0 - x^*\|_{b,\infty}^w, \quad k \in \mathbb{N}.$$

Proof. See Appendix F.4. ■

According to Theorem 25, any function Λ satisfying conditions (i)—(iii) can be used to estimate the convergence rate of the asynchronous iteration (43). Condition (iii) implies that the admissible choices for Λ depend on τ_k . This means that the rate at which the nodes execute their updates as well as the way communication delays tend to large both affect the guaranteed convergence rate of (43). For example, one can verify that under partial asynchronism, the function

$$\Lambda(t) = c^{\frac{t}{B+D+1}}$$

satisfies conditions (i)—(iii). In the following example, we use Theorem 25 and Corollary 4 to establish convergence rates for a particular class of totally asynchronous iterations, where the communication delays can grow unbounded at a linear rate.

Example 1 *For the asynchronous iteration (43), suppose that $\mathcal{K}_i = \mathbb{N}_0$ for each $i \in [m]$. Suppose also that there exist scalars $\alpha \in (0, 1)$ and $\beta \geq 0$ such that*

$$(1 - \alpha)k - \beta \leq s_{ij,k}, \quad i, j \in [m], k \in \mathcal{K}_i. \quad (45)$$

Note that (45) implies that the communication delays $\tau_{ij,k}$ belong to the interval $[0, \alpha k + \beta]$ and may therefore grow unbounded. In this example, as $t_i(k) = k$, we have

$$(1 - \alpha)k - \beta \leq \min_{i \in [m]} \min_{j \in [m]} s_{ij, t_i(k)}.$$

Thus, by Lemma 22, the iterates generated by the asynchronous iteration (43) satisfy

$$V_{k+1} \leq qV_k + p \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0,$$

with $V_k = \|x_k - x^*\|_{b,\infty}^w$, $q = 0$, $p = c$, and $\tau_k \leq \alpha k + \beta$. Since $q + p = c < 1$, it follows from Corollary 4 that the function

$$\Lambda(t) = \left(\frac{\alpha t}{1 - \alpha + \beta} + 1 \right)^{-\eta},$$

with $\eta = \ln(c)/\ln(1 - \alpha)$ satisfies conditions (i)–(iii) of Theorem 25. Therefore,

$$\|x_k - x^*\|_{b,\infty}^w \leq \left(\frac{\alpha k}{1 - \alpha + \beta} + 1 \right)^{-\eta} \|x_0 - x^*\|_{b,\infty}^w, \quad k \in \mathbb{N}_0.$$

We can see that under unbounded communication delays satisfying (45), the convergence rate of the asynchronous iteration (43) is of order $\mathcal{O}(1/k^\eta)$.

Although pseudo-contractions in the block-maximum norm converge when executed in a totally asynchronous manner, we note that in many applications, it is rare that the associated fixed-point iterations are maximum norm contractions. For instance, the gradient descent iterations for unconstrained optimization problems are maximum norm contractions if the Hessian matrix of the cost function is diagonally dominant (Bertsekas and Tsitsiklis, 2015, Section 3.1.3). The diagonal dominance assumption is quite strong and violated even by some strongly convex quadratic objective functions (Bertsekas and Tsitsiklis, 2015, Example 6.3.1).

Remark 26 Bertsekas and Tsitsiklis (2015) proved that contractions in the block-maximum norm converge when executed in a totally asynchronous manner. However, they did not quantify how bounds on the communication delays and update rates of agents affect the convergence rate of the iterates. There are very few results in the literature on convergence rates of asynchronous iterations involving block-maximum norm contractions. Notable exceptions are the works of Bertsekas and Tsitsiklis (1989) and Zeng et al. (2020), where the convergence rate of iterates was estimated under partial asynchronism. Theorems 23, 24 and 25 demonstrate that not only do we recover the asymptotic convergence results in Bertsekas and Tsitsiklis (2015) for maximum norm pseudo-contractions, but we also provide explicit bounds on the convergence rate of asynchronous iterations for various classes of bounded and unbounded communication delays and update rates.

5. Conclusions

We introduced a number of novel sequence results for asynchronous iterations that appear in the analysis of parallel and asynchronous algorithms. In contrast to previous analysis frameworks, which have used conservative bounds for the effects of asynchrony, our results attempt to capture the inherent structure in the asynchronous iterations. The results balance simplicity, applicability and power, and provide explicit bounds on how the amount of asynchrony affects the guaranteed convergence rates. To demonstrate the potential of

the sequence results, we illustrated how they can be used to improve our theoretical understanding of several important classes of asynchronous optimization algorithms. First, we derived better iteration complexity bounds for the proximal incremental aggregated gradient method, reducing the dependence of the convergence times on the maximum delay and problem condition number. Second, we provided tighter guarantees for the asynchronous stochastic gradient descent method that depend on the average delay rather than the maximal delay. Third, we gave an improved analysis of the ARock framework for asynchronous block-coordinate updates of Krasnosel'skiĭ–Mann iterations, proving a larger range of admissible step-sizes, faster convergence rates and better scaling properties with respect to the number of parallel computing elements. Finally, we gave a uniform treatment of asynchronous iterations involving block-norm contractions under partial and (several versions of) total asynchronism.

Appendix A. Proofs for Section 2

Lemma 27 *Suppose f is L -smooth and μ -strongly convex. Consider the delayed gradient method (3) with $\gamma \in \left(0, \frac{2}{\mu+L}\right]$ and let ω_k be the perturbation term defined in (5). Then,*

$$\omega_k \leq (\gamma^4 L^4 \tau^2 + 2\gamma^2 L^2 \tau) \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|^2\}.$$

Proof. Applying the Cauchy-Schwarz inequality, we obtain

$$\omega_k \leq 2\gamma \|\nabla f(x_k) - \nabla f(x_{k-\tau})\| \|x_k - \gamma \nabla f(x_k) - x^*\| + \gamma^2 \|\nabla f(x_k) - \nabla f(x_{k-\tau})\|^2.$$

We use L -smoothness of f to bound $\|\nabla f(x_k) - \nabla f(x_{k-\tau})\|$. Specifically,

$$\begin{aligned} \|\nabla f(x_k) - \nabla f(x_{k-\tau})\| &\leq L \|x_k - x_{k-\tau}\| \stackrel{(3)}{=} L \left\| \sum_{\ell=(k-\tau)_+}^{k-1} \gamma \nabla f(x_{\ell-\tau}) \right\| \\ &\leq \gamma L \sum_{\ell=(k-\tau)_+}^{k-1} \|\nabla f(x_{\ell-\tau})\|, \end{aligned}$$

where the last inequality follows from the triangle inequality. Since $\nabla f(x^*) = 0$, we have

$$\|\nabla f(x_k) - \nabla f(x_{k-\tau})\| \leq \gamma L^2 \sum_{\ell=(k-\tau)_+}^{k-1} \|x_{\ell-\tau} - x^*\| \leq \gamma L^2 \tau \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|\}.$$

From (1), $\|x_k - \gamma \nabla f(x_k) - x^*\| \leq \|x_k - x^*\|$ holds for the proposed range of γ . Thus,

$$\begin{aligned} \omega_k &\leq 2\gamma^2 L^2 \tau \|x_k - x^*\| \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|\} + \gamma^4 L^4 \tau^2 \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|^2\} \\ &\leq (\gamma^4 L^4 \tau^2 + 2\gamma^2 L^2 \tau) \max_{(k-2\tau)_+ \leq \ell \leq k} \{\|x_\ell - x^*\|^2\}. \end{aligned}$$

■

Lemma 28 *Suppose f is L -smooth and convex. Then, the iterates generated by (3) satisfy*

$$f(x_k) - f^* + \frac{1}{2L} \|\nabla f(x_{k-\tau})\|^2 - \frac{\gamma^2 L \tau}{2} \sum_{\ell=(k-\tau)_+}^{k-1} \|\nabla f(x_{\ell-\tau})\|^2 \leq \langle \nabla f(x_{k-\tau}), x_k - x^* \rangle.$$

Proof. According to Theorem 2.1.5 in Nesterov (2013), we have

$$\begin{aligned} f(x_k) &\leq f(x_{k-\tau}) + \langle \nabla f(x_{k-\tau}), x_k - x_{k-\tau} \rangle + \frac{L}{2} \|x_k - x_{k-\tau}\|^2 \\ &\leq f^* + \langle \nabla f(x_{k-\tau}), x_k - x^* \rangle - \frac{1}{2L} \|\nabla f(x_{k-\tau})\|^2 + \frac{L}{2} \|x_k - x_{k-\tau}\|^2. \end{aligned} \quad (46)$$

We can bound the final term by

$$\|x_k - x_{k-\tau}\|^2 \stackrel{(3)}{=} \left\| \sum_{\ell=(k-\tau)_+}^{k-1} \gamma \nabla f(x_{\ell-\tau}) \right\|^2 \leq \gamma^2 \tau \sum_{\ell=(k-\tau)_+}^{k-1} \|\nabla f(x_{\ell-\tau})\|^2, \quad (47)$$

where the inequality follows from the convexity of $\|\cdot\|^2$. Substituting (47) into (46), and rearranging the terms conclude the proof. \blacksquare

Appendix B. Proofs for Section 3

This section provides the proofs for the results presented in Section 3.

B.1 Proof of Lemma 2

The proof idea is inspired by the convergence theorem for totally asynchronous iterations (Bertsekas and Tsitsiklis, 2015, Proposition 6.2.1). We first use perfect induction to prove that

$$V_k \leq V_0, \quad k \in \mathbb{N}_0. \quad (48)$$

Since V_0 satisfies (48), the induction hypothesis is true for $k = 0$. Now, assume that (48) holds for all k up to some $K \in \mathbb{N}_0$. This means that $V_K \leq V_0$ and

$$\max_{(K-\tau_K)_+ \leq \ell \leq K} V_\ell \leq V_0.$$

Thus, from (13), we have

$$\begin{aligned} V_{K+1} &\leq (q+p)V_0 \\ &\leq V_0, \end{aligned}$$

where the second inequality follows from the fact that $q+p \in (0, 1)$. Therefore, $V_{K+1} \leq V_0$ and, hence, the induction proof is complete. Next, we prove that V_k converges to zero as $k \rightarrow +\infty$. To this end, we use induction to show that for each $m \in \mathbb{N}_0$, there exists $K_m \in \mathbb{N}_0$ such that

$$V_k \leq (q+p)^m V_0, \quad \forall k \geq K_m. \quad (49)$$

From (48), $V_k \leq (q+p)^0 V_0$ for all $k \geq 0$. Thus, the induction hypothesis holds for $m = 0$ (with $K_0 = 0$). Assuming (49) is true for a given m , we will show that there exists $K_{m+1} \in \mathbb{N}_0$ such that $V_k \leq (q+p)^{m+1} V_0$ for all $k \geq K_{m+1}$. From (14), one can find a sufficiently large $\bar{K}_m \geq K_m$ such that $k - \tau_k \geq K_m$ for $k \geq \bar{K}_m$. Since (49) holds by induction, we have

$$\max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell \leq (q+p)^m V_0, \quad \forall k \geq \bar{K}_m.$$

It follows from (13) that

$$\begin{aligned} V_{k+1} &\leq (q+p)(q+p)^m V_0 \\ &= (q+p)^{m+1} V_0, \quad \forall k \geq \bar{K}_m, \end{aligned}$$

which implies that

$$V_k \leq (q+p)^{m+1} V_0, \quad \forall k \geq \bar{K}_m + 1.$$

Set $K_{m+1} = \bar{K}_m + 1$. The induction proof is complete. In summary, we conclude that for each m , there exists K_m such that $V_k \leq (q+p)^m V_0$ for all $k \geq K_m$. Since $q+p < 1$, $(q+p)^m$ approaches zero as $m \rightarrow +\infty$. Hence, the sequence V_k asymptotically converges to zero.

B.2 Proof of Lemma 3

We will show by induction that

$$V_k \leq \Lambda(k) V_0, \quad k \in \mathbb{N}_0. \quad (50)$$

Since $\Lambda(0) = 1$, the induction hypothesis is true for $k = 0$. Assume for induction that (50) holds for all k up to some K . From (15), we have

$$\begin{aligned} V_{K+1} &\leq q\Lambda(K)V_0 + p \max_{(K-\tau_K)_+ \leq \ell \leq K} \Lambda(\ell)V_0 \\ &\leq q\Lambda(K)V_0 + p \max_{K-\tau_K \leq \ell \leq K} \Lambda(\ell)V_0. \end{aligned}$$

Since Λ is non-increasing on \mathbb{R} , we obtain $V_{K+1} \leq (q+p)\Lambda(K-\tau_K)V_0$. It follows from (16) that $V_{K+1} \leq \Lambda(K+1)V_0$. Therefore, (50) holds for $k = K+1$. The induction proof is complete.

B.3 Proof of Lemma 5

Let us define $W_k = 0$ for $k \in \{\dots, -2, -1\}$. Dividing both sides of (18) by Q_{k+1} , summing from $k = 0$ to K , and then using telescoping cancellation, we have

$$\begin{aligned} \sum_{k=0}^K \frac{X_k}{Q_{k+1}} + \frac{V_{K+1}}{Q_{K+1}} &\leq V_0 + \sum_{k=0}^K \sum_{\ell=(k-\tau_k)_+}^k \frac{p_k W_\ell}{Q_{k+1}} - \sum_{k=0}^K \frac{r_k W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}} \\ &= V_0 + \sum_{k=0}^K \sum_{\ell=k-\tau_k}^k \frac{p_k W_\ell}{Q_{k+1}} - \sum_{k=0}^K \frac{r_k W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}. \end{aligned}$$

Since $\tau_k \leq \tau$ for all $k \in \mathbb{N}_0$, we obtain

$$\begin{aligned}
 \sum_{k=0}^K \frac{X_k}{Q_{k+1}} + \frac{V_{K+1}}{Q_{K+1}} &\leq V_0 + \sum_{k=0}^K \sum_{\ell=k-\tau}^k \frac{p_k W_\ell}{Q_{k+1}} - \sum_{k=0}^K \frac{r_k W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}} \\
 &= V_0 + \sum_{\ell=0}^{\tau} \sum_{k=-\ell}^{K-\ell} \frac{p_{k+\ell} W_k}{Q_{k+\ell+1}} - \sum_{k=0}^K \frac{r_k W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}} \\
 &\leq V_0 + \sum_{\ell=0}^{\tau} \sum_{k=0}^K \frac{p_{k+\ell} W_k}{Q_{k+\ell+1}} - \sum_{k=0}^K \frac{r_k W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}} \\
 &= V_0 - \sum_{k=0}^K \left(\frac{r_k}{Q_{k+1}} - \sum_{\ell=0}^{\tau} \frac{p_{k+\ell}}{Q_{k+\ell+1}} \right) W_k + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}, \tag{51}
 \end{aligned}$$

where the second inequality comes from our assumption that $W_k \geq 0$ for $k \geq 0$ and $W_k = 0$ for $k < 0$. We are now ready to prove Lemma 5.

1. In this case, $q_k = 1$ and, hence, $Q_k = 1$ for each k . Thus, (51) simplifies to

$$\sum_{k=0}^K X_k + V_{K+1} \leq V_0 - \sum_{k=0}^K \left(r_k - \sum_{\ell=0}^{\tau} p_{k+\ell} \right) W_k + \sum_{k=0}^K e_k.$$

The assumption that

$$\sum_{\ell=0}^{\tau} p_{k+\ell} \leq r_k$$

holds for every $k \in \mathbb{N}_0$ implies that

$$\sum_{k=0}^K X_k + V_{K+1} \leq V_0 + \sum_{k=0}^K e_k.$$

Since $\{X_k\}$ and $\{V_k\}$ are non-negative sequences, it follows that

$$\sum_{k=0}^K X_k \leq V_0 + \sum_{k=0}^K e_k \quad \text{and} \quad V_{K+1} \leq V_0 + \sum_{k=0}^K e_k$$

hold for $K \in \mathbb{N}_0$.

2. In this case, $p_k = p$ and $r_k = r$. Thus, (51) simplifies to

$$\sum_{k=0}^K \frac{X_k}{Q_{k+1}} + \frac{V_{K+1}}{Q_{K+1}} \leq V_0 - \sum_{k=0}^K \left(\frac{r}{Q_{k+1}} - \sum_{\ell=0}^{\tau} \frac{p}{Q_{k+\ell+1}} \right) W_k + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}. \tag{52}$$

For each $k \in \mathbb{N}_0$ and $\ell \in [0, \tau]$, we have

$$\begin{aligned}
 Q_{k+\ell+1} &= Q_{k+1} \prod_{s=0}^{\ell-1} q_{k+s+1} \\
 &\geq Q_{k+1} q^\ell,
 \end{aligned}$$

where the inequality uses $q_k \geq q$ for $k \in \mathbb{N}_0$. Since $q \in (0, 1]$, we have $q^\ell \geq q^\tau$ for $\ell \in [0, \tau]$, implying that $Q_{k+\ell+1} \geq Q_{k+1}q^\tau$. Combining this inequality with (52) yields

$$\sum_{k=0}^K \frac{X_k}{Q_{k+1}} + \frac{V_{K+1}}{Q_{K+1}} \leq V_0 - (r - p(\tau + 1)q^{-\tau}) \sum_{k=0}^K \frac{W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}. \quad (53)$$

By Bernoulli's inequality, i.e., $(1 + x)^n \geq 1 + nx$ for any $n \in \mathbb{N}_0$ and any $x > -1$, we have

$$\begin{aligned} q^\tau &= (1 - (1 - q))^\tau \\ &\geq 1 - (1 - q)\tau. \end{aligned} \quad (54)$$

The assumption that $2\tau + 1 \leq 1/(1 - q)$ implies that $(1 - q) \leq 1/(2\tau + 1)$ and thereby

$$1 - (1 - q)\tau \geq \frac{\tau + 1}{2\tau + 1}.$$

By (54), it therefore holds that

$$q^\tau \geq \frac{\tau + 1}{2\tau + 1},$$

or equivalently, $q^{-\tau}(\tau + 1) \leq (2\tau + 1)$. From (53), we then have

$$\sum_{k=0}^K \frac{X_k}{Q_{k+1}} + \frac{V_{K+1}}{Q_{K+1}} \leq V_0 - (r - (2\tau + 1)p) \sum_{k=0}^{K-1} \frac{W_k}{Q_{k+1}} + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}.$$

The assumption that $2\tau + 1 \leq r/p$ allows us to drop the second term on the right-hand side. Thus,

$$\sum_{k=0}^K \frac{X_k}{Q_{k+1}} + \frac{V_{K+1}}{Q_{K+1}} \leq V_0 + \sum_{k=0}^K \frac{e_k}{Q_{k+1}}, \quad K \in \mathbb{N},$$

which concludes the proof.

Appendix C. Proofs for Subsection 4.1

This section provides the proofs for the results presented in Subsection 4.1. We first state two key lemmas which establish important recursions for the iterates generated by PIAG.

Lemma 29 *Suppose Assumptions 1 and 2 hold. Let $\{x_k\}$ be the sequence generated by Algorithm 2. Then, for any $x^* \in \mathcal{X}^*$ and every $k \in \mathbb{N}_0$, we have*

$$\begin{aligned} P(x_{k+1}) - P^* + \frac{1}{2\gamma} \|x_{k+1} - x^*\|^2 &\leq \frac{1}{2\gamma} \|x_k - x^*\|^2 + \frac{L(\tau + 1)}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_\ell\|^2 \\ &\quad - \frac{1}{2\gamma} \|x_{k+1} - x_k\|^2. \end{aligned} \quad (55)$$

Proof. At iteration $k \in \mathbb{N}_0$, the update rule in Algorithm 2 is of the form

$$g_k = \frac{1}{n} \sum_{i=1}^n \nabla f_i(x_{s_{i,k}}),$$

$$x_{k+1} = \text{prox}_{\gamma R}(x_k - \gamma g_k).$$

We write the proximal mapping more explicitly as

$$x_{k+1} = \underset{u \in \mathbb{R}^d}{\text{argmin}} \left\{ \frac{1}{2} \|u - (x_k - \gamma g_k)\|^2 + \gamma R(u) \right\}. \quad (56)$$

From the first-order optimality condition for the point u in the minimization problem (56), there is a sub-gradient $\xi_{k+1} \in \partial R(x_{k+1})$ such that

$$\langle x_{k+1} - x_k + \gamma(g_k + \xi_{k+1}), x^* - x_{k+1} \rangle \geq 0. \quad (57)$$

Since R is convex, we have

$$\begin{aligned} R(x_{k+1}) &\leq R(x^*) + \langle \xi_{k+1}, x_{k+1} - x^* \rangle \\ &\stackrel{(57)}{\leq} R(x^*) + \left\langle \frac{x_k - x_{k+1}}{\gamma} - g_k, x_{k+1} - x^* \right\rangle \\ &= R(x^*) + \frac{1}{2\gamma} (\|x_k - x^*\|^2 - \|x_{k+1} - x_k\|^2 - \|x_{k+1} - x^*\|^2) - \langle g_k, x_{k+1} - x^* \rangle, \end{aligned} \quad (58)$$

where the last equality follows from the fact

$$2\langle a - b, b - c \rangle = \|a - c\|^2 - \|a - b\|^2 - \|b - c\|^2, \quad a, b, c \in \mathbb{R}^d.$$

By Assumption 2, each component function f_i is L_i -smooth. Thus, from Lemma 1.2.3 in Nesterov (2013), we have

$$\begin{aligned} f_i(x_{k+1}) &\leq f_i(x_{s_{i,k}}) + \langle \nabla f_i(x_{s_{i,k}}), x_{k+1} - x_{s_{i,k}} \rangle + \frac{L_i}{2} \|x_{k+1} - x_{s_{i,k}}\|^2 \\ &\leq f_i(x^*) + \langle \nabla f_i(x_{s_{i,k}}), x_{k+1} - x^* \rangle + \frac{L_i}{2} \|x_{k+1} - x_{s_{i,k}}\|^2, \end{aligned}$$

where the second inequality follows from the convexity of f_i . Dividing both sides of the above inequality by n and then summing from $i = 1$ to n , we obtain

$$F(x_{k+1}) \leq F(x^*) + \langle g_k, x_{k+1} - x^* \rangle + \sum_{i=1}^n \frac{L_i}{2n} \|x_{k+1} - x_{s_{i,k}}\|^2. \quad (59)$$

By adding inequalities (58) and (59), rearranging the terms, and recalling that $P(x) = F(x) + R(x)$, we have

$$\begin{aligned} P(x_{k+1}) - P^* + \frac{1}{2\gamma} \|x_{k+1} - x^*\|^2 &\leq \frac{1}{2\gamma} \|x_k - x^*\|^2 + \underbrace{\sum_{i=1}^n \frac{L_i}{2n} \|x_{k+1} - x_{s_{i,k}}\|^2}_{\mathcal{H}} \\ &\quad - \frac{1}{2\gamma} \|x_{k+1} - x_k\|^2. \end{aligned} \quad (60)$$

Next, we find an upper bound on the term \mathcal{H} . We expand \mathcal{H} as follows:

$$\mathcal{H} = \sum_{i=1}^n \frac{L_i}{2n} \left\| \sum_{\ell=s_{i,k}}^k x_{\ell+1} - x_{\ell} \right\|^2 = \sum_{i=1}^n \frac{L_i(k-s_{i,k}+1)^2}{2n} \left\| \sum_{\ell=s_{i,k}}^k \frac{x_{\ell+1} - x_{\ell}}{k-s_{i,k}+1} \right\|^2.$$

The squared Euclidean norm ($\|\cdot\|^2$) is convex. Thus,

$$\mathcal{H} \leq \sum_{i=1}^n \frac{L_i(k-s_{i,k}+1)}{2n} \sum_{\ell=s_{i,k}}^k \|x_{\ell+1} - x_{\ell}\|^2.$$

Since $(k-\tau)_+ \leq s_{i,k} \leq k$ for all $i \in [n]$ and $k \in \mathbb{N}_0$, we have

$$\begin{aligned} \mathcal{H} &\leq \sum_{i=1}^n \frac{L_i(\tau+1)}{2n} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_{\ell}\|^2 \\ &\stackrel{(21)}{=} \frac{L(\tau+1)}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_{\ell}\|^2. \end{aligned} \quad (61)$$

Substituting (61) into the bound (60) concludes the proof. \blacksquare

Lemma 30 *Suppose Assumptions 1 and 2 hold. Let $\{\alpha_k\}$ be a sequence of non-negative numbers. Then, for every $k \in \mathbb{N}_0$, the sequence $\{x_k\}$ generated by Algorithm 2 satisfies*

$$\begin{aligned} \alpha_k(P(x_{k+1}) - P^*) &\leq \alpha_k(P(x_k) - P^*) + \frac{L\alpha_k}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_{\ell}\|^2 \\ &\quad - \alpha_k \left(\frac{1}{\gamma} - \frac{L\tau}{2} \right) \|x_{k+1} - x_k\|^2. \end{aligned} \quad (62)$$

Proof. According to Assumption 2, each component function f_i , $i \in [n]$, is L_i -smooth. It follows from the second inequality of (3.5) in Lemma 3.4 in Vanli et al. (2018) that

$$\begin{aligned} P(x_{k+1}) - P^* &\leq P(x_k) - P^* + \frac{L}{2} \sum_{\ell=(k-\tau)_+}^{k-1} \|x_{\ell+1} - x_{\ell}\|^2 - \left(\frac{1}{\gamma} - \frac{L(\tau+1)}{2} \right) \|x_{k+1} - x_k\|^2 \\ &= P(x_k) - P^* + \frac{L}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_{\ell}\|^2 - \left(\frac{1}{\gamma} - \frac{L\tau}{2} \right) \|x_{k+1} - x_k\|^2. \end{aligned}$$

Multiplying both sides of the above inequality by the non-negative number α_k proves the lemma. \blacksquare

C.1 Proof of Lemma 6

Since $\alpha_k = k + \alpha_0$ with $\alpha_0 \geq 0$, the sequence $\{\alpha_k\}$ is non-negative and satisfies $\alpha_{k+1} = \alpha_k + 1$ for every $k \in \mathbb{N}_0$. Adding inequalities (55) and (62), and rearranging the terms, we have

$$\begin{aligned} \alpha_{k+1}(P(x_{k+1}) - P^*) + \frac{1}{2\gamma} \|x_{k+1} - x^*\|^2 &\leq \alpha_k(P(x_k) - P^*) + \frac{1}{2\gamma} \|x_k - x^*\|^2 \\ &\quad + \frac{L(\alpha_k + \tau + 1)}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_\ell\|^2 \\ &\quad - \frac{1}{2} \left(\frac{2\alpha_k + 1}{\gamma} - \alpha_k L\tau \right) \|x_{k+1} - x_k\|^2. \end{aligned} \quad (63)$$

Multiplying both sides by 2γ and then letting $V_k = 2\gamma\alpha_k(P(x_k) - P^*) + \|x_k - x^*\|^2$ and $W_k = \|x_{k+1} - x_k\|^2$, we can rewrite (63) as

$$V_{k+1} \leq V_k + \gamma L(\alpha_k + \tau + 1) \sum_{\ell=(k-\tau)_+}^k W_\ell - (2\alpha_k + 1 - \gamma\alpha_k L\tau)W_k, \quad k \in \mathbb{N}_0.$$

The proof is complete.

C.2 Proof of Theorem 7

According to Lemma 6, the iterates generated by Algorithm 2 satisfy

$$V_{k+1} \leq V_k + p_k \sum_{\ell=(k-\tau)_+}^k W_\ell - r_k W_k, \quad k \in \mathbb{N}_0,$$

where $p_k = \gamma L(\alpha_k + \tau + 1)$, $r_k = 2\alpha_k + 1 - \gamma L\tau\alpha_k$, and $\alpha_k = k + \alpha_0$. To apply Lemma 5, we need to enforce that the convergence condition

$$\sum_{\ell=0}^{\tau} \gamma L(\alpha_{\ell+k} + \tau + 1) \leq 2\alpha_k + 1 - \gamma L\tau\alpha_k$$

is satisfied for every $k \in \mathbb{N}_0$. This inequality is equivalent to

$$\gamma L \leq \frac{2\alpha_k + 1}{\tau\alpha_k + \sum_{\ell=0}^{\tau} (\alpha_{\ell+k} + \tau + 1)}. \quad (64)$$

We will prove that if $\alpha_0 = \tau$ and $\gamma L(2\tau + 1) \leq 1$, then (64) holds for all $k \in \mathbb{N}_0$. Replacing $\alpha_k = k + \tau$ in (64), we have

$$\begin{aligned} \frac{2\alpha_k + 1}{\tau\alpha_k + \sum_{\ell=0}^{\tau} (\alpha_{\ell+k} + \tau + 1)} &= \frac{2k + 2\tau + 1}{(2\tau + 1)k + \frac{7}{2}\tau^2 + \frac{7}{2}\tau + 1} \\ &= \frac{1}{2\tau + 1} + \frac{(2\tau + 1)k + \frac{1}{2}\tau(\tau + 1)}{(2\tau + 1) \left((2\tau + 1)k + \frac{7}{2}\tau^2 + \frac{7}{2}\tau + 1 \right)}. \end{aligned}$$

The second term on the right-hand side is non-negative for any $\tau \in \mathbb{N}_0$ and $k \in \mathbb{N}_0$. Thus,

$$\frac{1}{2\tau + 1} \leq \frac{2\alpha_k + 1}{\tau\alpha_k + \sum_{\ell=0}^{\tau} (\alpha_{\ell+k} + \tau + 1)}, \quad k \in \mathbb{N}_0.$$

This shows that if

$$\gamma L \leq \frac{1}{2\tau + 1},$$

then (64) holds for $k \in \mathbb{N}_0$, and hence the convergence condition (19) in Lemma 5 is satisfied. It follows from part 1 of Lemma 5 that $V_k \leq V_0$ for $k \in \mathbb{N}$. Recalling that $V_k = 2\gamma\alpha_k(P(x_k) - P^*) + \|x_k - x^*\|^2$, we obtain

$$2\gamma(k + \tau)(P(x_k) - P^*) + \|x_k - x^*\|^2 \leq 2\gamma\tau(P(x_0) - P^*) + \|x_0 - x^*\|^2, \quad k \in \mathbb{N}.$$

By dropping the second term on the left-hand side and dividing both sides by $2\gamma(k + \tau)$, we finish the proof.

C.3 Proof of Lemma 9

According to Lemma 29 in Appendix C, the relation (55) holds for every $x^* \in \mathcal{X}^*$. Taking $x^* = \Pi_{\mathcal{X}^*}(x_k)$ in (55), we get

$$\begin{aligned} P(x_{k+1}) - P^* + \frac{1}{2\gamma}\|x_{k+1} - \Pi_{\mathcal{X}^*}(x_k)\|^2 &\leq \frac{1}{2\gamma}\|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2 \\ &+ \frac{L(\tau + 1)}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_\ell\|^2 - \frac{1}{2\gamma}\|x_{k+1} - x_k\|^2. \end{aligned} \quad (65)$$

By the projection property, we have $\|x_{k+1} - \Pi_{\mathcal{X}^*}(x_{k+1})\| \leq \|x_{k+1} - \Pi_{\mathcal{X}^*}(x_k)\|$. Combining this inequality with (65) yields

$$\begin{aligned} P(x_{k+1}) - P^* + \frac{1}{2\gamma}\|x_{k+1} - \Pi_{\mathcal{X}^*}(x_{k+1})\|^2 &\leq \frac{1}{2\gamma}\|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2 \\ &+ \frac{L(\tau + 1)}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_\ell\|^2 - \frac{1}{2\gamma}\|x_{k+1} - x_k\|^2. \end{aligned} \quad (66)$$

Adding inequalities (62) and (66), and setting $\alpha_k = \alpha$ for some $\alpha > 0$, we have

$$\begin{aligned} (1 + \alpha)(P(x_{k+1}) - P^*) + \frac{1}{2\gamma}\|x_{k+1} - \Pi_{\mathcal{X}^*}(x_{k+1})\|^2 &\leq \alpha(P(x_k) - P^*) + \frac{1}{2\gamma}\|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2 \\ &+ \frac{L(\alpha + \tau + 1)}{2} \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_\ell\|^2 \\ &- \frac{1}{2} \left(\frac{2\alpha + 1}{\gamma} - \alpha L\tau \right) \|x_{k+1} - x_k\|^2. \end{aligned} \quad (67)$$

Let $\theta = \frac{Q}{Q+1}$. Note that $\theta \in (0, 1)$. It follows from Assumption 3 that

$$\begin{aligned} (1 + \alpha)(P(x_{k+1}) - P^*) &= \theta(P(x_{k+1}) - P^*) + (1 - \theta + \alpha)(P(x_{k+1}) - P^*) \\ &\geq \frac{\mu\theta}{2}\|x_{k+1} - \Pi_{\mathcal{X}^*}(x_{k+1})\|^2 + (1 - \theta + \alpha)(P(x_{k+1}) - P^*). \end{aligned}$$

Combining the above inequality with (67) and then multiplying both sides by 2γ , we obtain

$$\begin{aligned} 2\gamma(1 - \theta + \alpha)(P(x_{k+1}) - P^*) + (1 + \gamma\mu\theta)\|x_{k+1} - \Pi_{\mathcal{X}^*}(x_{k+1})\|^2 \\ \leq 2\alpha\gamma(P(x_k) - P^*) + \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2 \\ + \gamma L(\alpha + \tau + 1) \sum_{\ell=(k-\tau)_+}^k \|x_{\ell+1} - x_\ell\|^2 \\ - (2\alpha + 1 - \gamma\alpha L\tau) \|x_{k+1} - x_k\|^2. \end{aligned} \quad (68)$$

By letting $\alpha = \frac{1}{\gamma L}$, $V_k = \frac{2}{L}(P(x_k) - P^*) + \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2$ and $W_k = \|x_{k+1} - x_k\|^2$, the inequality (68) can be rewritten as

$$(1 + \gamma\mu\theta)V_{k+1} \leq V_k + (1 + \gamma L(\tau + 1)) \sum_{\ell=(k-\tau)_+}^k W_\ell - \left(\frac{2}{\gamma L} + 1 - \tau\right) W_k.$$

Dividing both sides by $1 + \gamma\mu\theta$ completes the proof.

C.4 Proof of Theorem 10

According to Lemma 9, the iterates generated by Algorithm 2 satisfy

$$V_{k+1} \leq qV_k + p \sum_{\ell=(k-\tau)_+}^k W_\ell - rW_k$$

for every $k \in \mathbb{N}_0$, where

$$q = \frac{1}{1 + \gamma\mu\theta}, \quad p = \frac{1 + \gamma L(\tau + 1)}{1 + \gamma\mu\theta} \quad \text{and} \quad r = \frac{\frac{2}{\gamma L} + 1 - \tau}{1 + \gamma\mu\theta}.$$

To apply Lemma 5, we need to ensure that

$$2\tau + 1 \leq \min \left\{ \frac{1}{\gamma\mu\theta} + 1, \frac{\frac{2}{\gamma L} + 1 - \tau}{1 + \gamma L(\tau + 1)} \right\}.$$

This convergence condition is equivalent to

$$\begin{cases} 2\gamma\mu\theta\tau & \leq 1, \\ \gamma^2 L^2 (2\tau + 1)(\tau + 1) + 3\gamma L\tau & \leq 2. \end{cases} \quad (69)$$

Define $h = \gamma L(2\tau + 1)$. The inequalities (69) can be rewritten in terms of h , Q (recall that $\theta = \frac{Q}{Q+1}$), and τ as

$$\begin{cases} 2h\tau & \leq (Q + 1)(2\tau + 1), \\ h^2(\tau + 1) + 3h\tau & \leq 2(2\tau + 1). \end{cases} \quad (70)$$

For any fixed τ , the left-hand side of inequalities (70) is non-decreasing in $h \geq 0$ and smaller than the right-hand side for $h = 1$. Thus, the inequalities (70) hold for any $h \in [0, 1]$. This shows that if γ is set to

$$\gamma = \frac{h}{L(2\tau + 1)}, \quad h \in (0, 1],$$

then the convergence condition (69) is satisfied. Therefore, by part 2 of Lemma 5, $V_k \leq q^k V_0$ for $k \in \mathbb{N}$. Since $V_k = \frac{2}{L}(P(x_k) - P^*) + \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2$, it follows that

$$\begin{aligned} & \frac{2}{L}(P(x_k) - P^*) + \|x_k - \Pi_{\mathcal{X}^*}(x_k)\|^2 \\ & \leq \left(\frac{1}{1 + \gamma\mu\theta} \right)^k \left(\frac{2}{L}(P(x_0) - P^*) + \|x_0 - \Pi_{\mathcal{X}^*}(x_0)\|^2 \right) \\ & = \left(1 - \frac{1}{1 + (Q + 1)(2\tau + 1)/h} \right)^k \left(\frac{2}{L}(P(x_0) - P^*) + \|x_0 - \Pi_{\mathcal{X}^*}(x_0)\|^2 \right). \end{aligned}$$

Appendix D. Proofs for Subsection 4.2

In this section, we provide the proofs for the results presented in Subsection 4.2. We denote by \mathbb{E} the expectation with respect to the underlying probability space and by \mathbb{E}_k the conditional expectation with respect to the past, that is, with respect to $\{\xi_0, \dots, \xi_{k-1}\}$.

We first state a useful lemma which helps us to determine admissible step-sizes for the asynchronous SGD method.

Lemma 31 *Assume that α and β are non-negative real numbers. If $\alpha(\beta + 1) \leq 1$, then*

$$\alpha^2\beta^2 + \alpha \leq 1.$$

Proof. If $\alpha(\beta + 1) \leq 1$ holds, then $\alpha \leq 1/(\beta + 1)$ and $\alpha^2 \leq 1/(\beta + 1)^2$. Thus,

$$\begin{aligned} \alpha^2\beta^2 + \alpha & \leq \frac{\beta^2}{(\beta + 1)^2} + \frac{1}{\beta + 1} \\ & = \frac{\beta^2 + \beta + 1}{(\beta + 1)^2}. \end{aligned}$$

Since $\beta \geq 0$, one can verify that $\beta^2 + \beta + 1 \leq (\beta + 1)^2$, implying that the right-hand side of the above inequality is less than or equal to 1. Therefore, $\alpha^2\beta^2 + \alpha \leq 1$. \blacksquare

D.1 Proof of Lemma 13

When $\gamma_k = 0$, $x_{k+1} = x_k$. This implies that $\|x_{k+1} - x^*\|^2 = \|x_k - x^*\|^2$, and hence $V_{k+1} = V_k$. Below, we assume $\gamma_k > 0$.

By subtracting x^* from both sides of (29) and then taking norm squares, we have

$$\|x_{k+1} - x^*\|^2 = \|x_k - x^*\|^2 - 2\gamma_k \langle x_k - x^*, \nabla f(x_{k-\tau_k}, \xi_k) \rangle + \gamma_k^2 \|\nabla f(x_{k-\tau_k}, \xi_k)\|^2. \quad (71)$$

Taking conditional expectation on both sides of (71) with respect to $\{\xi_0, \dots, \xi_{k-1}\}$ leads to

$$\begin{aligned} \mathbb{E}_k[\|x_{k+1} - x^*\|^2] &= \|x_k - x^*\|^2 - 2\gamma_k \langle x_k - x^*, \nabla F(x_{k-\tau_k}) \rangle + \gamma_k^2 \mathbb{E}_k[\|\nabla f(x_{k-\tau_k}, \xi_k)\|^2] \\ &= \|x_k - x^*\|^2 - 2\gamma_k \langle x_k - x^*, \nabla F(x_{k-\tau_k}) \rangle \\ &\quad + \gamma_k^2 \mathbb{E}_k[\|\nabla f(x_{k-\tau_k}, \xi_k) - \nabla F(x_{k-\tau_k})\|^2] + \gamma_k^2 \|\nabla F(x_{k-\tau_k})\|^2, \end{aligned}$$

where the second equality uses the fact that $\mathbb{E}[\|x\|^2] = \mathbb{E}[\|x - \mathbb{E}[x]\|^2] + \|\mathbb{E}[x]\|^2$ for any random vector $x \in \mathbb{R}^d$. According to Assumption 5, stochastic gradients have bounded variance. Thus,

$$\mathbb{E}_k[\|x_{k+1} - x^*\|^2] \leq \|x_k - x^*\|^2 - 2\gamma_k \langle x_k - x^*, \nabla F(x_{k-\tau_k}) \rangle + \gamma_k^2 \sigma^2 + \gamma_k^2 \|\nabla F(x_{k-\tau_k})\|^2. \quad (72)$$

1. Convex case

Since F is convex and L -smooth, it follows from Theorem 2.1.5 in Nesterov (2013) that

$$F(x_{k-\tau_k}) + \langle x^* - x_{k-\tau_k}, \nabla F(x_{k-\tau_k}) \rangle + \frac{1}{2L} \|\nabla F(x_{k-\tau_k})\|^2 \leq F^*. \quad (73)$$

Moreover, from Lemma 1.2.3 in Nesterov (2013), we have

$$F(x_k) \leq F(x_{k-\tau_k}) + \langle x_k - x_{k-\tau_k}, \nabla F(x_{k-\tau_k}) \rangle + \frac{L}{2} \|x_k - x_{k-\tau_k}\|^2. \quad (74)$$

By adding inequalities (73) and (74) and then rearranging the terms, we get

$$F(x_k) - F^* \leq \langle x_k - x^*, \nabla F(x_{k-\tau_k}) \rangle + \frac{L}{2} \|x_k - x_{k-\tau_k}\|^2 - \frac{1}{2L} \|\nabla F(x_{k-\tau_k})\|^2.$$

Combining this inequality with (72) and taking the full expectation on both sides, we obtain

$$\begin{aligned} 2\gamma_k \mathbb{E}[F(x_k) - F^*] + \mathbb{E}[\|x_{k+1} - x^*\|^2] &\leq \mathbb{E}[\|x_k - x^*\|^2] + \gamma_k L \underbrace{\mathbb{E}[\|x_k - x_{k-\tau_k}\|^2]}_{\mathcal{H}} \\ &\quad - \left(\frac{\gamma_k}{L} - \gamma_k^2\right) \mathbb{E}[\|\nabla F(x_{k-\tau_k})\|^2] + \gamma_k^2 \sigma^2. \end{aligned} \quad (75)$$

Next, we find an upper bound on the term \mathcal{H} . From the first inequality in the proof of Lemma 15 in Koloskova et al. (2022), we have

$$\mathbb{E}[\|x_k - x_{k-\tau_k}\|^2] \leq 2\tau_k \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2 \mathbb{E}[\|\nabla F(x_{\ell-\tau_\ell})\|^2] + 2\sigma^2 \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2.$$

Substituting the above inequality into (75) yields

$$\begin{aligned}
 2\gamma_k \mathbb{E}[F(x_k) - F^*] + \mathbb{E}[\|x_{k+1} - x^*\|^2] &\leq \mathbb{E}[\|x_k - x^*\|^2] \\
 &+ 2\gamma_k \tau_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2 \mathbb{E}[\|\nabla F(x_{\ell-\tau_\ell})\|^2] \\
 &- \left(\frac{\gamma_k}{L} - \gamma_k^2\right) \mathbb{E}[\|\nabla F(x_{k-\tau_k})\|^2] \\
 &+ \left(\gamma_k^2 + 2\gamma_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2\right) \sigma^2. \tag{76}
 \end{aligned}$$

By letting $X_k = 2\gamma_k \mathbb{E}[F(x_k) - F^*]$, $V_k = \mathbb{E}[\|x_k - x^*\|^2]$, and $W_k = \gamma_k^2 \mathbb{E}[\|\nabla F(x_{k-\tau_k})\|^2]$, we can rewrite (76) as

$$X_k + V_{k+1} \leq V_k + 2\gamma_k \tau_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} W_\ell - \left(\frac{1}{\gamma_k L} - 1\right) W_k + \left(\gamma_k^2 + 2\gamma_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2\right) \sigma^2.$$

2. Strongly convex case

If F is μ -strongly convex, then

$$\frac{\mu}{2} \|x_k - x^*\|^2 \leq F(x_k) - F^*,$$

which implies that $\gamma_k \mu V_k \leq X_k$. Therefore,

$$V_{k+1} \leq (1 - \gamma_k \mu) V_k + 2\gamma_k \tau_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} W_\ell - \left(\frac{1}{\gamma_k L} - 1\right) W_k + \left(\gamma_k^2 + 2\gamma_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} \gamma_\ell^2\right) \sigma^2.$$

This completes the proof.

D.2 Proof of Theorem 14

For any $K \in \mathbb{N}_0$, define the sets $\mathcal{T}_K = \{k \leq K \mid \tau_k \leq \tau_{\text{th}}\}$ and $\bar{\mathcal{T}}_K = \{k \leq K \mid \tau_k > \tau_{\text{th}}\}$. Since $\gamma_k = \gamma$ for $\tau_k \leq \tau_{\text{th}}$ and $\gamma_k = 0$ for $\tau_k > \tau_{\text{th}}$, we have

$$\gamma_k = \begin{cases} \gamma, & k \in \mathcal{T}_K, \\ 0, & k \in \bar{\mathcal{T}}_K. \end{cases}$$

1. Convex case

For any $k \in \mathcal{T}_K$, according to Lemma 13, the iterates generated by Algorithm 3 satisfy

$$\begin{aligned}
 X_k + V_{k+1} &\leq V_k + 2\gamma\tau_k L \sum_{\ell=(k-\tau_k)_+}^{k-1} W_\ell - \left(\frac{1}{\gamma L} - 1\right) W_k + (1 + 2\gamma\tau_k L) \gamma^2 \sigma^2 \\
 &\leq V_k + 2\gamma\tau_{\text{th}} L \sum_{\ell=(k-\tau_{\text{th}})_+}^{k-1} W_\ell - \left(\frac{1}{\gamma L} - 1\right) W_k + (1 + 2\gamma\tau_{\text{th}} L) \gamma^2 \sigma^2 \\
 &= V_k + 2\gamma\tau_{\text{th}} L \sum_{\ell=(k-\tau_{\text{th}})_+}^k W_\ell - \left(\frac{1}{\gamma L} + 2\gamma\tau_{\text{th}} L - 1\right) W_k + (1 + 2\gamma\tau_{\text{th}} L) \gamma^2 \sigma^2,
 \end{aligned} \tag{77}$$

where the second inequality uses $\tau_k \leq \tau_{\text{th}}$ for $k \in \mathcal{T}_K$. For $k \in \overline{\mathcal{T}}_K$, $\gamma_k = 0$ and, hence, $X_k = 0$. Thus, by Lemma 13, we have

$$X_k + V_{k+1} = V_k, \quad k \in \overline{\mathcal{T}}_K.$$

Next, note that this implies that (77) also holds for $k \in \overline{\mathcal{T}}_K$ since W_ℓ is non-negative for all $\ell \in \mathbb{N}_0$, $W_k = 0$ for $k \in \overline{\mathcal{T}}_K$, and $(1 + 2\gamma\tau_{\text{th}} L) \gamma^2 \sigma^2 \geq 0$. Therefore,

$$X_k + V_{k+1} \leq V_k + p \sum_{\ell=(k-\tau_{\text{th}})_+}^k W_\ell - r W_k + e$$

with $p = 2\gamma\tau_{\text{th}} L$, $r = \frac{1}{\gamma L} + 2\gamma\tau_{\text{th}} L - 1$, and $e = (1 + 2\gamma\tau_{\text{th}} L) \gamma^2 \sigma^2$ is satisfied by the iterates for every k . To apply Lemma 5, we need to ensure that

$$2\gamma\tau_{\text{th}} L(\tau_{\text{th}} + 1) \leq \left(\frac{1}{\gamma L} + 2\gamma\tau_{\text{th}} L - 1\right).$$

This convergence condition is equivalent to

$$2\gamma^2\tau_{\text{th}}^2 L^2 + \gamma L \leq 1. \tag{78}$$

Using Lemma 31 in Appendix D with $\alpha = \gamma L$ and $\beta = \tau_{\text{th}}\sqrt{2}$, we choose the step-size γ as

$$\gamma \in \left(0, \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)}\right]$$

to guarantee the convergence condition (78). It follows from part 1 of Lemma 5 that

$$\sum_{k=0}^K X_k \leq V_0 + (K + 1)e, \quad K \in \mathbb{N}_0.$$

Since $V_0 = \|x_0 - x^*\|^2$, $X_k = 2\gamma_k \mathbb{E}[F(x_k) - F^*]$ and $e = (1 + 2\gamma\tau_{\text{th}} L) \gamma^2 \sigma^2$, we obtain

$$\begin{aligned}
 2 \sum_{k=0}^K \gamma_k \mathbb{E}[F(x_k) - F^*] &\leq \|x_0 - x^*\|^2 + (K + 1)(1 + 2\gamma\tau_{\text{th}} L) \gamma^2 \sigma^2 \\
 &\leq \|x_0 - x^*\|^2 + (K + 1)(1 + \sqrt{2})\gamma^2 \sigma^2,
 \end{aligned}$$

where the second inequality follows since γ is selected to make $\gamma L(\tau_{\text{th}}\sqrt{2} + 1) \leq 1$. By the convexity of F and definition of \bar{x}_K , we have

$$F(\bar{x}_K) - F^* = F\left(\frac{1}{\sum_{k=0}^K \gamma_k} \sum_{k=0}^K \gamma_k x_k\right) - F^* \leq \frac{1}{\sum_{k=0}^K \gamma_k} \sum_{k=0}^K \gamma_k (F(x_k) - F^*),$$

which implies that

$$\mathbb{E}[F(\bar{x}_K) - F^*] \leq \frac{\|x_0 - x^*\|^2 + (K+1)(1+\sqrt{2})\gamma^2\sigma^2}{2\sum_{k=0}^K \gamma_k}, \quad K \in \mathbb{N}_0. \quad (79)$$

Next, we will find a lower bound on $\sum_{k=0}^K \gamma_k$. As $\gamma_k = \gamma$ for $k \in \mathcal{T}_K$ and $\gamma_k = 0$ for $k \in \bar{\mathcal{T}}_K$, we have

$$\sum_{k=0}^K \gamma_k = \gamma |\mathcal{T}_K|. \quad (80)$$

We consider two cases:

1. If $2\tau_{\text{ave}} \leq \tau_{\text{max}}$, then $2\tau_{\text{ave}} \leq \tau_{\text{th}}$. As observed by Koloskova et al. (2022), since

$$\sum_{k=0}^K \tau_k = (K+1)\tau_{\text{ave}},$$

at most $\lfloor (K+1)/2 \rfloor$ of the terms on the left-hand side of the above equality can be larger than $2\tau_{\text{ave}}$. Thus, there are at least half of the iterations with the delay smaller than $2\tau_{\text{ave}}$ and, hence, τ_{th} . This implies that $(K+1)/2 \leq |\mathcal{T}_K|$.

2. If $\tau_{\text{max}} \leq 2\tau_{\text{ave}}$, then $\tau_{\text{max}} \leq \tau_{\text{th}}$. In this case, we have $\tau_k \leq \tau_{\text{th}}$ for all $k = 0, \dots, K$ since $\tau_k \leq \tau_{\text{max}}$ for $k \in \mathbb{N}_0$. This shows that $|\mathcal{T}_K| = K+1$.

Therefore, $(K+1)/2 \leq |\mathcal{T}_K|$. From (80), we then have

$$\sum_{k=0}^K \gamma_k \geq \frac{\gamma(K+1)}{2}.$$

Combining this inequality with (79) leads to

$$\mathbb{E}[F(\bar{x}_K) - F^*] \leq \frac{\|x_0 - x^*\|^2}{\gamma(K+1)} + (1+\sqrt{2})\gamma\sigma^2, \quad K \in \mathbb{N}_0.$$

2. Strongly convex case

Similar to the proof for the convex case,

$$V_{k+1} \leq (1 - \gamma\mu)V_k + 2\gamma\tau_{\text{th}}L \sum_{\ell=(k-\tau_{\text{th}})_+}^k W_\ell - \left(\frac{1}{\gamma L} + 2\gamma\tau_{\text{th}}L - 1\right)W_k + (1 + 2\gamma\tau_{\text{th}}L)\gamma^2\sigma^2 \quad (81)$$

is satisfied for $k \in \mathcal{T}_K$. For $k \in \overline{\mathcal{T}}_K$, $\gamma_k = 0$. Thus, according to Lemma 13, we have $V_{k+1} = V_k$ for $k \in \overline{\mathcal{T}}_K$. This implies that

$$V_{k+1} \leq V_k + 2\gamma\tau_{\text{th}}L \sum_{\ell=(k-\tau_{\text{th}})_+}^k W_\ell - \left(\frac{1}{\gamma L} + 2\gamma\tau_{\text{th}}L - 1 \right) W_k \quad (82)$$

holds for $k \in \overline{\mathcal{T}}_K$ since W_ℓ is non-negative for $\ell \in \mathbb{N}_0$ and $W_k = 0$ for $k \in \overline{\mathcal{T}}_K$. It follows from (81) and (82) that the iterates generated by Algorithm 3 satisfy

$$V_{k+1} \leq q_k V_k + 2\gamma\tau_{\text{th}}L \sum_{\ell=(k-\tau_{\text{th}})_+}^k W_\ell - \left(\frac{1}{\gamma L} + 2\gamma\tau_{\text{th}}L - 1 \right) W_k + e_k,$$

where $q_k = 1$ and $e_k = 0$ for $k \in \overline{\mathcal{T}}_K$, and $q_k = 1 - \gamma\mu$ and $e_k = (1 + 2\gamma\tau_{\text{th}}L)\gamma^2\sigma^2$ for $k \in \mathcal{T}_K$. Let $q = 1 - \gamma\mu$. It is clear that $q_k \geq q$ for every k . Thus, to apply Lemma 5, we need to enforce that

$$2\tau_{\text{th}} + 1 \leq \min \left\{ \frac{1}{\gamma\mu}, \frac{\frac{1}{\gamma L} + 2\gamma\tau_{\text{th}}L - 1}{2\gamma\tau_{\text{th}}L} \right\}.$$

This convergence condition is equivalent to

$$\begin{cases} \gamma\mu(2\tau_{\text{th}} + 1) & \leq 1, \\ 4\gamma^2\tau_{\text{th}}^2L^2 + \gamma L & \leq 1. \end{cases} \quad (83)$$

If $\gamma L(2\tau_{\text{th}} + 1) \leq 1$, the first inequality in (83) holds since $\mu \leq L$. Using Lemma 31 in Appendix D with $\alpha = \gamma L$ and $\beta = 2\tau_{\text{th}}$, we can see that if $\gamma L(2\tau_{\text{th}} + 1) \leq 1$, the second inequality in (83) also holds. Thus, the convergence condition (83) is satisfied for

$$\gamma \in \left(0, \frac{1}{L(2\tau_{\text{th}} + 1)} \right].$$

Therefore, by part 2 of Lemma 5, we have

$$V_{K+1} \leq Q_{K+1}V_0 + Q_{K+1} \sum_{k=0}^K \frac{e_k}{Q_{k+1}}, \quad K \in \mathbb{N}_0, \quad (84)$$

where $Q_k = \prod_{\ell=0}^{k-1} q_\ell$. Since $q_\ell = 1$ for $\ell \in \overline{\mathcal{T}}_K$ and $q_\ell = 1 - \gamma\mu$ for $\ell \in \mathcal{T}_K$, we have

$$Q_{K+1} = \prod_{\ell=0}^K q_\ell = \prod_{\ell \in \mathcal{T}_K} (1 - \gamma\mu) = (1 - \gamma\mu)^{|\mathcal{T}_K|}.$$

As discussed in the proof for the convex case, $(K + 1)/2 \leq |\mathcal{T}_K|$. Thus,

$$\begin{aligned} Q_{K+1} &\leq (1 - \gamma\mu)^{\frac{K+1}{2}} \\ &\leq \exp \left(-\frac{\gamma\mu(K+1)}{2} \right), \end{aligned} \quad (85)$$

where the second inequality follows from the fact that $1 - \alpha \leq \exp(-\alpha)$ for $\alpha \geq 0$. Next, we bound $Q_{K+1} \sum_{k=0}^K e_k / Q_{k+1}$. Since $e_k = 0$ for $k \in \overline{\mathcal{T}}_K$ and $e_k = (1 + 2\gamma\tau_{\text{th}}L)\gamma^2\sigma^2$ for $k \in \mathcal{T}_K$, we have

$$\begin{aligned} Q_{K+1} \sum_{k=0}^K \frac{e_k}{Q_{k+1}} &= (1 - \gamma\mu)^{|\mathcal{T}_K|} \sum_{k \in \mathcal{T}_K} \frac{(1 + 2\gamma\tau_{\text{th}}L)\gamma^2\sigma^2}{\prod_{\ell=0}^k q_\ell} \\ &= (1 + 2\gamma\tau_{\text{th}}L)\gamma^2\sigma^2 \sum_{k=0}^{|\mathcal{T}_K|-1} (1 - \gamma\mu)^k \\ &\leq (1 + 2\gamma\tau_{\text{th}}L)\gamma^2\sigma^2 \sum_{k=0}^{\infty} (1 - \gamma\mu)^k \\ &= \frac{(1 + 2\gamma\tau_{\text{th}}L)\gamma\sigma^2}{\mu}. \end{aligned} \tag{86}$$

Substituting (85) and (86) into (84) yields

$$\begin{aligned} V_{K+1} &\leq \exp\left(-\frac{\gamma\mu(K+1)}{2}\right) V_0 + \frac{(1 + 2\gamma\tau_{\text{th}}L)\gamma\sigma^2}{\mu}, \\ &\leq \exp\left(-\frac{\gamma\mu(K+1)}{2}\right) V_0 + \frac{2\gamma\sigma^2}{\mu}, \end{aligned}$$

where the second inequality is due to that $\gamma L(2\tau_{\text{th}} + 1) \leq 1$. Since $V_k = \mathbb{E}[\|x_k - x^*\|^2]$, it follows that

$$\mathbb{E}[\|x_K - x^*\|^2] \leq \exp\left(-\frac{\gamma\mu K}{2}\right) \|x_0 - x^*\|^2 + \frac{2\gamma\sigma^2}{\mu} \quad K \in \mathbb{N}.$$

This completes the proof.

D.3 Proof of Theorem 15

1. Convex case

According to Theorem 14, the iterates generated by the asynchronous SGD method with $\gamma \leq \frac{1}{L(\tau_{\text{th}}\sqrt{2}+1)}$ satisfy

$$\mathbb{E}[F(\bar{x}_K) - F^*] \leq \frac{\|x_0 - x^*\|^2}{\gamma(K+1)} + (1 + \sqrt{2})\gamma\sigma^2, \quad K \in \mathbb{N}_0. \tag{87}$$

We will find the total number of iterations K_ϵ necessary to achieve ϵ -optimal solution, i.e.,

$$\mathbb{E}[F(\bar{x}_K) - F^*] \leq \epsilon, \quad \text{for } K \geq K_\epsilon.$$

By selecting γ so that $(\sqrt{2} + 1)\gamma\sigma^2 \leq \epsilon/2$, the second term on the right-hand side of (87) is less than $\epsilon/2$. Set γ to

$$\gamma = \min \left\{ \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)}, \frac{\epsilon}{2(\sqrt{2} + 1)\sigma^2} \right\}. \tag{88}$$

We now choose K so that the first term on the right-hand side of (87) is less than $\epsilon/2$, i.e.,

$$\frac{\|x_0 - x^*\|^2}{\gamma(K+1)} \leq \frac{\epsilon}{2}.$$

This implies that

$$\begin{aligned} K \geq K_\epsilon &:= \frac{2\|x_0 - x^*\|^2}{\gamma\epsilon} - 1 \\ &\stackrel{(88)}{=} 2 \max \left\{ \frac{L(\tau_{\text{th}}\sqrt{2} + 1)}{\epsilon}, \frac{2(\sqrt{2} + 1)\sigma^2}{\epsilon^2} \right\} \|x_0 - x^*\|^2 - 1. \end{aligned}$$

2. Strongly convex case

It follows from Theorem 14 that the iterates generated by the asynchronous SGD method with $\gamma \leq \frac{1}{L(2\tau_{\text{th}}+1)}$ satisfy

$$\mathbb{E}[\|x_K - x^*\|^2] \leq \exp\left(-\frac{\gamma\mu K}{2}\right) \|x_0 - x^*\|^2 + \frac{2\gamma\sigma^2}{\mu} \quad K \in \mathbb{N}. \quad (89)$$

With the choice of

$$\gamma = \min \left\{ \frac{1}{L(2\tau_{\text{th}} + 1)}, \frac{\epsilon\mu}{4\sigma^2} \right\}, \quad (90)$$

the second term on the right-hand side of (89) is less than $\epsilon/2$. Next we choose K so that

$$\exp\left(-\frac{\gamma\mu K}{2}\right) \|x_0 - x^*\|^2 \leq \frac{\epsilon}{2}.$$

Taking logarithm of both sides and rearranging the terms gives

$$\begin{aligned} K \geq K_\epsilon &:= \frac{2}{\gamma\mu} \log \left(\frac{2\|x_0 - x^*\|^2}{\epsilon} \right) \\ &\stackrel{(90)}{=} 2 \max \left\{ \frac{L(2\tau_{\text{th}} + 1)}{\mu}, \frac{4\sigma^2}{\mu^2\epsilon} \right\} \log \left(\frac{2\|x_0 - x^*\|^2}{\epsilon} \right). \end{aligned}$$

D.4 Proof of Theorem 16

1. Convex case

According to Theorem 14, we have

$$\mathbb{E}[F(\bar{x}_{\mathcal{K}}) - F^*] \leq \frac{\|x_0 - x^*\|^2}{\gamma(\mathcal{K} + 1)} + (1 + \sqrt{2})\gamma\sigma^2.$$

By minimizing the right-hand side of the above inequality with respect to γ over the interval

$$\left(0, \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)} \right],$$

we obtain

$$\gamma^* = \min \left\{ \frac{1}{L(\tau_{\text{th}}\sqrt{2} + 1)}, \frac{\|x_0 - x^*\|}{\sigma\sqrt{\sqrt{2} + 1}\sqrt{\mathcal{K} + 1}} \right\}.$$

The rest of the proof is similar to the one for the serial SGD method derived in Lan (2012) and thus omitted.

2. Strongly convex case

By Theorem 14, we have

$$\mathbb{E}[\|x_{\mathcal{K}} - x^*\|^2] \leq \exp\left(-\frac{\gamma\mu\mathcal{K}}{2}\right) \|x_0 - x^*\|^2 + \frac{2\gamma\sigma^2}{\mu}.$$

The result follows from Lemma 1 in Karimireddy et al. (2020).

Appendix E. Proofs for Subsection 4.3

In this section, we provide the proofs for the results presented in Subsection 4.3.

E.1 Preliminaries

Let $I \in \mathbb{R}^{d \times d}$ be the identity matrix. We define the matrices $U_i \in \mathbb{R}^{d \times d_i}$, $i \in [m]$, for which $I = [U_1, \dots, U_m]$. Then, any vector $x = ([x]_1, \dots, [x]_m) \in \mathbb{R}^d$ can be represented as

$$x = \sum_{i=1}^m U_i [x]_i, \quad [x]_i \in \mathbb{R}^{d_i}, \quad i \in [m].$$

Since $[x_{k+1}]_j = [x_k]_j - \gamma S_j(\hat{x}_k)$ for $j = i_k$ and $[x_{k+1}]_j = [x_k]_j$ for $j \neq i_k$, the update formula of Algorithm 4 can be written as

$$x_{k+1} = x_k - \gamma U_{i_k} S_{i_k}(\hat{x}_k), \quad k \in \mathbb{N}_0. \quad (91)$$

Note that x_k depends on the random variable $\xi_{k-1} := \{i_0, \dots, i_{k-1}\}$ but not on i_j for any $j \geq k$. For convenience, we define $\xi_{-1} = \emptyset$. We use \mathbb{E} to denote the expectation over all random variables, and \mathbb{E}_k to denote the conditional expectation in terms of i_k given ξ_{k-1} .

For any random vector $x \in \mathbb{R}^d$, the variance can be decomposed as

$$\mathbb{E}[\|x - \mathbb{E}[x]\|^2] = \mathbb{E}[\|x\|^2] - \|\mathbb{E}[x]\|^2. \quad (92)$$

For any vectors $a, b \in \mathbb{R}^d$ and any constant $\eta > 0$, the inequalities

$$\langle a, b \rangle \leq \frac{\eta\|a\|^2}{2} + \frac{\|b\|^2}{2\eta}, \quad (93)$$

$$\|a + b\|^2 \leq (1 + \eta)\|a\|^2 + \left(1 + \frac{1}{\eta}\right)\|b\|^2, \quad (94)$$

$$-\|a\|^2 \leq -\frac{\|b\|^2}{1 + \eta} + \frac{\|a - b\|^2}{\eta}, \quad (95)$$

hold by the Cauchy-Schwarz inequality.

We need the following two lemmas in the convergence analysis of ARock. The first one provides an upper bound on the expectation of $\|x_k - \hat{x}_k\|^2$.

Lemma 32 *Let $\{x(k)\}$ be the sequence generated by Algorithm 4. Then, it holds that*

$$\mathbb{E}[\|x_k - \hat{x}_k\|^2] \leq \frac{\gamma^2 (\sqrt{m} + \sqrt{\tau})^2}{m^2} \sum_{\ell=(k-\tau)_+}^{k-1} \mathbb{E}[\|S(\hat{x}_\ell)\|^2], \quad k \in \mathbb{N}_0.$$

Proof. Let η be a positive constant. From (36), we have

$$\begin{aligned} \mathbb{E}[\|x_k - \hat{x}_k\|^2] &= \mathbb{E} \left[\left\| \sum_{j \in J_k} (x_{j+1} - x_j) \right\|^2 \right] \\ &\stackrel{(91)}{=} \gamma^2 \mathbb{E} \left[\left\| \sum_{j \in J_k} U_{i_j} S_{i_j}(\hat{x}_j) \right\|^2 \right] \\ &= \gamma^2 \mathbb{E} \left[\left\| \sum_{j \in J_k} \left(U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j) \right) + \frac{1}{m} \sum_{j \in J_k} S(\hat{x}_j) \right\|^2 \right] \\ &\stackrel{(94)}{\leq} \underbrace{\gamma^2 (1 + \eta) \mathbb{E} \left[\left\| \sum_{j \in J_k} \left(U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j) \right) \right\|^2 \right]}_{\mathcal{H}_1} \\ &\quad + \underbrace{\frac{\gamma^2}{m^2} \left(1 + \frac{1}{\eta} \right) \mathbb{E} \left[\left\| \sum_{j \in J_k} S(\hat{x}_j) \right\|^2 \right]}_{\mathcal{H}_2}. \end{aligned} \tag{96}$$

We will find upper bounds on the quantities \mathcal{H}_1 and \mathcal{H}_2 . We expand \mathcal{H}_1 as follows

$$\begin{aligned} \mathcal{H}_1 &= \sum_{j \in J_k} \mathbb{E} \left[\left\| U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j) \right\|^2 \right] \\ &\quad + \sum_{j, j' \in J_k, j > j'} 2 \mathbb{E} \left[\left\langle U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j), U_{i_{j'}} S_{i_{j'}}(\hat{x}_{j'}) - \frac{1}{m} S(\hat{x}_{j'}) \right\rangle \right] \\ &= \sum_{j \in J_k} \mathbb{E} \left[\left\| U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j) \right\|^2 \right], \end{aligned}$$

where the second equality uses the fact that

$$\begin{aligned} &\sum_{j, j' \in J_k, j > j'} \mathbb{E} \left[\left\langle U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j), U_{i_{j'}} S_{i_{j'}}(\hat{x}_{j'}) - \frac{1}{m} S(\hat{x}_{j'}) \right\rangle \right] \\ &= \sum_{j, j' \in J_k, j > j'} \mathbb{E} \left[\mathbb{E}_j \left[\left\langle U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j), U_{i_{j'}} S_{i_{j'}}(\hat{x}_{j'}) - \frac{1}{m} S(\hat{x}_{j'}) \right\rangle \right] \right] \\ &= \sum_{j, j' \in J_k, j > j'} \mathbb{E} \left[\left\langle \mathbb{E}_j \left[U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j) \right], U_{i_{j'}} S_{i_{j'}}(\hat{x}_{j'}) - \frac{1}{m} S(\hat{x}_{j'}) \right\rangle \right] \\ &= 0. \end{aligned}$$

We then bound \mathcal{H}_1 by

$$\begin{aligned}
 \mathcal{H}_1 &= \sum_{j \in J_k} \mathbb{E} \left[\mathbb{E}_j \left[\left\| U_{i_j} S_{i_j}(\hat{x}_j) - \frac{1}{m} S(\hat{x}_j) \right\|^2 \right] \right] \\
 &\stackrel{(92)}{=} \frac{1}{m} \sum_{j \in J_k} \mathbb{E} \left[\|S(\hat{x}_j)\|^2 \right] - \frac{1}{m^2} \sum_{j \in J_k} \mathbb{E} \left[\|S(\hat{x}_j)\|^2 \right] \\
 &\leq \frac{1}{m} \sum_{j \in J_k} \mathbb{E} \left[\|S(\hat{x}_j)\|^2 \right]. \tag{97}
 \end{aligned}$$

Next, we turn to \mathcal{H}_2 . We rewrite \mathcal{H}_2 as

$$\mathcal{H}_2 = |J_k|^2 \mathbb{E} \left[\left\| \sum_{j \in J_k} \frac{S(\hat{x}_j)}{|J_k|} \right\|^2 \right].$$

By convexity of the squared Euclidean norm $\|\cdot\|^2$, we get

$$\mathcal{H}_2 \leq |J_k| \sum_{j \in J_k} \mathbb{E} \left[\|S(\hat{x}_j)\|^2 \right]. \tag{98}$$

Substituting (97) and (98) into (96) yields

$$\mathbb{E}[\|x_k - \hat{x}_k\|^2] \leq \frac{\gamma^2}{m^2} \left((1 + \eta)m + \left(1 + \frac{1}{\eta}\right) |J_k| \right) \sum_{j \in J_k} \mathbb{E} \left[\|S(\hat{x}_j)\|^2 \right].$$

From Assumption 7.2, $(k - \tau)_+ \leq j \leq k - 1$ for any $j \in J_k$ and $|J_k| \leq \tau$ for all $k \in \mathbb{N}_0$. Thus,

$$\mathbb{E}[\|x_k - \hat{x}_k\|^2] \leq \frac{\gamma^2}{m^2} \left((1 + \eta)m + \left(1 + \frac{1}{\eta}\right) \tau \right) \sum_{j=(k-\tau)_+}^{k-1} \mathbb{E} \left[\|S(\hat{x}_j)\|^2 \right].$$

It is easy to verify that the optimal choice of η , which minimizes the right-hand-side of the above inequality, is $\eta = \sqrt{\tau/m}$. Using this choice of η and the change of variable $\ell = j$ completes the proof. \blacksquare

As a consequence of this lemma, we derive a bound on the expectation of $\langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle$.

Lemma 33 *The iterates $\{x(k)\}$ generated by Algorithm 4 satisfy*

$$\mathbb{E}[\langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle] \leq \frac{\gamma}{2} \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}} \right) \left(\frac{1}{\tau} \sum_{\ell=(k-\tau)_+}^{k-1} \mathbb{E}[\|S(\hat{x}_\ell)\|^2] + \mathbb{E}[\|S(\hat{x}_k)\|^2] \right).$$

Proof. Let η be a positive constant. From (93), we have

$$\mathbb{E}[\langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle] \leq \frac{\eta}{2} \mathbb{E}[\|x_k - \hat{x}_k\|^2] + \frac{1}{2\eta} \mathbb{E}[\|S(\hat{x}_k)\|^2].$$

It follows from Lemma 32 that

$$\mathbb{E}[\langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle] \leq \frac{\eta\gamma^2(\sqrt{m} + \sqrt{\tau})^2}{2m^2} \sum_{\ell=(k-\tau)_+}^{k-1} \mathbb{E}[\|S(\hat{x}_\ell)\|^2] + \frac{1}{2\eta} \mathbb{E}[\|S(\hat{x}_k)\|^2].$$

Substituting

$$\eta = \frac{m}{\gamma\sqrt{\tau}(\sqrt{m} + \sqrt{\tau})}$$

into the above inequality proves the statement of the lemma. \blacksquare

E.2 Proof of Lemma 17

We reuse the proof technique from Hannah and Yin (2017). Our point of departure with Hannah and Yin (2017) is to use Lemmas 32 and 33 to bound $\mathbb{E}_k[\|x_k - \hat{x}_k\|^2]$ and $\mathbb{E}_k[\langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle]$. Let x^* be a fixed point of the operator T . By subtracting x^* from both sides of (91) and then taking norm squares, we have

$$\|x_{k+1} - x^*\|^2 = \|x_k - x^*\|^2 - 2\gamma \langle x_k - x^*, U_{i_k} S_{i_k}(\hat{x}_k) \rangle + \gamma^2 \|U_{i_k} S_{i_k}(\hat{x}_k)\|^2. \quad (99)$$

By taking conditional expectation on both sides of (99) with respect to only the random variable i_k , we obtain

$$\begin{aligned} \mathbb{E}_k[\|x_{k+1} - x^*\|^2] &= \|x_k - x^*\|^2 - 2\frac{\gamma}{m} \langle x_k - x^*, S(\hat{x}_k) \rangle + \frac{\gamma^2}{m} \|S(\hat{x}_k)\|^2 \\ &= \|x_k - x^*\|^2 - 2\frac{\gamma}{m} \underbrace{\langle \hat{x}_k - x^*, S(\hat{x}_k) \rangle}_{\mathcal{H}} + 2\frac{\gamma}{m} \langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle + \frac{\gamma^2}{m} \|S(\hat{x}_k)\|^2. \end{aligned} \quad (100)$$

We will use \mathcal{H} to generate a $\|x_k - x^*\|^2$ term to help prove linear convergence. Since T is pseudo-contractive with contraction modulus c and $S(x^*) = 0$, it follows from Lemma 8 in Hannah and Yin (2017) that

$$-\mathcal{H} \leq -\frac{(1-c^2)}{2} \|\hat{x}_k - x^*\|^2 - \frac{1}{2} \|S(\hat{x}_k)\|^2. \quad (101)$$

We then use (95) to convert $\|\hat{x}_k - x^*\|^2$ to $\|x_k - x^*\|^2$ as follows

$$-\|\hat{x}_k - x^*\|^2 \leq -\left(\frac{1}{1+\eta}\right) \|x_k - x^*\|^2 + \frac{1}{\eta} \|x_k - \hat{x}_k\|^2,$$

where η is a positive constant. Combining this inequality and (101), we have

$$\begin{aligned} -\mathcal{H} &\leq -\frac{(1-c^2)}{2(1+\eta)} \|x_k - x^*\|^2 + \frac{(1-c^2)}{2\eta} \|x_k - \hat{x}_k\|^2 - \frac{1}{2} \|S(\hat{x}_k)\|^2 \\ &\leq -\frac{(1-c^2)}{2(1+\eta)} \|x_k - x^*\|^2 + \frac{1}{2\eta} \|x_k - \hat{x}_k\|^2 - \frac{1}{2} \|S(\hat{x}_k)\|^2. \end{aligned}$$

Substituting the above inequality into (100) yields

$$\begin{aligned} \mathbb{E}_k[\|x_{k+1} - x^*\|^2] &\leq \left(1 - \frac{\gamma(1 - c^2)}{m(1 + \eta)}\right) \|x_k - x^*\|^2 + 2\frac{\gamma}{m} \langle \hat{x}_k - x_k, S(\hat{x}_k) \rangle \\ &\quad + \frac{\gamma}{m\eta} \|x_k - \hat{x}_k\|^2 - \frac{\gamma}{m} (1 - \gamma) \|S(\hat{x}_k)\|^2. \end{aligned}$$

After taking the full expectation on both sides, it follows from Lemmas 32 and 33 that

$$\begin{aligned} \mathbb{E}[\|x_{k+1} - x^*\|^2] &\leq \left(1 - \frac{\gamma(1 - c^2)}{m(1 + \eta)}\right) \mathbb{E}[\|x_k - x^*\|^2] \\ &\quad + \frac{\gamma^2}{m\tau} \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right) \left(1 + \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right) \frac{\gamma}{\eta}\right) \sum_{\ell=(k-\tau)_+}^{k-1} \mathbb{E}[\|S(\hat{x}_\ell)\|^2] \\ &\quad - \frac{\gamma}{m} \left(1 - \gamma \left(1 + \frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right) \mathbb{E}[\|S(\hat{x}_k)\|^2]. \end{aligned} \quad (102)$$

Let $V_k = \mathbb{E}[\|x_k - x^*\|^2]$, $W_k = \mathbb{E}[\|S(\hat{x}_k)\|^2]$, and

$$\eta = \gamma \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right).$$

Then, the inequality (102) can be rewritten as

$$\begin{aligned} V_{k+1} &\leq \left(1 - \frac{\gamma(1 - c^2)}{m(1 + \gamma(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}))}\right) V_k + \frac{2\gamma^2}{m\tau} \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right) \sum_{\ell=(k-\tau)_+}^{k-1} W_\ell \\ &\quad - \frac{\gamma}{m} \left(1 - \gamma \left(1 + \frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right) W_k. \end{aligned}$$

This completes the proof.

E.3 Proof of Theorem 18

According to Lemma 17, the iterates generated by Algorithm 4 satisfy

$$V_{k+1} \leq qV_k + p \sum_{\ell=(k-\tau)_+}^k W_\ell - rW_k$$

for every $k \in \mathbb{N}_0$, where

$$q = 1 - \frac{\gamma(1 - c^2)}{m(1 + \gamma(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}))}, \quad p = \frac{2\gamma^2}{m\tau} \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right),$$

and

$$r = \frac{\gamma}{m} \left(1 - \gamma \left(1 + \frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right) + \frac{2\gamma^2}{m\tau} \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right).$$

To apply Lemma 5, we need to enforce that

$$2\tau + 1 \leq \min \left\{ \frac{1}{1-q}, \frac{r}{p} \right\}.$$

This convergence condition is equivalent to

$$\begin{cases} \gamma(1-c^2)(2\tau+1) & \leq m \left(1 + \gamma \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right), \\ \gamma \left(1 + 5 \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right) & \leq 1. \end{cases} \quad (103)$$

Using the change of variable

$$h = \gamma \left(1 + 5 \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right),$$

the inequalities (103) can be rewritten as

$$\begin{cases} h(1-c^2)(2\tau+1) & \leq m \left((5+h) \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right) + 1\right), \\ h & \leq 1. \end{cases} \quad (104)$$

Since $m \geq 1$ and $c \in (0, 1)$, we have

$$\begin{aligned} (1-c^2)(2\tau+1) & \leq (5+h)\tau + m \\ & = m \left((5+h) \left(\frac{\tau}{m}\right) + 1\right) \\ & \leq m \left((5+h) \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right) + 1\right). \end{aligned}$$

Thus, the inequalities (104) hold for any $h \in [0, 1]$. This shows that if the step-size γ is set to

$$\gamma = \frac{h}{1 + 5 \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)}, \quad h \in (0, 1],$$

then the convergence condition (103) is satisfied. Therefore, by part 2 of Lemma 5, $V_k \leq q^k V_0$ for all $k \in \mathbb{N}_0$. Since $V_k = \mathbb{E}[\|x_k - x^*\|^2]$, it follows that

$$\begin{aligned} \mathbb{E}[\|x_k - x^*\|^2] & \leq \left(1 - \frac{h(1-c^2)}{m \left(1 + (5+h) \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right)}\right)^k \|x_0 - x^*\|^2 \\ & \leq \left(1 - \frac{h(1-c^2)}{m \left(1 + 6 \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right)}\right)^k \|x_0 - x^*\|^2, \end{aligned}$$

where the second inequality is due to that $h \leq 1$. This matches the upper bound (37) on the convergence rate.

Finally, we derive the iteration complexity bound (38). Taking the logarithm of both sides of (37) yields

$$\log(\mathbb{E}[\|x_k - x^*\|^2]) \leq k \log \left(1 - \frac{h(1-c^2)}{m \left(1 + 6 \left(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}\right)\right)}\right) + \log(\|x_0 - x^*\|^2).$$

Since $\log(1+x) \leq x$ for any $x > -1$, it follows that

$$\log(\mathbb{E}[\|x_k - x^*\|^2]) \leq -\left(\frac{h(1-c^2)}{m(1+6(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}))}\right)k + \log(\|x_0 - x^*\|^2).$$

Therefore, for any k satisfying

$$-\left(\frac{h(1-c^2)}{m(1+6(\frac{\tau}{m} + \sqrt{\frac{\tau}{m}}))}\right)k + \log(\|x_0 - x^*\|^2) \leq \log(\epsilon), \quad (105)$$

we have $\log(\mathbb{E}[\|x_k - x^*\|^2]) \leq \log(\epsilon)$, implying that $\mathbb{E}[\|x_k - x^*\|^2] \leq \epsilon$. Rearranging terms in (105) completes the proof.

Appendix F. Proofs for Subsection 4.4

This section provides the proofs for the results presented in Subsection 4.4.

F.1 Proof of Lemma 22

For each $i \in [m]$, let κ_i and κ'_i be two arbitrary consecutive elements of \mathcal{K}_i . From (43), we have

$$[x_{k+1}]_i = T_i\left(\left[x_{s_{i1}, \kappa_i}\right]_1, \dots, \left[x_{s_{im}, \kappa_i}\right]_m\right), \quad k \in [\kappa_i, \kappa'_i).$$

As $t_i(k) = \kappa_i$ for $k \in [\kappa_i, \kappa'_i)$, we obtain

$$[x_{k+1}]_i = T_i\left(\left[x_{s_{i1}, t_i(k)}\right]_1, \dots, \left[x_{s_{im}, t_i(k)}\right]_m\right), \quad k \in [\kappa_i, \kappa'_i).$$

Since κ_i and κ'_i are two arbitrary consecutive elements of \mathcal{K}_i and $0 \in \mathcal{K}_i$ for each i , we can rewrite the asynchronous iteration (43) as

$$[x_{k+1}]_i = T_i\left(\left[x_{s_{i1}, t_i(k)}\right]_1, \dots, \left[x_{s_{im}, t_i(k)}\right]_m\right), \quad k \in \mathbb{N}_0. \quad (106)$$

Let $V_k = \|x_k - x^*\|_{b, \infty}^w$. From the definition of $\|\cdot\|_{b, \infty}^w$, we have

$$\begin{aligned} V_{k+1} &= \max_{i \in [m]} w_i \|[x_{k+1}]_i - [x^*]_i\|_i \\ &\stackrel{(106)}{=} \max_{i \in [m]} w_i \left\| T_i\left(\left[x_{s_{i1}, t_i(k)}\right]_1, \dots, \left[x_{s_{im}, t_i(k)}\right]_m\right) - [x^*]_i \right\|_i \\ &\leq \max_{i \in [m]} \left\| T\left(\left[x_{s_{i1}, t_i(k)}\right]_1, \dots, \left[x_{s_{im}, t_i(k)}\right]_m\right) - x^* \right\|_{b, \infty}^w, \end{aligned}$$

where the inequality follows from the fact that $w_i \|[x]_i\|_i \leq \|x\|_{b, \infty}^w$ for any $x \in \mathbb{R}^d$ and $i \in [m]$. Since T is pseudo-contractive with respect to the block-maximum norm with contraction modulus c , we obtain

$$\begin{aligned} V_{k+1} &\leq c \max_{i \in [m]} \left\| \left(\left[x_{s_{i1}, t_i(k)}\right]_1, \dots, \left[x_{s_{im}, t_i(k)}\right]_m\right) - x^* \right\|_{b, \infty}^w \\ &= c \max_{i \in [m]} \max_{j \in [m]} w_j \left\| \left[x_{s_{ij}, t_i(k)}\right]_j - [x^*]_j \right\|_j, \end{aligned}$$

where the equality follows from the definition of $\|\cdot\|_{b,\infty}^w$. Since $w_j\|[x]_j\|_j \leq \|x\|_{b,\infty}^w$ for any $x \in \mathbb{R}^d$ and $j \in [m]$, we have

$$\begin{aligned} V_{k+1} &\leq c \max_{i \in [m]} \max_{j \in [m]} \left\| x_{s_{ij,t_i(k)}} - x^* \right\|_{b,\infty}^w \\ &= c \max_{i \in [m]} \max_{j \in [m]} V_{s_{ij,t_i(k)}}. \end{aligned} \quad (107)$$

From (44), $k - \tau_k \leq s_{ij,t_i(k)}$ for all $i, j \in [m]$ and $k \in \mathbb{N}_0$. Thus, $s_{ij,t_i(k)} \in \{k - \tau_k, \dots, k\}$, and hence

$$V_{s_{ij,t_i(k)}} \leq \max_{k - \tau_k \leq \ell \leq k} V_\ell.$$

This together with (107) implies that

$$V_{k+1} \leq c \max_{k - \tau_k \leq \ell \leq k} V_\ell,$$

which is the desired result.

F.2 Proof of Theorem 23

According to Lemma 22, the iterates generated by (43) satisfy

$$V_{k+1} \leq qV_k + p \max_{(k - \tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0,$$

with $q = 0$, $p = c$, and

$$\tau_k = k - \min_{i \in [m]} \min_{j \in [m]} s_{ij,t_i(k)}.$$

As $t_i(k) \in \mathcal{K}_i$ for $i \in [m]$ and $k \in \mathbb{N}_0$, it follows from Assumption 8.1 that $\lim_{k \rightarrow \infty} t_i(k) = \infty$. Thus, by Assumption 8.2, $s_{ij,t_i(k)} \rightarrow \infty$ as $k \rightarrow \infty$ for all $i, j \in [m]$. This implies that

$$\lim_{k \rightarrow \infty} k - \tau_k = \infty.$$

Therefore, since $c < 1$, using Lemma 2 with $q = 0$ and $p = c$ completes the proof.

F.3 Proof of Theorem 24

From Assumption 9.2, $k - D \leq s_{ij,k}$ for all $i, j \in [m]$ and $k \in \mathcal{K}_i$. Since $t_i(k) \in \mathcal{K}_i$ for $k \in \mathbb{N}_0$, we have

$$t_i(k) - D \leq s_{ij,t_i(k)}.$$

By Assumption 9.1, $k - B \leq t_i(k)$ for each i and all $k \in \mathbb{N}_0$. Thus, $k - B - D \leq s_{ij,t_i(k)}$, implying that

$$\tau_k = k - \min_{i \in [m]} \min_{j \in [m]} s_{ij,t_i(k)} \leq B + D.$$

It follows from Lemma 22 that the iterates generated by (43) satisfy

$$V_{k+1} \leq c \max_{(k-\tau_k)_+ \leq \ell \leq k} V_\ell, \quad k \in \mathbb{N}_0,$$

with $V_k = \|x_k - x^*\|_{b,\infty}^w$ and $\tau_k \leq B + D$. As $c < 1$, using Lemma (1) with $q = 0$, $p = c$ and $\tau = B + D$ leads to

$$V_k \leq c^{\frac{k}{B+D+1}} V_0, \quad k \in \mathbb{N}_0.$$

The proof is complete.

F.4 Proof of Theorem 25

According to Lemma 22, the iterates generated by (43) satisfy (15) with $V_k = \|x_k - x^*\|_{b,\infty}^w$, $q = 0$, $p = c$, and τ_k defined in (44). Since $q + p = c < 1$, it follows from Lemma 3 that $V_k \leq \Lambda(k)V_0$. This completes the proof.

References

- A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. *Advances in Neural Information Processing Systems*, 24(1):1–9, 2011.
- Y. Arjevani, O. Shamir, and N. Srebro. A tight convergence analysis for stochastic gradient descent with delayed updates. *Algorithmic Learning Theory*, 1(1):111–132, 2020.
- R. Z. Aviv, I. Hakimi, A. Schuster, and K. Y. Levy. Asynchronous distributed learning: Adapting to gradient delays without prior knowledge. *International Conference on Machine Learning*, 1(1):436–445, 2021.
- A. Aytikin, H. R. Feyzmahdavian, and M. Johansson. Analysis and implementation of an asynchronous optimization algorithm for the parameter server. *arXiv preprint arXiv:1610.05507*, 2016.
- G. M. Baudet. Asynchronous iterative methods for multiprocessors. *Journal of the ACM (JACM)*, 25(2):226–244, 1978.
- A. Beck. *First-order Methods in Optimization*. SIAM, 2017.
- D. P. Bertsekas. Distributed asynchronous computation of fixed points. *Mathematical Programming*, 27(1):107–120, 1983.
- D. P. Bertsekas. *Convex Optimization Algorithms*. Athena Scientific, 2015.
- D. P. Bertsekas and D. E. Baz. Distributed asynchronous relaxation methods for convex network flow problems. *SIAM Journal on Control and Optimization*, 25(1):74–85, 1987.
- D. P. Bertsekas and J. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 2015.

- D. P. Bertsekas and J. N. Tsitsiklis. Convergence rate and termination of asynchronous iterative algorithms. *International Conference on Supercomputing*, 1(1):461–470, 1989.
- D. Blatt, A. O. Hero, and H. Gauchman. A convergent incremental gradient method with a constant step size. *SIAM Journal on Optimization*, 18(1):29–51, 2007.
- L. Cannelli, F. Facchinei, V. Kungurtsev, and G. Scutari. Asynchronous parallel algorithms for nonconvex optimization. *Mathematical Programming*, 184(1):121–154, 2020.
- D. Chazan and W. Miranker. Chaotic relaxation. *Linear Algebra and its Applications*, 2(2):199–222, 1969.
- Y. K. Cheung, R. Cole, and Y. Tao. Fully asynchronous stochastic coordinate descent: A tight lower bound on the parallelism achieving linear speedup. *Mathematical Programming*, 1(1):1–63, 2020.
- A. Cohen, A. Daniely, Y. Drori, T. Koren, and M. Schain. Asynchronous stochastic optimization robust to arbitrary delays. *Advances in Neural Information Processing Systems*, 34(1):9024–9035, 2021.
- J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13(1), 2012.
- M. Eisen, A. Mokhtari, and A. Ribeiro. Decentralized quasi-Newton methods. *IEEE Transactions on Signal Processing*, 65(10):2613–2628, 2017.
- O. Fercoq and P. Richtárik. Accelerated, parallel, and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- H. R. Feyzmahdavian and M. Johansson. On the convergence rates of asynchronous iterations. *53rd IEEE Conference on Decision and Control*, 1(1):153–159, 2014.
- H. R. Feyzmahdavian, M. Johansson, and T. Charalambous. Contractive interference functions and rates of convergence of distributed power control laws. *IEEE Transactions on Wireless Communications*, 11(12):4494–4502, 2012.
- H. R. Feyzmahdavian, A. Aytakin, and M. Johansson. A delayed proximal gradient method with linear convergence rate. *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 1(1):1–6, 2014.
- H. R. Feyzmahdavian, A. Aytakin, and M. Johansson. An asynchronous mini-batch algorithm for regularized stochastic optimization. *IEEE Transactions on Automatic Control*, 61(12):3740–3754, 2016.
- R. M. Gower, N. Loizou, X. Qian, A. Sailanbayev, E. Shulgin, and P. Richtárik. SGD: General analysis and improved rates. *International conference on machine learning*, 1(1):5200–5209, 2019.

- M. Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. On the convergence rate of incremental aggregated gradient algorithms. *SIAM Journal on Optimization*, 27(2):1035–1048, 2017.
- M. T. Hale, A. Nedić, and M. Egerstedt. Asynchronous multi-agent primal-dual optimization. *IEEE Transactions on Automatic Control*, 62(9):4421–4435, 2017.
- M. Han and K. Daudjee. Giraph unchained: Barrierless asynchronous parallel execution in pregel-like graph processing systems. *Proceedings of the VLDB Endowment*, 8(9):950–961, 2015.
- R. Hannah and W. Yin. More iterations per second, same quality—why asynchronous algorithms may drastically outperform traditional ones. *arXiv preprint arXiv: 1708.05136*, 2017.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Science & Business Media, 2009.
- Q. Ho, J. Cipar, H. Cui, J. K. Kim, S. Lee, P. B. Gibbons, G. A. Gibson, G. R. Ganger, and E. P. Xing. More effective distributed ML via a stale synchronous parallel parameter server. *Advances in Neural Information Processing Systems*, 1(1):1223–1231, 2013.
- X. Huang, K. Yuan, X. Mao, and W. Yin. An improved analysis and rates for variance reduction under without-replacement sampling orders. *Advances in Neural Information Processing Systems*, 34(1):3232–3243, 2021.
- S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh. SCAFFOLD: Stochastic controlled averaging for federated learning. *International Conference on Machine Learning*, 1(1):5132–5143, 2020.
- A. Koloskova, S. U. Stich, and M. Jaggi. Sharper convergence guarantees for asynchronous SGD for distributed and federated learning. *Advances in Neural Information Processing Systems*, 35(1):17202–17215, 2022.
- G. Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.
- R. Leblond, F. Pedregosa, and S. Lacoste-Julien. ASAGA: Asynchronous parallel SAGA. *Artificial Intelligence and Statistics*, 1(1):46–54, 2017.
- R. Leblond, F. Pedregosa, and S. Lacoste-Julien. Improved asynchronous parallel optimization analysis for stochastic incremental methods. *The Journal of Machine Learning Research*, 19(1):3140–3207, 2018.
- M. Li, D. G. Andersen, and A. Smola. Distributed delayed proximal gradient methods. *NIPS Workshop on Optimization for Machine Learning*, 5(1):1–5, 2013a.
- 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*, 6(1):1–10, 2013b.

- X. Lian, Y. Huang, Y. Li, and J. Liu. Asynchronous parallel stochastic gradient for nonconvex optimization. *Advances in Neural Information Processing Systems*, 28(1):2737–2745, 2015.
- J. Liu and S. J. Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.
- J. Liu, S. J. Wright, C. Ré, V. Bittorf, and S. Sridhar. An asynchronous parallel stochastic coordinate descent algorithm. *International Conference on Machine Learning*, 1(1):469–477, 2014.
- H. Mania, X. Pan, D. Papailiopoulos, B. Recht, K. Ramchandran, and M. I. Jordan. Perturbed iterate analysis for asynchronous stochastic optimization. *SIAM Journal on Optimization*, 27(4):2202–2229, 2017.
- M. Mehyar, D. Spanos, J. Pongsajapan, S. H. Low, and R. M. Murray. Asynchronous distributed averaging on communication networks. *IEEE/ACM Transactions on Networking*, 15(1):512–520, 2007.
- K. Mishchenko, F. Iutzeler, and J. Malick. A distributed flexible delay-tolerant proximal gradient algorithm. *SIAM Journal on Optimization*, 30(1):933–959, 2020.
- K. Mishchenko, F. Bach, M. Even, and B. E. Woodworth. Asynchronous SGD beats mini-batch SGD under arbitrary delays. *Advances in Neural Information Processing Systems*, 35(1):420–433, 2022.
- C. C. Moallemi and B. Van Roy. Convergence of min-sum message-passing for convex optimization. *IEEE Transactions on Information Theory*, 56(1):2041–2050, 2010.
- I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *Mathematical Programming*, 175(1):69–107, 2019.
- Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Science & Business Media, 2013.
- F. Niu, B. Recht, C. Ré, and S. J. Wright. HOGWILD!: A lock-free approach to parallelizing stochastic gradient descent. *Advances in Neural Information Processing Systems*, 24(1):693–701, 2011.
- Z. Peng, Y. Xu, M. Yan, and W. Yin. ARock: An algorithmic framework for asynchronous parallel coordinate updates. *SIAM Journal on Scientific Computing*, 38(5):2851–2879, 2016.
- B. T. Polyak. *Introduction to Optimization*. Optimization Software, Inc., Publications Division, New York, 1987.
- M. Schmidt, N. Le Roux, and F. Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1):83–112, 2017.

- G. Stathopoulos and C. N. Jones. An inertial parallel and asynchronous forward–backward iteration for distributed convex optimization. *Journal of Optimization Theory and Applications*, 182(3):1088–1119, 2019.
- S. U. Stich and S. P. Karimireddy. The error-feedback framework: Better rates for SGD with delayed gradients and compressed updates. *The Journal of Machine Learning Research*, 21(1):9613–9648, 2020.
- T. Sun, R. Hannah, and W. Yin. Asynchronous coordinate descent under more realistic assumption. *International Conference on Neural Information Processing Systems*, 1(1):6183–6191, 2017.
- T. Sun, Y. Sun, D. Li, and Q. Liao. General proximal incremental aggregated gradient algorithms: Better and novel results under general scheme. *Advances in Neural Information Processing Systems*, 32(1):996–1006, 2019.
- P. Tseng and S. Yun. Incrementally updated gradient methods for constrained and regularized optimization. *Journal of Optimization Theory and Applications*, 160(3):832–853, 2014.
- J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.
- L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- N. D. Vanli, M. Gürbüzbalaban, and A. Ozdaglar. Global convergence rate of proximal incremental aggregated gradient methods. *SIAM Journal on Optimization*, 28(2):1282–1300, 2018.
- H. T. Wai, W. Shi, C. Uribe, A. Nedić, and A. Scaglione. Accelerating incremental gradient optimization with curvature information. *Computational Optimization and Applications*, 76(2):347–380, 2020.
- S. J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.
- L. Xiao and T. Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- Y. Zeng, F. Feng, and W. Yin. AsyncQVI: Asynchronous-parallel Q-value iteration for discounted Markov decision processes with near-optimal sample complexity. *International Conference on Artificial Intelligence and Statistics*, 1(1):713–723, 2020.
- H. Zhang, Y. H. Dai, L. Guo, and W. Peng. Proximal-like incremental aggregated gradient method with linear convergence under Bregman distance growth conditions. *Mathematics of Operations Research*, 46(1):61–81, 2021.