



RECENT ADVANCES IN STOCHASTIC APPROXIMATION WITH APPLICATIONS TO OPTIMIZATION AND FIXED POINT PROBLEMS

RAJEEVA L. KARANDIKAR¹, M. VIDYASAGAR^{2,*}

¹Chennai Mathematical Institute, Chennai, India

²Indian Institute of Technology, Hyderabad, India

Abstract. We begin by briefly surveying some results on the convergence of the Stochastic Gradient Descent (SGD) Method, proved in a recent paper. These results are based on viewing SGD as a version of Stochastic Approximation (SA). Ever since its introduction in the classic paper of Robbins and Monro in 1951, SA has become a standard tool for finding a solution of an equation of the form $\mathbf{f}(\theta) = \mathbf{0}$, when only noisy measurements of $\mathbf{f}(\cdot)$ are available. In most situations, every component of the putative solution θ_t is updated at each step t . In some applications in Reinforcement Learning (RL), only one component of θ_t is updated at each t . This is known as asynchronous SA. In this paper, we study Block Asynchronous SA (BASA), in which, at each step t , some but not necessarily all components of θ_t are updated. The theory presented here embraces both conventional (synchronous) SA as well as asynchronous SA, and all in-between possibilities. We provide sufficient conditions for the convergence of BASA, and also prove bounds on the rate of convergence of θ_t to the solution. For the case of conventional SGD, these results reduce to those proved in our companion paper. Then we apply these results to the problem of finding a fixed point of a map with only noisy measurements. This problem arises frequently in RL. We prove sufficient conditions for convergence as well as estimates for the rate of convergence.

Keywords. Block asynchronous updating; Q-learning; Rates of convergence; Reinforcement learning; Stochastic approximation.

2020 MSC. 62L20, 60G17, 93D05.

1. INTRODUCTION

1.1. Background. Ever since its introduction in the classic paper of Robbins and Monro [30], Stochastic Approximation (SA) has become a standard tool in many problems in applied mathematics. It is worth noting that the phrase “Stochastic Approximation” was coined in [30]. As stated in [30], the original problem formulation in SA was to find a solution to an equation of

*Corresponding author.

E-mail address: rlk@cmi.ac.in (R.L. Karandikar), m.vidyasagar@iith.ac.in (M. Vidyasagar)

Received May 3, 2024; Accepted August 12, 2024.

the form¹

$$f(\theta) = c,$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$, c is a specified constant, and one has access only to noisy measurements of the function. Obviously, one can redefine f and assume that $c = 0$, without loss of generality. Almost at once, the approach was extended to finding a stationary point of a \mathcal{C}^1 -function $J(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ in [20], and to the case where $J(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}$ in [4]. Other early contributions are [10, 11]. In the early papers, SA was analyzed under extremely stringent assumptions on the function, and on the measurement error. With the passage of time, subsequent researchers have substantially relaxed the assumptions.

Over the years, SA has become a standard tool for analyzing the behavior of stochastic algorithms in a variety of areas, out of which two topics are the focus in the present paper, namely: optimization, and finding a fixed point of a contractive map, which arises frequently in Reinforcement Learning (RL). The aim of the present paper is two-fold: First, we survey some known results in the theory of SA, including some results due to the present authors. Second, we present some new results on so-called Block Asynchronous SA, or BASA.

1.2. Problem Formulation. Suppose $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is some function. It is desired to find a solution to the equation $\mathbf{f}(\theta^*) = \mathbf{0}$, when only noisy measurements of $\mathbf{f}(\cdot)$ are available. An iterative approach is adopted to solve this equation. Let t denote the iteration count, and choose the initial guess θ_0 either in a deterministic or a random fashion. At time (or step) $t + 1$, the available measurement is $\mathbf{f}(\theta_t) + \xi_{t+1}$, where ξ_{t+1} is variously referred to as the measurement error or the “noise.” Both phrases are used interchangeably in this paper. The current guess θ_t is updated via the formula

$$\theta_{t+1} = \theta_t + \alpha_t \circ [\mathbf{f}(\theta_t) + \xi_{t+1}], \quad (1.1)$$

where $\alpha_t \in (0, \infty)^d$ is called the **step size vector**, and \circ denotes the Hadamard product.² If $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a map and it is desired to find a fixed point of $\mathbf{g}(\cdot)$, when we can define $\mathbf{f}(\theta) = \mathbf{g}(\theta) - \theta$. This causes (1.1) to become

$$\theta_{t+1} = (\mathbf{1}_d - \alpha_t) \circ \theta_t + \alpha_t \circ [\mathbf{g}(\theta_t) + \xi_{t+1}], \quad (1.2)$$

where $\mathbf{1}_d$ denotes the column vector of d ones. In this case, it is customary to restrict α_t to belong to $(0, 1)^d$ instead of $(0, \infty)^d$. Then each component of θ_{t+1} is a convex combination of the corresponding components of θ_t and the noisy measurement of $\mathbf{g}(\theta_t)$. If $J : \mathbb{R}^d \rightarrow \mathbb{R}$ is a \mathcal{C}^1 -function, and it is desired to find a stationary point of it, then we can define $\mathbf{f}(\theta) = -\nabla J(\theta)$, in which case (1.1) becomes

$$\theta_{t+1} = \theta_t + \alpha_t \circ [-\nabla J(\theta_t) + \xi_{t+1}]. \quad (1.3)$$

The choice $\mathbf{f}(\theta) = -\nabla J(\theta)$ instead of $\nabla J(\theta)$ is used when the objective is to *minimize* $J(\cdot)$, and $J(\cdot)$ is convex, at least in a neighborhood of the minimum. If the objective is to maximize $J(\cdot)$, then one would choose $\mathbf{f}(\theta) = \nabla J(\theta)$.

¹For the convenience of the reader, all results cited from the literature are stated in the notation used in the present paper, which may differ from the original paper.

²Recall that if \mathbf{a}, \mathbf{b} are vectors of equal dimension, then their **Hadamard product** $\mathbf{c} = \mathbf{a} \circ \mathbf{b}$ is defined by $c_i := a_i b_i$ for all i .

What is described above is the “core” problem formulation. Several variations are possible, depending on the objective of the analysis, the nature of the step size vector, and the nature of the error vector ξ_{t+1} . Some of the most widely studied variations are described next.

Objectives of the Analysis. Historically, the majority of the literature is devoted to showing that the iterations converge *in expectation* to a solution of the equation $\mathbf{f}(\theta) = \mathbf{0}$ (or its modification for fixed point and stationarity problems). This is the objective in [21] and other subsequent papers. In recent times, the emphasis has shifted towards proving that the iterations converge *almost surely* to the desired limit. Since any stochastic algorithm such as (1.3) generates a *single sample path*, it is very useful to know that almost every run of the algorithm leads to the desired outcome.

Another possibility is **convergence in probability**. Suppose $\theta_t \rightarrow \theta^*$ in probability, and define

$$q(t, \varepsilon) := \Pr\{\|\theta_t - \theta^*\|_2 > \varepsilon\}. \quad (1.4)$$

The objective is to derive suitable conditions under which, $q(t, \varepsilon) \rightarrow 0$ as $t \rightarrow \infty$ for each $\varepsilon > 0$, and if possible, to derive explicit upper bounds for $q(t, \varepsilon)$. Some authors refer to such bounds as “high probability bounds.” The advantage of bounds on $q(t, \varepsilon)$ is that they are applicable for *all* t (or at least, for all *sufficiently large* t), and not just when $t \rightarrow \infty$. For this reason, some authors refer to the derivation of such bounds as **finite-time SA**. Some contributions in this direction are [3, 8, 28, 31, 35]. We do not discuss FTSA in the paper. The interested reader is referred to the above-cited papers and the references therein.

Step Size Sequences. Next we discuss various options for the step size vector α_t , which is allowed to be random. In all cases, it is assumed that there is a *scalar deterministic* sequence $\{\beta_t\}$ taking values in $(0, \infty)$, or in $(0, 1)^d$ in the case of (1.2). We will discuss three commonly used variants of SA, namely: synchronous (also called fully synchronous), asynchronous, and block asynchronous. In **synchronous SA**, one chooses $\alpha_t = \beta_t \mathbf{1}_d$. Thus, in (1.1), *the same* step size β_t is applied to *every* component of θ_t . In block asynchronous SA (or BASA), there are d different $\{0, 1\}$ -valued stochastic processes, denoted by $\kappa_t^i, i \in [d]$, called the “update” processes. Then the i -th component of θ_t is updated only if $\kappa_t^i = 1$. To put it another way, define the “update set” as $S_t := \{i \in [d] : \kappa_t^i = 1\}$. Then $\alpha_t^i = 0$ if $i \notin S_t$. However, this raises the question as to what α_t^i is for $i \in S_t$. Two options are suggested in the literature, known as the “global” clock and the “local” clock respectively. This distinction was first suggested in [5]. If a global clock is used, then $\alpha_t^i = \beta_t$. To define the step size when a local clock is used, first define

$$v_t^i := \sum_{\tau=0}^t \kappa_\tau^i. \quad (1.5)$$

Thus v_t^i counts the number of times that θ_t^i is updated, and is referred to as the “counter” process. Then the step size is defined as

$$\alpha_t^i := \beta_{v_t^i}. \quad (1.6)$$

The distinction between global and local clocks can be briefly summarized as follows: When a global clock is used, every component of θ_t that gets updated has exactly the same step size, namely β_t , while the other components have a step size of zero. When a local clock is used, among the components of θ_t that get updated at time t , different components may have different step sizes. An important variant of BASA is asynchronous SA (ASA). This phrase

was apparently first used in [33], in the context of proving the convergence of the Q -learning algorithm from Reinforcement Learning (RL). In ASA, *exactly one* component of θ_t is updated at each t . This can be represented as follows: Let $\{N_t\}$ be an integer-valued stochastic process taking values in $[d]$. Then, at time t , the update set S_t is the singleton $\{N_t\}$. The counter process v_t^i is now defined via

$$v_t^i = \sum_{\tau=0}^t I_{\{N_\tau=i\}},$$

where I denotes the indicator process. The step size can either be β_t if a global clock is used, or $\beta_{v_t^i}$ if a local clock is used. In [5], the author analyzes the convergence of ASA with both global as well as local clocks. In the Q -learning algorithm introduced in [36], the update is asynchronous (one component at a time) and a global clock is used. In [33], where the phrase ASA was first introduced, the convergence of ASA is proved under some assumptions which include Q -learning as a special case. Accordingly, the author uses a global clock in the formulation of ASA. In [12], the authors use a local clock to study the rate of convergence of Q -learning.

Error Vector. Next we discuss the assumptions made on the error vector ξ_{t+1} . To state the various assumptions precisely, let θ_0^t denote $(\theta_0, \dots, \theta_t)$, and define α_0^t and ξ_1^t analogously; note that there is no ξ_0 . Let \mathcal{F}_t denote the σ -algebra generated by $\theta_0, \alpha_0^t, \xi_1^t$, and observe that $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$. Thus $\{\mathcal{F}_t\}_{t \geq 0}$ is a filtration; Now (1.1) makes it clear that θ_t is measurable with respect to \mathcal{F}_t , denoted by $\theta_t \in \mathcal{M}(\mathcal{F}_t)$. Given an \mathbb{R}^d -valued random variable X , let $E_t(X)$ denote $E(X|\mathcal{F}_t)$, the conditional expectation of X with respect to \mathcal{F}_t , and let $CV_t(X)$ denote the conditional variance of X , defined as

$$CV_t(X) = E_t(\|X - E_t(X)\|_2^2) = E_t(X^2) - [E_t(X)]^2.$$

An important ingredient in SA theory is the set of assumptions imposed on the two entities $E_t(\xi_{t+1})$ and $CV_t(\xi_{t+1})$. We begin with $E_t(\xi_{t+1})$. The simplest assumptions are that

$$E_t(\xi_{t+1}) = \mathbf{0}, \forall t, \quad (1.7)$$

and that there exists a constant M such that

$$CV_t(\xi_{t+1}) \leq M, \forall t. \quad (1.8)$$

where the equality and the bound hold almost surely. To avoid tedious repetition, the phrase “almost surely” is omitted hereafter, unless it is desirable to state it explicitly. Equation (1.7) implies that $\{\xi_t\}$ is a martingale difference sequence with respect to the filtration $\{\mathcal{F}_t\}$. Equation (1.7) further means that $\mathbf{f}(\theta_t) + \xi_{t+1}$ provides an *unbiased* measurement of $\mathbf{f}(\theta_t)$. In (1.9), the bound on $CV(\xi_{t+1})$ is not just uniform over t , *but also uniform over θ_t* . Over time, the assumptions on both $E_t(\xi_{t+1})$ and $CV_t(\xi_{t+1})$ have been relaxed by successive authors. The most general set of conditions to date are found in [18],³ and are as follows: There exist sequences of constants μ_t and M_t such that

$$\|E_t(\xi_{t+1})\|_2 \leq \mu_t(1 + \|\theta_t\|_2), \forall t. \quad (1.9)$$

$$CV_t(\xi_{t+1}) \leq M_t(1 + \|\theta_t\|_2^2), \forall t. \quad (1.10)$$

In [18], the following are established:

³This paper is currently under final review by *Journal of Optimization Theory and Applications*.

(1) Suppose

$$\sum_{t=0}^{\infty} \alpha_t^2 < \infty, \sum_{t=0}^{\infty} \alpha_t \mu_t < \infty, \sum_{t=0}^{\infty} \alpha_t^2 M_t^2 < \infty.$$

Then the iterations $\{\theta_t\}$ are bounded almost surely.

(2) If, in addition, $\sum_{t=0}^{\infty} \alpha_t = \infty$, then θ_t converges almost surely to the unique solution of (1.1).

Thus, by suitably tuning the step size sequence, bounds of the form (1.9) and (1.10) can be accommodated. The literature review in [18, Section 1.1] contains details of the various intermediate stages between (1.7)–(1.8) and (1.9)–(1.10), and the relevant publications. A condensed version of it is reproduced in Section 2.1. The reader is also directed to [24] for a partial survey that is up to date until its date of publication, 2003.

Methods of Analysis. There are two broad approaches to the analysis of SA, which might be called the ODE approach and the martingale approach. In the ODE approach, it is shown that, as the step sizes $\alpha_t \rightarrow 0$, the stochastic sample paths of (1.1) “converge” to the (deterministic) solution trajectories of the associated ODE $\dot{\theta} = \mathbf{f}(\theta)$. This approach is introduced in [9, 21, 26]. Book-length treatments of the ODE approach can be found in [2, 7, 22, 23]. The Kushner-Clark condition [22] is not a directly verifiable condition, but one needs to fall back on martingale or similar assumptions (such as ‘mixingale’) on noise to verify it. The martingale method was pioneered in [14], and independently discovered and enhanced in [29]. In this approach, the stochastic process $\{\theta_t\}$ is directly analyzed without recourse to any ODE. Conclusions about the behavior of this stochastic process are drawn using the theory of supermartingales. The two methods complement each other. A typical theorem based on the ODE approach states that *if* the iterations remain bounded almost surely, then convergence takes place. Often the boundedness (also called “stability”) can be established using other methods. Also, the ODE approach can address the situation where the equation has multiple solutions. In contrast, in the martingale approach, both the boundedness and the convergence of the iterations can be established simultaneously. An important paper in the ODE approach is [6], in which the boundedness of the iterations is a conclusion and not a hypothesis.

1.3. Contributions of the Paper. After the survey of the Stochastic Gradient method, the emphasis in the paper is on the finding the solution of a fixed-point equation of the following form: Suppose \mathbf{h} maps the sequence space $(\mathbb{R}^d)^{\mathbb{N}}$ into itself. The objective is to find a fixed point $\mathbf{x}^* \in (\mathbb{R}^d)^{\mathbb{N}}$ such that

$$\mathbf{h}_t(\mathbf{x}^*) = \mathbf{x}_t^*, \forall t \geq 0. \quad (1.11)$$

This part of the paper consists of an analysis of Block (or Batch) Asynchronous SA, or BASA, for finding a solution to (1.11). Suppose $\mathbf{h}(\cdot)$ is a memoryless contraction, in the sense that

$$\mathbf{h}_t(\mathbf{x}) = \mathbf{g}(\mathbf{x}_t)$$

for some map $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ which is a contraction in the ℓ_{∞} -norm. Then the formulation reduces to (1.2). But we also the more general case where \mathbf{h} has memory, delays, etc. Towards this end, we begin by analyzing the convergence of “intermittently updated” processes of the form

$$w_{t+1} = (1 - \alpha_t \kappa_t) w_t + \alpha_t \kappa_t \xi_{t+1},$$

where $\{w_t\}$ is an \mathbb{R} -valued stochastic process, $\{\xi_t\}$ is the measurement error, $\{\alpha_t\}$ is a $(0, 1)$ -valued “step size” process, and $\{\kappa_t\}$ is a $\{0, 1\}$ -valued “update” process. For this formulation,

we derive sufficient conditions for convergence, as well as bounds on the *rate* of convergence. We study both the use of both a local clock as well as a global clock, a distinction first introduced in [5]. This formulation is a precursor to the full BASA formulation of (1.2), where again we derive both sufficient conditions for convergence, and bounds on the *rate* of convergence.

1.4. Scope and Organization of the Paper. This paper contains a survey of some results due to the present authors, and some new results. In Section 2, various results from [18] are stated without proof; these results pertain to the convergence of the synchronous SA algorithm, when the error signal ξ_{t+1} satisfies the bounds (1.9) and (1.10). These are the most general assumptions to date. In Section 3, we survey some applications of these convergence results to the stochastic gradient method. The results in [18] make the least restrictive assumptions on the measurement error. These two sections comprise the survey part of the paper.

In Section 4, we commence presenting some new results. Specifically, we study Block (or Batch) Asynchronous SA, denoted by BASA, as described in (1.2). The focus is on finding a fixed point of a map $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ which is a contraction in the ℓ_∞ -norm, or a scaled version thereof. While this problem arises in Reinforcement Learning in several situations, finding fixed points is a pervasive application of stochastic approximation. The novelties here are that (i) we permit a completely general model for choosing the coordinates of θ_t to be updated at time t , and (ii) we also derive bounds on the *rate* of convergence.

2. SYNCHRONOUS STOCHASTIC APPROXIMATION

2.1. Historical Review. We begin with the classical results, starting with [30] which introduced the SA algorithm for the scalar case where $d = 1$. However, we state it here for the multidimensional case. In that paper, the update equation is (1.1), and the error ξ_{t+1} is assumed to satisfy the following assumptions (though this notation is not used in that paper)

$$E_t(\xi_{t+1}) = \mathbf{0}, CV_t(\xi_{t+1}) \leq M^2 \quad (2.1)$$

for some finite constant M . The first assumption implies that $\{\xi_{t+1}\}$ is a martingale difference sequence, and also that $\mathbf{f}(\theta_t) + \xi_{t+1}$ is an *unbiased* measurement of $\mathbf{f}(\theta_t)$. The second assumption means that the conditional variance of the error is globally bounded, both as a function of θ_t and as a function of t . With the assumptions in (2.1), along with some assumptions on the function $\mathbf{f}(\cdot)$, it is shown in [30] that θ_t converges to a solution of $\mathbf{f}(\theta^*) = \mathbf{0}$, provided the step size sequence satisfies the **Robbins-Monro (RM)** conditions

$$\sum_{t=0}^{\infty} \alpha_t^2 < \infty, \sum_{t=0}^{\infty} \alpha_t = \infty. \quad (2.2)$$

This approach was extended in [20] to finding a stationary point of a \mathcal{C}^1 function $J : \mathbb{R} \rightarrow \mathbb{R}$, that is, a solution to $\nabla J(\theta) = \mathbf{0}$,⁴ using an *approximate gradient* of $J(\cdot)$. The specific formulation used in [20] is

$$h_{t+1} := \frac{J(\theta_t + c_t \Delta + \xi_{t+1}^+) - J(\theta_t - c_t \Delta + \xi_{t+1}^-)}{2c_t} \approx \nabla J(\theta_t). \quad (2.3)$$

⁴Strictly speaking, we should use $J'(\theta)$ for the scalar case. But we use vector notation to facilitate comparison with later formulas.

where c_t is called the **increment**, Δ is some fixed number, and ξ_{t+1}^+ , ξ_{t+1}^- are the measurement errors. This terminology “increment” is not standard but is used here. As is standard in such a setting, it is assumed that $\nabla J(\cdot)$ is globally Lipschitz-continuous with constant L . For simplicity, it is common to assume that these sequences are i.i.d. and also independent of each other, with zero mean and finite variance M^2 . We too do the same. In order to make the expression a better and better approximation to the true $\nabla J(\theta_t)$, the increment c_t must approach zero as $t \rightarrow \infty$. Note that there are *two* sources of error in (2.3). First, even if the errors ξ_{t+1}^\pm are zero, the first-order difference is not exactly equal to the gradient $\nabla J(\theta_t)$. Second, the presence of the measurement errors ξ_{t+1}^\pm introduces an additional error term. To analyze this, let us define

$$\mathbf{z}_t = E_t(\mathbf{h}_{t+1}), \mathbf{x}_t = \mathbf{z}_t - \nabla J(\theta_t), \zeta_{t+1} = \mathbf{h}_{t+1} - \mathbf{z}_t. \quad (2.4)$$

In this case, the error term satisfies

$$\|E_t(\zeta_{t+1})\|_2 \leq Lc_t, CV_t(\zeta_{t+1}) \leq M^2/(2c_t^2). \quad (2.5)$$

These conditions are more general than in (2.1). For this situation, in the scalar case, it was shown in [20] that θ_t converges to a stationary point of $J(\cdot)$ if the Kiefer-Wolfowitz-Blum (KWB) conditions

$$c_t \rightarrow 0, \sum_{t=0}^{\infty} (\alpha_t^2/c_t^2) < \infty, \sum_{t=0}^{\infty} \alpha_t c_t < \infty, \sum_{t=0}^{\infty} \alpha_t = \infty \quad (2.6)$$

are satisfied. This approach was extended to the multidimensional case in [4], and it is shown that the same conditions also ensure convergence when $d > 1$. Note that the conditions automatically imply the finiteness of the sum of α_t^2 .

Now we summarize subsequent results. It can be seen from Theorem 2.5 below that in the present paper, the error ξ_{t+1} is assumed to satisfy the following assumptions:

$$\|E_t(\xi_{t+1})\|_2 \leq \mu_t(1 + \|\theta_t\|_2), \quad (2.7)$$

$$CV_t(\xi_{t+1}) \leq M_t^2(1 + \|\theta_t\|_2^2), \quad (2.8)$$

where θ_t is the current iteration. It can be seen that the above assumptions extend (2.5) in several ways. First, the conditional expectation is allowed to grow as an affine function of $\|\theta_t\|_2$, for each fixed t . Second, the conditional variance is also allowed to grow as a quadratic function of $\|\theta_t\|_2$, for each fixed t . Third, while the coefficient μ_t is required to approach zero, the coefficient M_t can grow without bound as a function of t . We are not aware of any other paper that makes such general assumptions. However, there are several papers wherein the assumptions on ξ_{t+1} are intermediate between (2.1) and (2.5). We attempt to summarize a few of them next. For the benefit of the reader, we state the results using the notation of the present paper.

In [21], the author considered a recursion of the form

$$\theta_{t+1} = \theta_t - \alpha_t \nabla J(\theta_t) + \alpha_t \xi_{t+1} + \alpha_t \beta_{t+1},$$

where $\beta_t \rightarrow \mathbf{0}$ as $t \rightarrow \infty$. Here, the sequence $\{\xi_{t+1}\}$ is **not** assumed to be a martingale difference sequence. Rather, it is assumed to satisfy a different set of conditions, referred to as the Kushner-Clark conditions; see [21, A5]. It is then shown that if the error sequence $\{\xi_{t+1}\}$ satisfies (2.1), i.e., is a martingale difference sequence, then Assumption (A5) holds. Essentially the same formulation is studied in [27]. The same formulation is also studied [7, Section 2.2], where (2.1) holds, and $\beta_t \rightarrow \mathbf{0}$ as $t \rightarrow \infty$. In [32], it is assumed only that $\limsup_t \beta_t < \infty$. In all

cases, it is shown that θ_t converges to a solution of $\mathbf{f}(\theta^*) = \mathbf{0}$, *provided* the iterations remain bounded almost surely. Therefore, the boundedness of the iterations is established via separate arguments.

In all of the above references, the bound on $CV_t(\xi_{t+1})$ is as in (2.1). We are aware of only one paper when the bound on $CV_t(\xi_{t+1})$ is akin to that in (2.8). In [16], the authors study smooth convex optimization. They assume that the estimated gradient is *unbiased*, so that $\mu_t = 0$ for all t . However, an analog of (2.8) is assumed to hold, which is referred to as “state-dependent noise.” See [16, Assumption (SN)]. In short, there is no paper wherein the assumptions on the error are as general as in (2.7) and (2.8).

2.2. Convergence Theorems. In this subsection, we state without proof some results from [18] on the convergence of SA, when the measurement error satisfies (2.7) and (2.8), which are the most general assumptions to date. In addition to proving convergence, we also provide a general framework for estimating the *rate* of convergence. The applications of these convergence theorems to stochastic gradient descent (SGD) are discussed in Section 3.

The theorems proved in [18] make use of the following classic “almost supermartingale theorem” of Robbins-Siegmund [29, Theorem 1]. The result is also proved as [2, Lemma 2, Section 5.2]. Also see a recent survey paper as [13, Lemma 4.1]. The theorem states the following:

Lemma 2.1. *Suppose $\{z_t\}, \{f_t\}, \{g_t\}, \{h_t\}$ are stochastic processes taking values in $[0, \infty)$, adapted to some filtration $\{\mathcal{F}_t\}$, satisfying*

$$E_t(z_{t+1}) \leq (1 + f_t)z_t + g_t - h_t \text{ a.s., } \forall t, \quad (2.9)$$

where, as before, $E_t(z_{t+1})$ is a shorthand for $E(z_{t+1}|\mathcal{F}_t)$. Then, on the set

$$\Omega_0 := \left\{ \omega : \sum_{t=0}^{\infty} f_t(\omega) < \infty \right\} \cap \left\{ \omega : \sum_{t=0}^{\infty} g_t(\omega) < \infty \right\},$$

we have that $\lim_{t \rightarrow \infty} z_t$ exists, and in addition, $\sum_{t=0}^{\infty} h_t(\omega) < \infty$. In particular, if $P(\Omega_0) = 1$, then $\{z_t\}$ is bounded almost surely, and $\sum_{t=0}^{\infty} h_t(\omega) < \infty$ almost surely.

The first convergence result, namely Theorem 2.4 below, is a fairly straight-forward, but useful, extension of Lemma 2.1. It is based on a concept which is introduced in [14] but without giving it a name. The formal definition is given in [34, Definition 1]:

Definition 2.2. A function $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is said to **belong to Class \mathcal{B}** if $\eta(0) = 0$, and in addition


$$\inf_{\varepsilon \leq r \leq M} \eta(r) > 0, \forall 0 < \varepsilon < M < \infty.$$

Note $\eta(\cdot)$ is *not* assumed to be monotonic, or even to be continuous. However, if $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, then $\eta(\cdot)$ belongs to Class \mathcal{B} if and only if (i) $\eta(0) = 0$, and (ii) $\eta(r) > 0$ for all $r > 0$. Such a function is called a “class P function” in [15]. Thus a Class \mathcal{B} function is slightly more general than a function of Class \mathcal{P} .

As example of a function of Class \mathcal{B} is given next:

Example 2.3. Define a function $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ by

$$\phi(\theta) = \begin{cases} \theta, & \text{if } \theta \in [0, 1], \\ e^{-(\theta-1)}, & \text{if } \theta > 1. \end{cases}$$

FIGURE 1. An illustration of a function in Class \mathcal{B}

Then ϕ belongs to Class \mathcal{B} . A sketch of the function $\phi(\cdot)$ is given in Figure 1. Note that, if we were to change the definition to:

$$\phi(\theta) = \begin{cases} \theta, & \text{if } \theta \in [0, 1], \\ 2e^{-(\theta-1)}, & \text{if } \theta > 1, \end{cases}$$

then $\phi(\cdot)$ would be discontinuous at $\theta = 1$, but it would still belong to Class \mathcal{B} . Thus a function need not be continuous to belong to Class \mathcal{B} .

Now we present our first convergence theorem, which is an extension of Lemma 2.1. This theorem is used to establish the convergence of stochastic gradient methods for nonconvex functions, as discussed in Section 3. It is [18, Theorem 1].

Theorem 2.4. *Suppose $\{z_t\}, \{f_t\}, \{g_t\}, \{h_t\}, \{\alpha_t\}$ are $[0, \infty)$ -valued stochastic processes defined on some probability space (Ω, Σ, P) , and adapted to some filtration $\{\mathcal{F}_t\}$. Suppose further that*

$$E_t(z_{t+1}) \leq (1 + f_t)z_t + g_t - \alpha_t h_t \text{ a.s., } \forall t. \quad (2.10)$$

Define

$$\Omega_0 := \{\omega \in \Omega : \sum_{t=0}^{\infty} f_t(\omega) < \infty \text{ and } \sum_{t=0}^{\infty} g_t(\omega) < \infty\}, \quad (2.11)$$

$$\Omega_1 := \{\sum_{t=0}^{\infty} \alpha_t(\omega) = \infty\}. \quad (2.12)$$

Then

- (1) Suppose that $P(\Omega_0) = 1$. Then the sequence $\{z_t\}$ is bounded almost surely, and there exists a random variable W defined on (Ω, Σ, P) such that $z_t(\omega) \rightarrow W(\omega)$ almost surely.
- (2) Suppose that, in addition to $P(\Omega_0) = 1$, it is also true that $P(\Omega_1) = 1$. Then

$$\liminf_{t \rightarrow \infty} h_t(\omega) = 0 \quad \forall \omega \in \Omega_0 \cap \Omega_1. \quad (2.13)$$

Further, suppose there exists a function $\eta(\cdot)$ of Class \mathcal{B} such that $h_t(\omega) \geq \eta(z_t(\omega))$ for all $\omega \in \Omega_0$. Then $z_t(\omega) \rightarrow 0$ as $t \rightarrow \infty$ for all $\omega \in \Omega_0$.

Next we study a linear stochastic recurrence relation. Despite its simplicity, it is a key tool in establishing the convergence of Stochastic Gradient Descent (SGD) studied in Section 3. Suppose θ_0 is an \mathbb{R}^d -valued random variable, and that $\{\zeta_t\}_{t \geq 1}$ is an \mathbb{R}^d -valued stochastic process. Define $\{\theta_t\}_{t \geq 1}$ recursively by

$$\theta_{t+1} = (1 - \alpha_t)\theta_t + \alpha_t \zeta_{t+1}, t \geq 0, \quad (2.14)$$

where $\{\alpha_t\}_{t \geq 0}$ is another $[0, 1]$ -valued stochastic process. Define $\{\mathcal{F}_t\}$ to be the filtration where \mathcal{F}_t is the σ -algebra generated by $\theta_0, \alpha_0^t, \zeta_1^t$. Note that (2.14) is of the form (1.2) with $\mathbf{g}(\theta) \equiv \mathbf{0}$. Hence $\mathbf{g}(\cdot)$ has the unique fixed point $\mathbf{0}$, and we would want that $\theta_t \rightarrow \mathbf{0}$ as $t \rightarrow \infty$. Theorem 2.5 below is a ready consequence of applying [18, Theorem 3] to the function $J(\theta) = (1/2)\|\theta\|_2^2$.

Theorem 2.5. *Suppose there exist sequences of constants $\{\mu_t\}, \{M_t\}$ such that, for all $t \geq 0$ we have*

$$\|E_t(\zeta_{t+1})\|_2 = \|\eta_t\|_2 \leq \mu_t(1 + \|\theta_t\|_2), \quad (2.15)$$

$$CV_t(\zeta_{t+1}) = E_t(\|\psi_{t+1}\|_2^2) \leq M_t^2(1 + \|\theta_t\|_2^2). \quad (2.16)$$

Under these conditions, if

$$\sum_{t=0}^{\infty} \alpha_t^2 < \infty, \sum_{t=0}^{\infty} \mu_t \alpha_t < \infty, \sum_{t=0}^{\infty} M_t^2 \alpha_t^2 < \infty, \quad (2.17)$$

then $\{\theta_t\}$ is bounded, and $\|\theta_t\|_2$ converges to an \mathbb{R}^d -valued random variable. If in addition,

$$\sum_{t=0}^{\infty} \alpha_t = \infty, \quad (2.18)$$

then $\theta_t \rightarrow \mathbf{0}$.

Next, we state an extension of Theorem 2.4 that provides an estimate on rates of convergence. For the purposes of this paper, we use the following definition inspired by [25].

Definition 2.6. Suppose $\{Y_t\}$ is a stochastic process, and $\{f_t\}$ is a sequence of positive numbers. We say that

- (1) $Y_t = O(f_t)$ if $\{Y_t/f_t\}$ is bounded almost surely.
- (2) $Y_t = \Omega(f_t)$ if Y_t is positive almost surely, and $\{f_t/Y_t\}$ is bounded almost surely.
- (3) $Y_t = \Theta(f_t)$ if Y_t is both $O(f_t)$ and $\Omega(f_t)$.
- (4) $Y_t = o(f_t)$ if $Y_t/f_t \rightarrow 0$ almost surely as $t \rightarrow \infty$.

The next theorem is a modification of Theorem 2.4 that provides bounds on the *rate* of convergence. It is [18, Theorem 2].

Theorem 2.7. *Suppose $\{z_t\}, \{f_t\}, \{g_t\}, \{\alpha_t\}$ are stochastic processes defined on some probability space (Ω, Σ, P) , taking values in $[0, \infty)$, adapted to some filtration $\{\mathcal{F}_t\}$. Suppose further that*

$$E_t(z_{t+1}) \leq (1 + f_t)z_t + g_t - \alpha_t z_t \quad \forall t, \quad (2.19)$$

where

$$\sum_{t=0}^{\infty} f_t(\omega) < \infty, \sum_{t=0}^{\infty} g_t(\omega) < \infty, \sum_{t=0}^{\infty} \alpha_t(\omega) = \infty.$$

Then $z_t = o(t^{-\lambda})$ for every $\lambda \in (0, 1]$ such that (i) there exists a $T < \infty$ such that

$$\alpha_t(\omega) - \lambda t^{-1} \geq 0 \quad \forall t \geq T, \quad (2.20)$$

and (ii)

$$\sum_{t=0}^{\infty} (t+1)^\lambda g_t(\omega) < \infty, \sum_{t=0}^{\infty} [\alpha_t(\omega) - \lambda t^{-1}] = \infty. \quad (2.21)$$

With this motivation, we present a refinement of Theorem 2.5. Again, this is obtained by applying [18, Theorem 4] to the function $J(\theta) = (1/2)\|\theta\|_2^2$.

Theorem 2.8. *Let various symbols be as in Theorem 2.5. Further, suppose there exist constants $\gamma > 0$ and $\delta \geq 0$ such that⁵*

$$\mu_t = O(t^{-\gamma}), M_t = O(t^\delta),$$

where we take $\gamma = 1$ if $\mu_t = 0$ for all sufficiently large t , and $\delta = 0$ if M_t is bounded. Choose the step-size sequence $\{\alpha_t\}$ as $O(t^{-(1-\phi)})$ and $\Omega(t^{-(1-c)})$ where ϕ is chosen to satisfy

$$0 < \phi < \min\{0.5 - \delta, \gamma\},$$

and $c \in (0, \phi]$. Define

$$\nu := \min\{1 - 2(\phi + \delta), \gamma - \phi\}. \quad (2.22)$$

Then $\|\theta_t\|_2^2 = o(t^{-\lambda})$ for every $\lambda \in (0, \nu)$. In particular, if $\mu_t = 0$ for all t and M_t is bounded with respect to t , then we can take $\nu = 1 - 2\phi$.

3. APPLICATIONS TO STOCHASTIC GRADIENT DESCENT

In this section, we reprise some relevant results from [18] on the convergence of the Stochastic Gradient Method. Specifically, we analyze the convergence of the Stochastic Gradient Descent (SGD) algorithm in the form

$$\theta_{t+1} = \theta_t - \alpha_t \mathbf{h}_{t+1}, \quad (3.1)$$

where \mathbf{h}_{t+1} is a stochastic gradient. For future use, let us define

$$\mathbf{z}_t = E_t(\mathbf{h}_{t+1}), \mathbf{x}_t = \mathbf{z}_t - \nabla J(\theta_t), \zeta_{t+1} = \mathbf{h}_{t+1} - \mathbf{z}_t. \quad (3.2)$$

The last equation in (3.2) implies that $E_t(\zeta_{t+1}) = \mathbf{0}$. Therefore

$$E_t(\|\mathbf{h}_{t+1}\|_2^2) = \|\mathbf{z}_t\|_2^2 + E_t\|\zeta_{t+1}\|_2^2. \quad (3.3)$$

We make two assumptions about the stochastic gradient: **Assumption:** There exist sequences of constants $\{\mu_t\}$ and $\{M_t\}$ such that

$$\|\mathbf{x}_t\|_2 \leq \mu_t[1 + \|\nabla J(\theta_t)\|_2], \quad \forall \theta_t \in \mathbb{R}^d, \quad \forall t, \quad (3.4)$$

$$E_t(\|\zeta_{t+1}\|_2^2) \leq M_t^2[1 + J(\theta_t)], \quad \forall \theta_t \in \mathbb{R}^d, \quad \forall t. \quad (3.5)$$

As mentioned above, these are the least restrictive assumptions in the literature.

In order to analyze the convergence of (3.1), we make two standing assumptions on $J(\cdot)$, namely:

(S1) $J(\cdot)$ is \mathcal{C}^1 , and $\nabla J(\cdot)$ is globally Lipschitz-continuous with constant L .

(S2) $J(\cdot)$ is bounded below, and the infimum is attained. Thus

$$J^* := \inf_{\theta \in \mathbb{R}^d} J(\theta)$$

is well-defined, and $J^* > -\infty$. Moreover, the set

$$S_J := \{\theta : J(\theta) = J^*\}$$

is nonempty. Note that hereafter we take $J^* = 0$.

⁵Since $t^{-\gamma}$ is undefined when $t = 0$, we really mean $(t+1)^{-\gamma}$. The same applies elsewhere also.

Aside from these standing assumptions, we introduce four other conditions. Note that not all conditions are assumed in every theorem

(GG) There exists a constant $H < \infty$ such that $\|\nabla J(\theta)\|_2^2 \leq HJ(\theta)$, $\forall \theta \in \mathbb{R}^d$.

(PL) There exists a constant K such that $\|\nabla J(\theta)\|_2^2 \geq KJ(\theta)$, $\forall \theta \in \mathbb{R}^d$.

(KL') There exists a function $\psi(\cdot)$ of Class \mathcal{B} such that $\|\nabla J(\theta)\|_2 \geq \psi(J(\theta))$, $\forall \theta \in \mathbb{R}^d$.

(NSC) There exists a function $\eta(\cdot)$ of Class \mathcal{B} such that $\rho(\theta) \leq \eta(J(\theta))$, $\forall \theta \in \mathbb{R}^d$, where

$$\rho(\theta) := \inf_{\phi \in S_J} \|\theta - \phi\|_2.$$

In the above (GG) stands for ‘‘Gradient Growth.’’ It is satisfied with $H = 2L$ whenever $J(\cdot)$ is convex, but can also hold otherwise. Condition (PL) stands for ‘‘Polyak-Lojasiewicz,’’ while (KL') stands for ‘‘modified Kurdyka-Lojasiewicz.’’ Finally, (NSC) stands for ‘‘Near Strong Convexity.’’ A good discussion of (PL) and (KL) (as opposed to (KL')) can be found in [19], while [18, Section 6] explains the difference between (KL) and (KL'), as well Condition (NSC).

With this background, we first state a theorem on the convergence of the SGD, but without any conclusions as to the rate of convergence. It is [18, Theorem 3].

Theorem 3.1. *Suppose the objective function $J(\cdot)$ satisfies the standing assumptions (S1) and (S2) together with (GG), and that the stochastic gradient \mathbf{h}_{t+1} satisfies (3.4) and (3.5). With these assumptions, we have the following conclusions;*

(1) Suppose

$$\sum_{t=0}^{\infty} \alpha_t^2 < \infty, \sum_{t=0}^{\infty} \alpha_t \mu_t < \infty, \sum_{t=0}^{\infty} \alpha_t^2 M_t^2 < \infty. \quad (3.6)$$

Then $\{\nabla J(\theta_t)\}$ and $\{J(\theta_t)\}$ are bounded, and in addition, $J(\theta_t)$ converges to some random variable as $t \rightarrow \infty$.

(2) If in addition $J(\cdot)$ satisfies (KL'), and

$$\sum_{t=0}^{\infty} \alpha_t = \infty, \quad (3.7)$$

then $J(\theta) \rightarrow 0$ and $\nabla J(\theta_t) \rightarrow \mathbf{0}$ as $t \rightarrow \infty$.

(3) Suppose that in addition to (KL'), $J(\cdot)$ also satisfies (NSC), and that (3.6) and (3.7) both hold. Then $\rho(\theta_t) \rightarrow 0$ as $t \rightarrow \infty$.

Theorem 3.1 does not say anything about the *rate* of convergence. By strengthening the hypothesis from (PL) to (KL'), we can derive explicit bounds on the rate. It is [17, Theorem 4].

Theorem 3.2. *Let various symbols be as in Theorem 3.1. Suppose $J(\cdot)$ satisfies the standing assumptions (S1) through (S3), and also property (PL), and that (3.6) and (3.7) hold. Further, suppose there exist constants $\gamma > 0$ and $\delta \geq 0$ such that⁶*

$$\mu_t = O(t^{-\gamma}), M_t = O(t^{\delta}),$$

where we take $\gamma = 1$ if $\mu_t = 0$ for all sufficiently large t , and $\delta = 0$ if M_t is bounded. Choose the step-size sequence $\{\alpha_t\}$ as $O(t^{-(1-\phi)})$ and $\Omega(t^{-(1-C)})$ where ϕ and C are chosen to satisfy

$$0 < \phi < \min\{0.5 - \delta, \gamma\}, C \in (0, \phi].$$

⁶Since $t^{-\gamma}$ is undefined when $t = 0$, we really mean $(t+1)^{-\gamma}$. The same applies elsewhere also.

Define

$$\mathbf{v} := \min\{1 - 2(\phi + \delta), \gamma - \phi\}. \quad (3.8)$$

Then $\|\nabla J(\theta_t)\|_2^2 = o(t^{-\lambda})$ and $J(\theta_t) = o(t^{-\lambda})$ for every $\lambda \in (0, \mathbf{v})$. In particular, by choosing ϕ very small, it follows that $\|\nabla J(\theta_t)\|_2^2 = o(t^{-\lambda})$ and $J(\theta_t) = o(t^{-\lambda})$ whenever

$$\lambda < \min\{1 - 2\delta, \gamma\}. \quad (3.9)$$

Corollary 3.3. *Suppose all hypotheses of Theorem 3.2 hold. In particular, if $\mu_t = 0$ for all large enough t in (3.4), and M_t in (3.5) is bounded with respect to t , then $\|\nabla J(\theta_t)\|_2^2 = o(t^{-\lambda})$ and $J(\theta_t) = o(t^{-\lambda})$ for all $\lambda < 1$.*

It is worthwhile to compare the content of Corollary 3.3 with the bounds from [1]. In that paper, it is assumed that $\mathbf{z}_t := E_t(\mathbf{h}_{t+1}) = \nabla J(\theta_t)$, and that $CV_t(\mathbf{h}_{t+1}) \leq M^2$ for some finite constant M ; see [1, Eq. (2)]. In the present notation, this is the same as saying that $\mu_t = 0$ for all t , and that $M_t = M$ for all t . Thus the assumption is that the stochastic gradient \mathbf{h}_{t+1} is unbiased and has conditional variance that is uniformly bounded with respect to t and θ_t . With these assumptions on the stochastic gradient, it is shown in [1] that for an arbitrary convex objective function, the best achievable rate $\|\nabla J(\theta_t)\|_2 = O(t^{-1/2})$, or equivalently, $\|\nabla J(\theta_t)\|_2^2 = O(t^{-1})$. Thus the bounds in Corollary 3.3 are tight for any class of functions satisfying the hypotheses therein, which includes both convex as well as a class of nonconvex functions.

4. BLOCK ASYNCHRONOUS STOCHASTIC APPROXIMATION

Until now, we have reviewed some results from a companion paper [18]. This section and the next contain original results due to the authors that are not reported anywhere else. Suppose one wishes to solve (1.2), that is, to find a fixed point of a given map $\mathbf{g}(\cdot)$. As mentioned earlier, when every component of θ_t is updated at each t , this is the standard version of SA, referred to by us as “synchronous” SA, though the term is not very standard. When *exactly one* component of θ_t is updated at each t , this is known as “Asynchronous” SA, a term first introduced in [33]. In this section, we study the solution of (1.2) using “Block Asynchronous” SA, whereby, At each step t , *some but not necessarily all* components of θ_t are updated. This is denoted by the acronym BASA. Clearly both Synchronous SA and Asynchronous SA are special cases of BASA.

4.1. Intermittent Updating: Convergence and Rates. The key distinguishing feature of BASA is that each component of θ_t gets updated in an “intermittent” fashion. Before tackling the convergence of BASA in \mathbb{R}^d , in the present subsection we state and prove results analogous to Theorems 2.5 and 2.8 for the scalar case with intermittent updating.

The problem setup is as follows: The recurrence relationship is

$$w_{t+1} = (1 - \alpha_t \kappa_t) w_t + \alpha_t \kappa_t \xi_{t+1}, \quad (4.1)$$

where $\{w_t\}$ is an \mathbb{R} -valued stochastic process of interest, $\{\xi_t\}$ is the measurement error (or “noise”), $\{\alpha_t\}$ is a $(0, 1)$ -valued stochastic process called the “step size” process, and $\{\kappa_t\}$ is a $\{0, 1\}$ -valued stochastic process called the “update” process. Clearly, if $\kappa_t = 0$, then $w_{t+1} = w_t$, irrespective of the value of α_t ; therefore w_{t+1} is updated only at those t for which $\kappa_t = 1$. This

is the rationale for the name. With the update process $\{\kappa_t\}$, as before we associate a “counter” process $\{v_t\}$, defined by

$$v_t = \sum_{s=0}^t \kappa_s. \quad (4.2)$$

Thus v_t is the number of times up to and including time t at which w_t is updated. We also define

$$v^{-1}(\tau) := \min\{t \in \mathbb{N} : v_t = \tau\}, \quad \forall \tau \geq 1. \quad (4.3)$$

Then $v^{-1}(\cdot)$ is well-defined, and

$$v(v^{-1}(\tau)) = \tau, v^{-1}(v_t) \leq t, v^{-1}(\tau) \leq \tau - 1. \quad (4.4)$$

The last inequality arises from the fact that there are $t + 1$ terms in (4.2). Also, $\kappa_t = 1$ only when $t = v^{-1}(\tau)$ for some τ , and is zero for other values of t . Hence, in (4.1), if $t = v^{-1}(\tau)$ for some τ , then w_t gets updated to w_{t+1} , and

$$w_{t+1} = w_{t+2} = \cdots = w_{v^{-1}(\tau+1)}, \quad (4.5)$$

at which time w gets updated again. Thus w_t is a “piecewise-constant” process, remaining constant between updates. This suggests that we can transform the independent variable from t to τ . Define

$$x_\tau := w_{v^{-1}(\tau)}, \zeta_{\tau+1} := \xi_{v^{-1}(\tau)+1}, \quad \forall \tau \geq 1, \quad (4.6)$$

with the convention that $x_1 = w_0$. Note that the convention is consistent whether $v_0 = 1$ or not (as can be easily verified). Also we define $b_\tau := \alpha_t \kappa_t$, whenever $t = v^{-1}(\tau)$ for some τ . With these definitions, (4.1) is equivalent to

$$x_{\tau+1} = (1 - b_\tau)x_\tau + b_\tau \zeta_{\tau+1}, \quad \forall \tau \geq 1, \quad (4.7)$$

Note that, in (4.7), b_τ is a random variable for all $\tau \geq 1$, and that there is no b_0 . To analyze the behavior of (4.7), we introduce some preliminary concepts. Let \mathcal{F}_t be the σ -algebra generated by w_0, κ_0^t, ξ_1^t . With the change in time indices, define $\{\mathcal{G}_\tau\}$, where $\mathcal{G}_\tau = \mathcal{F}_{v^{-1}(\tau)}$, whenever $t = v^{-1}(\tau)$ for some τ . Then it is easy to see that $\{\mathcal{G}_\tau\}$ is also a filtration, and that

$$E(x_\tau | \mathcal{G}_\tau) = E_t(w_t | \mathcal{F}_t)$$

whenever $t = v^{-1}(\tau)$ for some τ . Hence we can mimic the earlier notation and denote $E(X | \mathcal{G}_\tau)$ by $E_\tau(X)$. Also, if it is assumed that original step size α_t belongs to $\mathcal{M}(\mathcal{F}_t)$, then $b_\tau \in \mathcal{M}(\mathcal{F}_t) = \mathcal{M}(\mathcal{F}_{v^{-1}(\tau)}) = \mathcal{M}(\mathcal{G}_\tau)$. The assumption implies that, while the step α_t may be random, it only makes use of the information available up to and including step t .

Now we present a general convergence result for (4.7). Observe that $\{w_t\}$ is a “piecewise-constant version” of $\{x_\tau\}$. Hence if some conclusions are established for the x -process, they are also established for the w -process, after adjusting for the time change from t to τ .

Theorem 4.1. *Consider the recursion (4.7). Suppose there exist constants μ_t, M_t such that*

$$|E_t(\xi_{t+1})| \leq \mu_t(1 + |w_t|) \quad \forall t \geq 0, \quad (4.8)$$

$$CV_t(\xi_{t+1}) \leq M_t^2(1 + w_t^2), \quad \forall t \geq 0. \quad (4.9)$$

Define

$$f_\tau = b_\tau^2(1 + 2\mu_{v^{-1}(\tau)}^2 + M_{v^{-1}(\tau)}^2) + 3b_\tau\mu_{v^{-1}(\tau)}, \quad (4.10)$$

$$g_\tau = b_\tau^2(2\mu_{v^{-1}(\tau)}^2 + M_{v^{-1}(\tau)}^2) + b_\tau\mu_{v^{-1}(\tau)}. \quad (4.11)$$

Then we have the following conclusions:

(1) If

$$\sum_{\tau=1}^{\infty} f_{\tau} < \infty, \sum_{\tau=1}^{\infty} g_{\tau} < \infty, \quad (4.12)$$

then x_{τ} is bounded almost surely.

(2) If, in addition to (4.12), we also have

$$\sum_{\tau=1}^{\infty} b_{\tau} = \infty, \quad (4.13)$$

then $x_{\tau} \rightarrow 0$ as $\tau \rightarrow \infty$.

(3) If both (4.12) and (4.13) hold, then $x_{\tau} = o(\tau^{-\lambda})$ for every $\lambda < 1$ such that

$$\sum_{\tau=1}^{\infty} (\tau+1)^{\lambda} g_{\tau} < \infty, \quad (4.14)$$

$$\sum_{\tau=1}^{\infty} [b_{\tau} - \lambda \tau^{-1}] = \infty, \quad (4.15)$$

and in addition, there exists a $T < \infty$ such that

$$b_{\tau} - \lambda \tau^{-1} \geq 0 \quad \forall \tau \geq T. \quad (4.16)$$

Proof. The proof consists of reformulating the bounds on the error ξ_{t+1} in such a way that Theorems 2.5 and 2.7 apply. By assumption, we have that

$$|E_t(\xi_{t+1})| \leq \mu_t(1 + |w_t|), \quad \forall t.$$

In particular, when $t = v^{-1}(\tau)$, we have that $\zeta_{\tau+1} = \xi_{t+1}$, and

$$|E_{\tau}(\zeta_{\tau+1})| = |E_t(\xi_{t+1})| \leq \mu_t(1 + |w_t|) = \mu_{v^{-1}(\tau)}(1 + |x_{\tau}|).$$

It follows in an entirely analogous manner that $CV_{\tau}(\zeta_{\tau+1}) \leq M_{v^{-1}(\tau)}(1 + x_{\tau}^2)$. With these observations, we see that Theorems 2.5 and 2.7 apply to (4.7), with the only changes being that (i) the stochastic process is scalar-valued and not vector-valued, (ii) the time index is denoted by τ and not t , and (iii) μ_t, M_t are replaced by $\mu_{v^{-1}(\tau)}, M_{v^{-1}(\tau)}$ respectively. Now the conclusions of the theorem follow from Theorems 2.5 and 2.7. \square

Now, for the convenience of the reader, we reprise the two commonly used approaches for choosing the step size, known as a “global clock” and a “local clock” respectively. This distinction was apparently first introduced in [5]. In each case, there is a *deterministic* sequence $\{\beta_t\}_{t \geq 0}$ of step sizes. If a global clock is used, then $\alpha_t = \beta_t$ at each update, so that $b_{\tau} = \beta_{v^{-1}(\tau)}$. If a local clock is used, then $\alpha_t = \beta_{v_t}$, so that then $b_{\tau} = \beta_{\tau-1}$. The extra -1 in the subscript is to ensure consistency in notation. To illustrate, suppose $\kappa_t = 1$ for all t . Then $v_t = t + 1$, and $v^{-1}(\tau) = \tau - 1$.

Now we begin our analysis of (4.7) with the two types of clocks. Now that Theorem 4.1 is established, the challenge is to determine when (4.13) through (4.16) (as appropriate) hold for the two choices of step sizes, namely global vs. local clocks.

Towards this end, we introduce a few assumptions regarding the update process.

(U1) $v_t \rightarrow \infty$ as $t \rightarrow \infty$ almost surely.

(U2) There exists a random variable r such that

$$\frac{v_t}{t} \rightarrow r \text{ as } t \rightarrow \infty, \text{ a.s..} \quad (4.17)$$

Observe that both assumptions are sample-pathwise. Thus (U2) implies (U1).

We begin by stating the convergence results when a local clock is used.

Theorem 4.2. *Suppose a local clock is used, so that $\alpha_t = \beta_{v_t}$, so that $b_\tau = \beta_{\tau-1}$. Suppose further that Assumption (U1) holds, and moreover*

- (a) $\{\mu_t\}$ is nonincreasing; that is, $\mu_{t+1} \leq \mu_t$, $\forall t$.
- (b) M_t is uniformly bounded, say by M .

With these assumptions,

(1) If

$$\sum_{t=0}^{\infty} \beta_t^2 < \infty, \sum_{t=0}^{\infty} \beta_t \mu_t < \infty, \quad (4.18)$$

then $\{x_\tau\}$ is bounded almost surely, and $\{w_t\}$ is bounded almost surely.

(2) If, in addition

$$\sum_{t=0}^{\infty} \beta_t = \infty, \quad (4.19)$$

then $x_\tau \rightarrow 0$ as $t \rightarrow \infty$ almost surely, and $w_t \rightarrow 0$ as $t \rightarrow \infty$ almost surely.

- (3) Suppose $\beta_t = O(t^{-(1-\phi)})$, for some $\phi > 0$, and $\beta_t = \Omega(t^{-(1-C)})$ for some $C \in (0, \phi]$. Suppose that $\mu_t = O(t^{-\varepsilon})$ for some $\varepsilon > 0$. Then $x_\tau \rightarrow 0$ as $\tau \rightarrow \infty$, and $w_t \rightarrow 0$ as $t \rightarrow \infty$, for all $\phi < \min\{0.5, \varepsilon\}$. Further, $x_\tau = o(\tau^{-\lambda})$, and $w_t = o((v_t)^{-\lambda})$ for all $\lambda < \varepsilon - \phi$. In particular, if $\mu_t = 0$ for all t , then $x_\tau = o(\tau^{-\lambda})$, and $w_t = o((v_t)^{-\lambda})$ for all $\lambda < 1$.
- (4) If Assumption (U2) holds instead of (U1), then in the previous item, $w_t = o((v_t)^{-\lambda})$ can be replaced by $w_t = o(t^{-\lambda})$.

Proof. The proof consists of showing that, under the stated hypotheses, the appropriate conditions in (4.12) through (4.16) hold.

Recall that $b_\tau = \beta_{\tau-1}$. Also, by Assumption (U1), $v_t \rightarrow \infty$ as $t \rightarrow \infty$, almost surely. Hence $v^{-1}(\tau)$ is well-defined for all $\tau \geq 1$.

Henceforth all arguments are along a particular sample path, and we omit the phrase “almost surely,” and also do not display the argument $\omega \in \Omega$.

We first prove Item 1 of the theorem. Recall the definitions of f_τ and g_τ from (4.10) and (4.11) respectively. Item 1 is established if it is shown that (4.12) holds. For this purpose, note that $\mu_s \leq \mu_t$ if $s > t$, and $M_t \leq M$ for all t . We analyze each of the three terms comprising f_τ . First,

$$\sum_{\tau=1}^{\infty} b_\tau^2 = \sum_{\tau=1}^{\infty} \beta_{\tau-1}^2 = \sum_{t=0}^{\infty} \beta_t^2 < \infty.$$

Next, since $M_t \leq M$ for all t , we have that

$$\sum_{\tau=1}^{\infty} b_\tau^2 M_{v^{-1}(\tau)}^2 \leq M^2 \sum_{\tau=1}^{\infty} b_\tau^2 < \infty.$$

Finally,

$$\sum_{\tau=1}^{\infty} b_\tau \mu_{v^{-1}(\tau)} \leq \sum_{\tau=1}^{\infty} \beta_{\tau-1} \mu_{\tau-1} = \sum_{t=0}^{\infty} \beta_t \mu_t < \infty.$$

Here we use the fact that $v^{-1}(\tau) \geq \tau - 1$, so that $\mu_{v^{-1}(\tau)} \leq \mu_{\tau-1}$. Thus it follows from (4.10) that $\{f_\tau\} \in \ell_1$, which is the first half of (4.12). Next, since $\{b_\tau \mu_{v^{-1}(\tau)}\} \in \ell_1$, so is $\{b_\tau^2 \mu_{v^{-1}(\tau)}^2\}$. Hence it follows from (4.11) that $\{g_\tau\} \in \ell_1$, which is the second half of (4.12). This establishes that $\{x_\tau\}$ is bounded, which in turn implies that $\{w_t\}$ is bounded.

To prove Item 2, note that $\sum_{\tau=1}^{\infty} b_\tau = \sum_{\tau=0}^{\infty} \beta_\tau = \infty$. Hence (4.13) holds, and $x_\tau \rightarrow 0$ as $\tau \rightarrow \infty$, which in turn implies that $w_t \rightarrow 0$ as $t \rightarrow \infty$.

Finally we come to the rates of convergence. Recall that $\mu_t = O(t^{-\varepsilon})$ while M_t is bounded by M . Also, β_t is chosen to be $O(t^{-(1-\phi)})$ and $\Omega(t^{-(1-C)})$. From the above, it is clear that

$$f_\tau = O(\tau^{-2+2\phi}) + O(\tau^{-1+\phi-\varepsilon}).$$

Hence (4.12) holds if

$$-2 + 2\phi < -1 \text{ and } -1 + \phi - \varepsilon < -1, \text{ or } \phi < \min\{0.5, \varepsilon\}.$$

Next, from the definition of g_τ in (4.11), it follows that

$$(v^{-1}(\tau) + 1)^\lambda g_\tau \leq (v^{-1}(\tau + 1))^\lambda g_\tau = O(\tau^{-1+\phi-\varepsilon+\lambda}).$$

Hence (4.14) holds if

$$-1 + \phi - \varepsilon + \lambda < -1 \implies \lambda < \varepsilon - \phi.$$

Combining everything shows that $x_\tau = o(\tau^{-\lambda})$ whenever

$$\phi < \min\{0.5, \varepsilon\}, \lambda < \varepsilon - \phi.$$

If $\mu_t = 0$ for all t , then ε can be chosen to be arbitrarily large. However, the limiting factor is that the argument in Theorem 2.7 holds only for $\lambda \leq 1$. Hence $x_\tau = o(\tau^{-\lambda})$ whenever

$$\phi < 0.5, \lambda < 1.$$

Now suppose Assumption (U2) holds, and fix some $\varepsilon > 0$. Then along almost all sample paths, for sufficiently large T we have that $v_t/t \geq r - \varepsilon$ for all $t \geq T$. Thus, whenever $t \geq T$, we have that

$$v_t \geq rt \implies o((v_t)^{-\lambda}) \leq o((rt)^{-\lambda}) = o(t^{-\lambda}).$$

Thus w_t has the same rate of convergence as x_τ . \square

Since the analysis can commence after a finite number of iterations, it is easy to see that Assumption (a) above can be replaced by the following: $\{\mu_t\}$ is eventually nonincreasing; that is, there exists a $T < \infty$ such that $\mu_{t+1} \leq \mu_t$, $\forall t \geq T$.

Next we state a result when a global clock is used. Theorem 4.3 below is not directly comparable to Theorem 4.2 above. Specifically, in Theorem 4.2, the bias coefficient μ_t is assumed to be non increasing, and the variance bound M_t^2 is assumed to be bounded uniformly with respect to t . However, the step sizes are constrained only by the requirement that various summations are finite. In contrast, in Theorem 4.3, there are no assumptions regarding μ_t and \mathcal{M}_t , but the step size sequence $\{\beta_t\}$ is assumed to be nonincreasing.

Theorem 4.3. *Suppose a global clock is used, so that $\alpha_t = \beta_t$ whenever $t = v^{-1}(\tau)$ for some τ and as a result $b_\tau = \beta_{v^{-1}(\tau)}$. Suppose further that Assumption (U2) holds. Finally, suppose that β_t is nonincreasing, so that $\beta_{t+1} \leq \beta_t$ for all t . $\beta_{t+1} \leq \beta_t$, for all t . Under these assumptions,*

(1) If (4.18) holds, and in addition

$$\sum_{t=0}^{\infty} \beta_t^2 M_t^2 < \infty, \quad (4.20)$$

then $\{w_t\}$ is bounded almost surely.

(2) If, in addition, (4.19) holds, then $w_t \rightarrow 0$ as $t \rightarrow \infty$ almost surely.

(3) Suppose in addition that $\beta_t = O(t^{-(1-\phi)})$, for some $\phi > 0$, and $\beta_t = \Omega(t^{-(1-C)})$ for some $C \in (0, \phi]$. Suppose that $\mu_t = O(t^{-\varepsilon})$ for some $\varepsilon > 0$, and $M_t = O(t^\delta)$ for some $\delta \geq 0$. Then $w_t \rightarrow 0$ as $t \rightarrow \infty$ whenever $\phi < \min\{0.5 - \delta, \varepsilon\}$. Moreover, $w_t = o(t^{-\lambda})$ for all $\lambda < \varepsilon - \phi$. In particular, if $\mu_t = 0$ for all t , then $w_t = o(t^{-\lambda})$ for all $\lambda < 1$.

The proof of Theorem 4.3 makes use of the following auxiliary lemma.

Lemma 4.4. Suppose the update process $\{\kappa_t\}$ satisfies Assumption (U2). Suppose $\{\beta_t\}$ is an \mathbb{R}_+ -valued sequence of deterministic constants such that $\beta_{t+1} \leq \beta_t$ for all t , and in addition, (4.19) holds. Then

$$\sum_{\tau=1}^{\infty} \beta_{v^{-1}(\tau)} = \sum_{t=0}^{\infty} \beta_t \kappa_t = \infty. \quad (4.21)$$

Proof. We begin by showing that there exists an integer M such that, whenever $2^k > M$, we have

$$\frac{1}{2^k} \left(\sum_{t=2^{k+1}}^{2^{k+1}} \kappa_t \right) \geq \frac{r}{2}. \quad (4.22)$$

By assumption, the ratio $v_t/t \rightarrow r$ as $t \rightarrow \infty$, where r could depend on the sample path (though the dependence on ω is not displayed). So we can define $\varepsilon = r/2$, and choose an integer M such that

$$\left| \frac{1}{T} \sum_{t=0}^{T-1} \kappa_t - r \right| = \left| \frac{1}{T} \sum_{t=0}^{T-1} (\kappa_t - r) \right| < \frac{\varepsilon}{3}, \quad \forall T \geq M.$$

Thus, if $2^k > M$, we have that

$$\begin{aligned} \left| \frac{1}{2^k} \sum_{t=2^{k+1}}^{2^{k+1}} (\kappa_t - r) \right| &\leq \left| \frac{1}{2^k} \sum_{t=1}^{2^{k+1}} (\kappa_t - r) \right| + \left| \frac{1}{2^k} \sum_{t=1}^{2^k} (\kappa_t - r) \right| \\ &< \frac{2}{3}\varepsilon + \frac{1}{3}\varepsilon = \varepsilon = \frac{r}{2}. \end{aligned}$$

Next, suppose that $\beta_{t+1} \leq \beta_t$ for all t . (If this holds only for all *sufficiently large* t , we just start all the summations from the time when the above holds.)

$$\begin{aligned} \sum_{t=0}^{\infty} \beta_t \kappa_t &\geq \sum_{k=1}^{\infty} \left(\sum_{t=2^{k+1}}^{2^{k+1}} \beta_t \kappa_t \right) \geq \sum_{k=1}^{\infty} \left(\sum_{t=2^{k+1}}^{2^{k+1}} \beta_{2^{k+1}} \kappa_t \right) \\ &= \sum_{k=1}^{\infty} \beta_{2^{k+1}} \left(\sum_{t=2^{k+1}}^{2^{k+1}} \kappa_t \right) \geq \sum_{k=1}^{\infty} \beta_{2^{k+1}} 2^k \frac{r}{2} = \frac{r}{4} \sum_{k=1}^{\infty} \beta_{2^{k+1}} 2^{k+1} \\ &= \frac{r}{4} \sum_{k=1}^{\infty} \sum_{t=2^{k+1}+1}^{2^{k+2}} \beta_{2^{k+1}} \geq \frac{r}{4} \sum_{k=1}^{\infty} \sum_{t=2^{k+1}+1}^{2^{k+2}} \beta_t = \frac{r}{4} \sum_{k=5}^{\infty} \beta_t = \infty. \end{aligned}$$

This is the desired conclusion. \square

Proof. Of Theorem 4.3: Recall that a global clock is used, so that $b_\tau = \beta_{v^{-1}(\tau)}$. Hence

$$\begin{aligned} \sum_{\tau=1}^{\infty} f_\tau &= \sum_{\tau=1}^{\infty} [\beta_{v^{-1}(\tau)}^2 + \beta_{v^{-1}(\tau)}^2 M_{v^{-1}(\tau)}^2 + \beta_{v^{-1}(\tau)} \mu_{v^{-1}(\tau)}] \\ &= \sum_{t=0}^{\infty} [\beta_t^2 + \beta_t M_t^2 + \beta_t \mu_t] < \infty \end{aligned}$$

Via entirely similar reasoning, it follows that $\{g_\tau\} \in \ell_1$. Hence (4.12) holds, and Item 1 follows.

To prove Item 2, it is necessary to establish (4.13), which in this case becomes

$$\sum_{\tau=1}^{\infty} \beta_{v^{-1}(\tau)} = \sum_{\tau=0}^{\infty} b_\tau = \infty.$$

This is (4.13). Hence Item 2 follows.

Finally we come to the rates of convergence. The only difference is that now $M_t = O(t^\delta)$ whereas it was bounded in Theorem 4.2. To avoid tedious repetition, we indicate only the changed steps. The only change is that now

$$f_\tau = O(\tau^{-2+2\phi}) + O(\tau^{-2+2\phi+2\delta}) + O(\tau^{-1+\phi-\varepsilon}).$$

Hence (4.12) holds if

$$-2+2\phi < -1, -2+2\phi+2\delta < -1, \text{ and } -1+\phi-\varepsilon < -1,$$

or

$$\phi < \min\{0.5 - \delta, \varepsilon\}.$$

Next, from the definition of g_τ in (4.11), it follows that

$$(v^{-1}(\tau) + 1)^\lambda g_\tau \leq (v^{-1}(\tau + 1))^\lambda g_\tau = O(\tau^{-1+\phi-\varepsilon+\lambda}).$$

Hence (4.14) holds if

$$-1 + \phi - \varepsilon + \lambda < -1 \implies \lambda < \varepsilon - \phi.$$

Hence $x_\tau = o(\tau^{-\lambda})$ and $w_t = o(t^{-\lambda})$ whenever

$$\phi < \min\{0.5 - \delta, \varepsilon\}, \lambda < \varepsilon - \phi.$$

If $\mu_t = 0$ for all t , then we can choose ε to be arbitrarily large, and we are left with

$$\phi < 0.5 - \delta, \lambda < 1.$$

□

4.2. Boundedness of Iterations. Next, we give a precise statement of the class of fixed point problems to be studied. In this subsection, it is shown that the iterations are bounded (almost surely), while in the next subsection, the convergence of the iterations is established, together with the rate of convergence. The boundedness of the iterations is established under far more general conditions than the convergence. More details are given at the appropriate place.

Let \mathbb{N} denote the set of natural numbers including zero, and let $\mathbf{h} : \mathbb{N} \times (\mathbb{R}^d)^\mathbb{N} \rightarrow (\mathbb{R}^d)^\mathbb{N}$ denote a **measurement function**. Thus \mathbf{h} maps \mathbb{R}^d -valued sequences into \mathbb{R}^d -valued sequences. The objective is to determine a fixed point of this map when only noisy measurements of \mathbf{h} are available at each time t . Specifically, define

$$\eta_t = \mathbf{h}(t, \theta_0^t). \tag{4.23}$$

Suppose that, at time $t + 1$, the learner has access to a vector $\eta_t + \xi_{t+1}$, where ξ_{t+1} denotes the measurement error. The objective is to determine a sequence $\pi^* \in (\mathbb{R}^d)^\mathbb{N}$ (if it exists) such that

$$\mathbf{h}(\pi^*) = \pi^*,$$

using only the noise-corrupted measurements of η_t .

To facilitate this, a few assumptions are made regarding the map \mathbf{h} . First, the map \mathbf{h} is assumed to be **nonanticipative**⁷ and to have **finite memory**. The nonanticipativeness of \mathbf{h} means that

$$\theta_0^\infty, \phi_0^\infty \in (\mathbb{R}^d)^\mathbb{N}, \theta_0^t = \phi_0^t \implies \mathbf{h}(\tau, \theta_0^\infty) = \mathbf{h}(\tau, \phi_0^\infty), 0 \leq \tau \leq t. \quad (4.24)$$

In other words, $\mathbf{h}(t, \theta_0^\infty)$ depends only on θ_0^t . The finite memory of \mathbf{h} means that there exists a finite constant Δ which does not depend on t , such that $\mathbf{h}(t, \theta_0^t)$ further depends only on $\theta_{t-\Delta+1}^t$. With slightly sloppy notation, this can be written as

$$\mathbf{h}(t, \theta_0^t) = \mathbf{h}(t, \theta_{t-\Delta+1}^t), \forall t \geq \Delta, \forall \theta_0^\infty \in (\mathbb{R}^d)^\mathbb{N}. \quad (4.25)$$

This formulation incorporates the possibility of “delayed information” of the form

$$\eta_{t,i} = g_i(\theta_1(t - \Delta_1(t)), \dots, \theta_d(t - \Delta_d(t))), \quad (4.26)$$

where $\Delta_1(t), \dots, \Delta_d(t)$ are delays that could depend on t . The only requirement is that each $\Delta_j(t) \leq \Delta$ for some finite Δ . This formulation is analogous to [33, Eq. (2)] and [5, Eq. (1.4)], which is slightly more general in that they require only that $t - \Delta_i(t) \rightarrow \infty$ as $t \rightarrow \infty$, for each index $i \in [d]$. In particular, if \mathbf{h} is “memoryless” in the sense that, for some function $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, we have

$$\mathbf{h}(t, \theta_0^t) = \mathbf{g}(\theta_t), \quad (4.27)$$

then we can take $\Delta = 1$. Note that, if \mathbf{h} is of the form (4.27), then the problem at hand becomes one of finding a fixed point in \mathbb{R}^d of the map \mathbf{g} , gives noisy measurements of \mathbf{g} at each time step.

To proceed further, it is assumed that the measurement function satisfies the following assumption:

(F1) There exist an integer $\Delta \geq 1$ and a constant $\gamma \in (0, 1)$ such that

$$\|\mathbf{h}(t, \psi_{t-\Delta+1}^t) - \mathbf{h}(t, \phi_{t-\Delta+1}^t)\|_\infty \leq \gamma \|\psi_{t-\Delta+1}^t - \phi_{t-\Delta+1}^t\|_\infty, \forall t \geq \Delta, \forall \psi_0^\infty, \phi_0^\infty \in (\mathbb{R}^d)^\mathbb{N}. \quad (4.28)$$

This assumption means that the map $\theta_{t-\Delta+1}^t \mapsto \mathbf{h}(t, \theta_{t-\Delta+1}^t)$ is a contraction with respect to $\|\cdot\|_\infty$. In case $\Delta = 1$ and \mathbf{h} is of the form (4.27), Assumption (F1) says that the map \mathbf{g} is a contraction.

Now we discuss a few implications of Assumption (F1).

(F2) By repeatedly applying (4.28) over blocks of width Δ , one can conclude that

$$\|\mathbf{h}(t, \psi_{t-\Delta+1}^t) - \mathbf{h}(t, \phi_{t-\Delta+1}^t)\|_\infty \leq \gamma^{\lfloor t/\Delta \rfloor} \|\psi_0^{\Delta-1} - \phi_0^{\Delta-1}\|_\infty, \forall \psi_0^\infty, \phi_0^\infty \in (\mathbb{R}^d)^\mathbb{N}. \quad (4.29)$$

Therefore, for every sequence ϕ_0^∞ , the iterations $\mathbf{h}(t, \phi_0^t)$ converge to a unique fixed point π^* . In particular, if we let $(\pi^*)_0^\infty$ denote the sequence whose value is π^* for every t , then it follows that

$$\|\mathbf{h}(t, (\pi^*)_0^t) - \pi^*\|_\infty \leq C_0 \gamma^{\lfloor t/\Delta \rfloor}, \forall t, \quad (4.30)$$

for some constant C_0 .

⁷In control and system theory, such a function is also referred to as “causal.”

(F3) The following also follows from Assumption (F1): There exist constants $\rho < 1$ and $c'_1 > 0$ such that

$$\|\mathbf{h}(t, \phi_0^t)\|_\infty \leq \rho \max\{c'_1, \|\phi_0^t\|_\infty\}, \forall \phi \in (\mathbb{R}^d)^\mathbb{N}, t \geq 0. \quad (4.31)$$

In order to determine π^* in (F2), we use BASA. Specifically, we choose θ_0 as we wish (either deterministically or at random). At time t , we update θ_t to θ_{t+1} according to

$$\theta_{t+1} = \theta_t + \alpha_t \circ [\eta_t + \xi_{t+1}], \quad (4.32)$$

where α_t is the *vector of step sizes* belonging to $[0, 1)^d$, ξ_{t+1} is the measurement noise vector belonging to \mathbb{R}^d , and \circ denotes the Hadamard product. We are interested in studying two questions:

(Q1) Under what conditions is the sequence of iterations $\{\theta_t\}$ bounded almost surely?

(Q2) Under what conditions does the sequence of iterations $\{\theta_t\}$ converge to π^* as $t \rightarrow \infty$?

Question (Q1) is addressed in this subsection, whereas Question (Q2) is addressed in the next.

In order to study the above two questions, we make some assumptions about various entities in (4.32). Let \mathcal{F}_t denote the σ -algebra generated by the random variables θ_0 , ξ_1^t , and $\alpha_{0,i}^{t,i}$ for $i \in [d]$. Then it is clear that $\{\mathcal{F}_t\}$ is a filtration. As before, we denote $E(X|\mathcal{F}_t)$ by $E_t(X)$.

The first set of assumptions is on the noise.

(N1) There exists a finite constant c'_1 and a sequence of constants $\{\mu_t\}$ such that

$$\|E_t(\xi_{t+1})\|_2 \leq c'_1 \mu_t (1 + \|\theta_0^t\|_\infty), \forall t \geq 0. \quad (4.33)$$

(N2) There exists a finite constant c'_2 and a sequence of constants $\{M_t\}$ such that

$$CV_t(\xi_{t+1}) \leq c'_2 M_t^2 (1 + \|\theta_0^t\|_\infty^2), \forall t \geq 0, \quad (4.34)$$

where, as before,

$$CV_t(\xi_{t+1}) = E_t(\|\xi_{t+1} - E_t(\xi_{t+1})\|_2^2)$$

Before proceeding further, let us compare the conditions (4.33) and (4.34) with their counterparts (2.15) and (2.16) in Theorem 2.5. It can be seen that the above two requirements are more liberal (i.e., less restrictive) than in Theorem 2.5, because the quantity $\|\theta_t\|_2$ is replaced by $\|\theta_0^t\|_\infty$. Hence, in (4.33) and (4.34), the bounds are more loose. However, Theorems 4.5 and 4.10 in the next subsection apply only to *contractive* mappings. Hence Theorems 4.5 and 4.10 complement Theorem 2.5, and do not subsume it.

The next set of assumptions is on the step size sequence.

(S1) The random step size sequences $\{\alpha_{t,i}\}$ and the sequences $\{\mu_t\}$, $\{M_t^2\}$ and satisfy (almost surely)

$$\sum_{t=0}^{\infty} \alpha_{t,i}^2 < \infty, \sum_{t=0}^{\infty} M_t^2 \alpha_{t,i}^2 < \infty, \sum_{t=0}^{\infty} \mu_t \alpha_{t,i} < \infty, \forall i \in [d]. \quad (4.35)$$

(S2) The random step size sequence $\{\alpha_{t,i}\}$ satisfies (almost surely)

$$\sum_{t=0}^{\infty} \alpha_{t,i} = \infty, \text{ a.s., } \forall i \in [d]. \quad (4.36)$$

With these assumptions in place, we state the main result of this subsection, namely, the almost sure boundedness of the iterations. In the next subsection, we state and prove the convergence of the iterations, under more restrictive assumptions.

Theorem 4.5. *Suppose that Assumptions (N1) and (N2) about the noise sequence, (S1) and (S2) about the step size sequence, and (F1) about the function \mathbf{h} hold, and that θ_{t+1} is defined via (4.32). Then $\sup_t \|\theta_t\|_\infty < \infty$ almost surely.*

The proof of the theorem is fairly long and involves several preliminary results and observations.

To aid in proving the results, we introduce a sequence of “renormalizing constants.” This is similar to the technique used in [33]. For $t \geq 0$, define

$$\Lambda_t := \max\{\|\theta_0^t\|_\infty, c'_1\}, \quad (4.37)$$

where c'_1 is defined in (4.23). With this definition, it follows from (4.31) that $\eta_t = \mathbf{h}(t, \theta_0^t)$ satisfies

$$\|\eta_t\|_\infty \leq \rho \Lambda_t, \quad \forall t. \quad (4.38)$$

Define $\zeta_{t+1} = \Lambda_t^{-1} \xi_{t+1}$ for all $t \geq 0$. Now observe that $\Lambda_t^{-1} \leq c_1^{-1}$, and $\Lambda_t^{-1} \leq (\|\theta_0^t\|_\infty)^{-1}$. Hence

$$\|E_t(\zeta_{t+1,i})\|_\infty \leq c'_1 \mu_t (c_1^{-1} + 1) =: c_2 \mu_t, \quad (4.39)$$

where $c_2 = c'_1 (c_1^{-1} + 1)$. In particular, the above implies that

$$|E_t(\zeta_{t+1,i})| \leq c_2 \mu_t, \quad \forall t \geq 0. \quad (4.40)$$

Similarly

$$CV_t(\zeta_{t+1,i}) \leq c_3 M_t^2, \quad \forall t \geq 0, \quad (4.41)$$

for some constant c_3 .

If we compare (4.39) with (4.33), and (4.40) with (4.34), we see that the bounds for the “modified” error ζ_{t+1} are simpler than those for ξ_{t+1} . Specifically, the right side of both (4.39) and (4.40) are bounded with respect to θ_0^t for each t , though they may be unbounded as functions of t . In contrast, the right sides of (4.33) and (4.34) are permitted to be functions of $\|\theta_0^t\|_\infty$.

Though the next result is quite obvious, we state it separately, because it is used repeatedly in the sequel.

Lemma 4.6. *For $i \in [d]$ and $0 \leq s \leq k < \infty$, define the doubly-indexed stochastic process*

$$D_i(s, k+1) = \sum_{t=s}^k \left[\prod_{r=t+1}^k (1 - \alpha_{r,i}) \right] \alpha_{t,i} \zeta_{t+1,i}, \quad (4.42)$$

where an empty product is taken as 1. Then $\{D_i(s, k)\}$ satisfies the recursion

$$D_i(s, k+1) = (1 - \alpha_{k,i}) D_i(s, k) + \alpha_{k,i} \zeta_{k+1,i}, \quad D_i(s, s) = 0. \quad (4.43)$$

In the other direction, (4.42) gives a closed-form solution for the recursion (4.43).

Recall that \mathbb{N} denotes the set of non-negative integers $\{0, 1, 2, \dots\}$. The next lemma is basically the same as [33, Lemma 2].

Lemma 4.7. *There exists $\Omega_1 \subset \Omega$ with $P(\Omega_1) = 1$ and $r_1^* : \Omega_1 \times (0, 1) \rightarrow \mathbb{N}$ such that*

$$|D_i(s, k+1)(\omega)| \leq \varepsilon, \quad \forall k \geq s \geq r_1^*(\omega, \varepsilon). \quad (4.44)$$

Proof. Let $\varepsilon > 0$ be given. It follows from Lemma 4.6 that D_i satisfies the recursion

$$D_i(0, t+1) = (1 - \alpha_{t,i})D_i(0, t) + \alpha_{t,i}\zeta_{t+1,i}$$

with $D_i(0, 0) = 0$. Let us fix an index $i \in [d]$, and invoke (4.40) and (4.41). Then it follows from (4.41) that

$$CV_t(\zeta_{t+1,i}) \leq c_3 M_t^2,$$

and (4.40) also holds. Now, if Assumptions (S1) and (S2) also hold, then all the hypotheses needed to apply Theorem 2.5 are in place. Therefore $D_i(0, k+1)$ converges to zero almost surely. This holds for each $i \in [d]$. Therefore, if we define

$$\Omega_1 = \{\omega \in \Omega : D_i(0, k+1)(\omega) \rightarrow 0 \text{ as } t \rightarrow \infty \forall i \in [d]\},$$

then $P(\Omega_1) = 1$. We can see that for $\omega \in \Omega_1$ we can choose $r_1^*(\omega, \varepsilon)$ such that $\forall k \geq r_1^*(\omega, \varepsilon), i \in [d]$ we have

$$|D_i(0, k+1)(\omega)| \leq \frac{1}{2}\varepsilon.$$

To proceed further, we suppress the argument ω in the interests of clarity. Observe from (4.42) that, whenever $s \leq k$ we have

$$D_i(s, k+1) = \sum_{t=s}^k \left[\prod_{r=t+1}^k (1 - \alpha_{r,i}) \right] \alpha_{t,i} \zeta_{t+1,i} \quad (4.45)$$

$$= \sum_{t=0}^k \left[\prod_{r=t+1}^k (1 - \alpha_{r,i}) \right] \alpha_{t,i} \zeta_{t+1,i} - \sum_{t=0}^{s-1} \left[\prod_{r=t+1}^k (1 - \alpha_{r,i}) \right] \alpha_{t,i} \zeta_{t+1,i} \quad (4.46)$$

$$= D_i(0, k+1) - \left[\prod_{r=s}^k (1 - \alpha_{r,i}) \right] \sum_{t=0}^{s-1} \left[\prod_{r=t+1}^{s-1} (1 - \alpha_{r,i}) \right] \alpha_{t,i} \zeta_{t+1,i} \quad (4.47)$$

$$= D_i(0, k+1) - \left[\prod_{r=s}^k (1 - \alpha_{r,i}) \right] D_i(0, s). \quad (4.48)$$

Since $1 - \alpha_{r,i} \in (0, 1)$ for all r, i , it follows that the product also belongs to $(0, 1)$. Therefore

$$|D_i(s, k+1)| \leq |D_i(0, k+1)| + |D_i(0, s)| \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

This is the desired conclusion. \square

Lemma 4.8. *There exists $\Omega_2 \subset \Omega$ with $P(\Omega_2) = 1$ and $r_2^* : \Omega_1 \times \mathbb{N} \times (0, 1) \rightarrow \mathbb{N}$ such that*

$$\prod_{s=j}^k (1 - \alpha_{s,i}(\omega)) \leq \varepsilon, \forall k \geq r_2^*(\omega, j, \varepsilon), i \in [d], \omega \in \Omega_2. \quad (4.49)$$

Proof. In view of the assumption (S2), if we define

$$\Omega_2 = \left\{ \omega \in \Omega : \sum_{s=j}^{\infty} \alpha_{s,i}(\omega) = \infty \forall i \in [d] \right\},$$

then $P(\Omega_2) = 1$. For all $\omega \in \Omega_2$, we have

$$\sum_{s=j}^{\infty} \alpha_{s,i}(\omega) = \infty.$$

Using the elementary inequality $(1 - x) \leq \exp\{-x\}$ for all $x \in [0, \infty)$, it follows that

$$\prod_{s=j}^k (1 - \alpha_{t,i}(\omega)) \leq \exp \left\{ - \sum_{s=j}^k \alpha_{t,i}(\omega) \right\}.$$

Hence for $\omega \in \Omega_2$, $\prod_{s=j}^k (1 - \alpha_{t,i}(\omega))$ converges to zero as $k \rightarrow \infty$. Thus we can choose $r_2^*(\omega, j, \varepsilon)$ with the required property. \square

In the rest of this section, we will fix $\omega \in \Omega_1 \cap \Omega_2$, the functions r_1^*, r_2^* obtained in Lemma 4.7 and Lemma 4.8 respectively and prove that if (F1) holds, then $\|\theta_t(\omega)\|_\infty$ is bounded, which proves Theorem 4.1.

Let us rewrite the updating rule (4.32) as

$$\theta_{t+1,i} = (1 - \alpha_{t,i})\theta_{t,i} + \alpha_{t,i}(\eta_{t,i} + \Lambda_t \zeta_{t+1,i}), i \in [d], t \geq 0, \quad (4.50)$$

By recursively invoking (4.50) for $k \in [0, t]$, we get

$$\theta_{t+1,i} = A_{t+1,i} + B_{t+1,i} + C_{t+1,i} \quad (4.51)$$

where

$$A_{t+1,i} = \left[\prod_{k=0}^t (1 - \alpha_{k,i}) \right] \theta_{0,i}, \quad (4.52)$$

$$B_{t+1,i} = \sum_{k=0}^t \left[\prod_{r=k+1}^t (1 - \alpha_{r,i}) \right] \alpha_{k,i} \eta_{k,i}, \quad (4.53)$$

$$C_{t+1,i} = \sum_{k=0}^t \left[\prod_{r=k+1}^t (1 - \alpha_{r,i}) \right] \alpha_{k,i} \Lambda_k \zeta_{k+1,i}. \quad (4.54)$$

Lemma 4.9. For $i \in [d]$,

$$|C_{t+1,i}| \leq \Lambda_t \sup_{0 \leq r \leq t} |D_i(r, t+1)|. \quad (4.55)$$

Proof. We begin by establishing an alternate expression for $C_{k,i}$, namely

$$C_{t+1,i} = \Lambda_0 D_i(0, t+1) + \sum_{k=1}^t (\Lambda_k - \Lambda_{k-1}) D_i(k, t+1), \quad (4.56)$$

where $D_i(\cdot, \cdot)$ is defined in (4.42). For this purpose, observe from Lemma 4.6 that $C_{t+1,i}$ satisfies

$$C_{t+1,i} = \Lambda_t \alpha_{t,i} \zeta_{t+1,i} + (1 - \alpha_{t,i}) C_{t,i} = \Lambda_t D_i(t, t+1) + (1 - \alpha_{t,i}) C_{t,i}, \quad (4.57)$$

because $\alpha_{t,i} \zeta_{t+1,i} = D_i(t, t+1)$ due to (4.43) with $s = t$. The proof of (4.56) is by induction. It is evident from (4.54) that

$$C_{1,i} = \Lambda_0 \alpha_{0,1} \zeta_{1,i} = \Lambda_0 D_i(0, 1).$$

Thus (4.56) holds when $t = 0$. Now suppose by way of induction that

$$C_{t,i} = \Lambda_0 D_i(0, t) + \sum_{k=1}^{t-1} (\Lambda_k - \Lambda_{k-1}) D_i(k, t). \quad (4.58)$$

Using this assumption, and the recursion (4.57), we establish (4.56).

Substituting from (4.58) into (4.57) gives

$$C_{t+1,i} = \Lambda_t D_i(t, t+1) + \Lambda_0 (1 - \alpha_{t,i}) D_i(0, t) + (1 - \alpha_{t,i}) \sum_{k=1}^{t-1} (\Lambda_k - \Lambda_{k-1}) D_i(k, t). \quad (4.59)$$

Now (4.42) implies that

$$(1 - \alpha_{t,i})D_i(k, t) = D_i(k, t+1) - \alpha_{t,i}\zeta_{t+1,i} = D_i(k, t+1) - D_i(t, t+1).$$

Therefore the summation in (4.59) becomes

$$\begin{aligned} \sum_{k=1}^{t-1} (\Lambda_k - \Lambda_{k-1})(1 - \alpha_{t,i})D_i(k, t) &= \sum_{k=1}^{t-1} (\Lambda_k - \Lambda_{k-1})D_i(k, t) \\ &\quad - D_i(t, t+1) \sum_{k=1}^{t-1} (\Lambda_k - \Lambda_{k-1}) = S_1 + S_2 \text{ say.} \end{aligned}$$

Then S_2 is just a telescoping sum and equals

$$S_2 = -\Lambda_{t-1}D_i(t, t+1) + \Lambda_0D_i(t, t+1).$$

The second term in (4.59) equals

$$\Lambda_0(1 - \alpha_{t,i})D_i(0, t) = \Lambda_0[D_i(0, t+1) - \alpha_{t,i}\zeta_{t+1,i}] = \Lambda_0D_i(0, t+1) - \Lambda_0D_i(t, t+1).$$

Putting everything together and observing that the term $\Lambda_0D_i(t, t+1)$ cancels out gives

$$C_{t+1,i} = \Lambda_0D_i(0, t+1) + (\Lambda_t - \Lambda_{t-1})D_i(t, t+1) + \sum_{k=1}^{t-1} (\Lambda_k - \Lambda_{k-1})D_i(k, t).$$

This is the same as (4.59) with $t+1$ replacing t . This completes the induction step and thus (4.56) holds. Using the fact that $\Lambda_t \geq \Lambda_{t-1}$, the desired bound (4.55) follows readily. \square

Proof. (Of Theorem 4.1) As per the statement of the theorem, we assume that (F1) holds. We need to prove that

$$\sup_{t \geq 0} \Lambda_t < \infty.$$

Define

$$\delta = \min\left\{\frac{1-\rho}{2\rho}, \frac{1}{2}\right\},$$

and observe that, as a consequence, we have that $\rho(1+2\delta) \leq 1$. Choose $r_1^* = r_1^*(\delta)$ as in Lemma 4.7 such that

$$|D_i(s, k+1)| \leq \delta \quad \forall k \geq s \geq r_1^*, \quad \forall i \in [d].$$

It is now shown that

$$\Lambda_t \leq (1+2\delta)\Lambda_{r_1^*} \quad \forall t, \quad \forall i \in [d]. \quad (4.60)$$

By the monotonicity of $\{\Lambda_t\}$, it is already known that $\Lambda_t \leq \Lambda_{r_1^*}$ for $t \leq r_1^*$. Hence, once (4.60) is established, it will follow that

$$\sup_{0 \leq t < \infty} \Lambda_t \leq (1+2\delta)\Lambda_{r_1^*}.$$

The proof of (4.60) is by induction on t . Accordingly, suppose (4.60) holds for $t \leq k$. Using (4.55), we have

$$|C_{k+1,i}| \leq \delta \Lambda_k \leq \Lambda_{r_1^*} \delta (1+2\delta). \quad (4.61)$$

It is easy to see from its definition that

$$|A_{k+1,i}| \leq \Lambda_{r_1^*} \left[\prod_{s=0}^k (1 - \alpha_{s,i}) \right]$$

Using the induction hypothesis that $\Lambda_t \leq (1 + 2\delta)\Lambda_{r_1^*}$ for $t \leq k$, we have

$$\begin{aligned}
|B_{k+1,i}| &\leq \sum_{s=0}^k \left[\prod_{r=s+1}^k (1 - \alpha_{r,i}) \right] \alpha_{s,i} |\eta_{s,i}| \\
&\leq \sum_{s=0}^k \left[\prod_{r=s+1}^k (1 - \alpha_{r,i}) \right] \alpha_{s,i} \rho \Lambda_s \\
&\leq \rho (1 + 2\delta) \Lambda_{r_1^*} \sum_{s=0}^k \left[\prod_{r=s+1}^k (1 - \alpha_{r,i}) \right] \alpha_{s,i} \\
&\leq \Lambda_{r_1^*} \sum_{s=0}^k \left[\prod_{r=s+1}^k (1 - \alpha_{r,i}) \right] \alpha_{s,i},
\end{aligned}$$

because $\rho(1 + 2\delta) \leq 1$. Also, the following identity is easy to prove by induction.

$$\left[\prod_{s=0}^k (1 - \alpha_{s,i}) \right] + \sum_{s=0}^k \left[\prod_{r=s+1}^k (1 - \alpha_{r,i}) \right] \alpha_{s,i} = 1 \quad \forall k < \infty \quad (4.62)$$

Combining these bounds gives

$$|A_{k+1,i}| + |B_{k+1,i}| \leq \Lambda_{r_1^*}.$$

Combining this with (4.51) and (4.61) leads to

$$\theta_{k+1,i} \leq \Lambda_{r_1^*} (1 + \delta(1 + 2\delta)) \leq \Lambda_{r_1^*} (1 + 2\delta).$$

Therefore $\|\theta_{k+1}\|_\infty \leq \Lambda_{r_1^*} (1 + 2\delta)$, and

$$\Lambda_{k+1} = \max\{\|\theta_{k+1}\|_\infty, \Lambda_k\} \leq \Lambda_{r_1^*} (1 + 2\delta).$$

This proves the induction hypothesis and completes the proof of Theorem 4.1. \square

4.3. Convergence of Iterations with Rates. In this subsection, we further study the iteration sequence (4.32), under a variety of Block (or Batch) updating schemes, corresponding to various choices of the step sizes. Whereas the almost sure boundedness of the iterations is established in the previous subsection, in this subsection we prove that the iterations converge to the desired fixed point π^* . Then we also find bounds on the rate of convergence.

We study three specific methods for choosing the step size vector α_t in (4.32). Within the first two methods, we further divide into local clocks and global clocks. However, in the third method, we permit only the use of a global clock, for reasons to be specified.

4.3.1. Convergence Theorem. The overall plan is to follow up Theorem 4.5, which establishes the almost sure boundedness of the iterations, with a stronger result showing that the iterations converge almost surely to π^* , the fixed point of the map \mathbf{h} . This convergence is established under the same assumptions as in Theorem 4.5. In particular, the step size sequence is assumed to satisfy (S1) and (S2). Having done this, we then study conditions under which (S1) and (S2) hold for each of the three methods for choosing the step sizes.

Theorem 4.10. *Suppose that Assumptions (N1) and (N2) about the noise sequence, (S1) and (S2) about the step size sequence, and (F1) about the function \mathbf{h} hold, and that θ_{t+1} is defined via (4.32). Then $\theta_t \rightarrow \pi^*$ as $t \rightarrow \infty$ almost surely, where π^* is defined in (F2).*

Proof. From (4.51), we have an expression for $\theta_{t+1,i}$, where $A_{t+1,i}$, $B_{t+1,i}$ and $C_{t+1,i}$ are given by (4.52), (4.53) and (4.54) respectively. Also, by changing notation from k to t and s to k in (4.62), and multiplying both sides by π_i^* , we can write

$$\pi_i^* = \left[\prod_{k=0}^t (1 - \alpha_{k,i}) \right] \pi_i^* + \left\{ \sum_{k=0}^t \left[\prod_{r=k+1}^t (1 - \alpha_{r,i}) \right] \alpha_{k,i} \right\} \pi_i^*, \forall t.$$

Substituting from these formulas gives

$$\theta_{t+1,i} - \pi_i^* = \bar{A}_{t+1,i} + \bar{B}_{t+1,i} + C_{t+1,i}, \quad (4.63)$$

where

$$\bar{A}_{t+1,i} = \prod_{k=0}^t (1 - \alpha_{k,i}) (\theta_{0,i} - \pi_i^*), \quad (4.64)$$

$$\bar{B}_{t+1,i} = \left[\prod_{r=k+1}^t (1 - \alpha_{r,i}) \right] \alpha_{k,i} (\eta_{k,i} - \pi_i^*), \quad (4.65)$$

and $C_{t+1,i}$ is as in (4.54). It is shown in turn that each of these quantities approaches zero as $t \rightarrow \infty$.

First, from Assumption (S2), it follows that⁸

$$\prod_{k=0}^t (1 - \alpha_{k,i}) \rightarrow 0 \text{ as } t \rightarrow \infty.$$

Since $\theta_{0,i} - \pi_i^*$ is a constant along each sample path, $\bar{A}_{t+1,i}$ approaches zero.

Second, by combining (4.29) and (4.30) in Property (F2), it follows that

$$|\eta_{t,i} - \pi_i^*| \leq \gamma^{\lfloor t/\Delta \rfloor} \|\theta_0^\Delta - (\pi^*)^\Delta\|_\infty \leq C_1 \gamma^{\lfloor t/\Delta \rfloor}$$

for some constant C_1 (which depends on the sample path). Thus

$$\sum_{r=0}^{\infty} |\eta_{r,i} - \pi_i^*| < \infty$$

along almost all sample paths. Now it follows from (4.65) that

$$\begin{aligned} |\bar{B}_{t+1,i}| &\leq \left[\prod_{r=k+1}^t (1 - \alpha_{r,i}) \right] \alpha_{k,i} |\eta_{k,i} - \pi_i^*| \\ &\leq \left[\prod_{r=k+1}^t (1 - \alpha_{r,i}) \right] \alpha_{k,i} C_1 \gamma^{\lfloor t/\Delta \rfloor} =: L_{t+1,i}. \end{aligned} \quad (4.66)$$

Let $L_{t+1,i}$ denote the right side of this inequality. Then it follows from Lemma 4.6 that $L_{t+1,i}$ satisfies the recursion

$$L_{t+1,i} = (1 - \alpha_{t,i}) L_{t,i} + \alpha_{t,i} C_1 \gamma^{\lfloor t/\Delta \rfloor}. \quad (4.67)$$

The convergence of $L_{t+1,i}$ to zero can be proved using Theorem 4.1. Since the quantity $C_1 \gamma^{\lfloor t/\Delta \rfloor}$ is deterministic, its mean is itself and its variance is zero. So in (4.8) and (4.9), we can define

$$\mu_t^L := C_1 \gamma^{\lfloor t/\Delta \rfloor}, M_t^L := 0 \forall t.$$

⁸We omit the phrase “almost surely” in these arguments.

We can substitute these definitions into (4.10) and (4.11), and define

$$f_\tau^L = b_\tau^2(1 + 2\mu_{v^{-1}(\tau)}^2) + 3b_\tau\mu_{v^{-1}(\tau)}, \quad (4.68)$$

$$g_\tau^L = b_\tau^2(2\mu_{v^{-1}(\tau)}^2) + b_\tau\mu_{v^{-1}(\tau)}. \quad (4.69)$$

Since $\alpha_t \in [0, 1]$ and the sequence $\{\mu_t^L\}$ is summable (because $\gamma < 1$), and $M_t^L \equiv 0$, (4.12) is satisfied. Also, by Assumption (S2), (4.13) is satisfied. Hence $L_{t+1,i} \rightarrow 0$ as $t \rightarrow \infty$, which in turn implies that $\bar{B}_{t+1,i} \rightarrow 0$ as $t \rightarrow \infty$.

Finally, we come to $C_{t+1,i}$. It is evident from (4.54) and Lemma 4.6 that $C_{t+1,i}$ satisfies the recursion

$$C_{t+1,i} = (1 - \alpha_{t,i})C_{t,i} + \alpha_{t,i}\Lambda_t\zeta_{t,i}. \quad (4.70)$$

Now observe that Λ_t is bounded, and the rescaled error signal $\zeta_{t+1,i}$ satisfies (4.40) and (4.41). Hence, if Λ^* is a bound for Λ_t , then it follows from (4.40) and (4.41) that

$$|E_t(\Lambda_t\zeta_{t+1,i})| \leq c_2\Lambda^*\mu_t, \forall t \geq 0, CV_t(\Lambda_t\zeta_{t+1,i}) \leq c_3\Lambda^*M_t^2, \forall t \geq 0, \quad (4.71)$$

Hence, when Assumptions (S1) and (S2) hold, it follows from Theorem 4.1 that $C_{t+1,i} \rightarrow 0$ as $t \rightarrow \infty$. \square

4.3.2. Various Types of Updating and Rates of Convergence. Next, we describe three different ways of choosing the update processes $\{\kappa_{t,i}\}$.

Bernoulli Updating: For each $i \in [d]$, choose a rate $b_i \in (0, 1]$, and let $\{\kappa_{t,i}\}$ be a Bernoulli process such that

$$\Pr\{\kappa_{t,i} = 1\} = b_i, \forall t.$$

Moreover, the processes $\{\kappa_{t,i}\}$ and $\{\kappa_{t,j}\}$ are independent whenever $i \neq j$. Let $v_{t,i}$, the counter process for coordinate i , be defined as usual. Then it is easy to see that $v_{t,i}/t \rightarrow b_i$ as $t \rightarrow \infty$, for each $i \in [d]$. Thus Assumption (U2) is satisfied for each $i \in [d]$.

Markovian Updating: Suppose $\{Y_t\}$ is a sample path of an irreducible Markov process on the state space $[d]$. Define the update process $\{\kappa_{t,i}\}$ by

$$\kappa_{t,i} = I_{\{Y_t=i\}} = \begin{cases} 1, & \text{if } Y_t = i, \\ 0, & \text{if } Y_t \neq i. \end{cases}$$

Let μ denote the stationary distribution of the Markov process. Then the ratio $v_{t,i}/t \rightarrow \mu_i$ as $t \rightarrow \infty$, for each $i \in [d]$. Hence once again Assumption (U2) holds.

Batch Markovian Updating: This is an extension of the above. Instead of a single Markovian sample path, there are N different sample paths, denoted by $\{Y_t^n\}$ where $n \in [N]$. Each sample path $\{Y_t^n\}$ comes an irreducible Markov process over the state space $[d]$, and the dynamics of different Markov processes could be different (though there does not seem to be any advantage to doing this). The update process is now given by

$$\kappa_{t,i} = \sum_{n \in [N]} I_{\{Y_t^n=i\}}.$$

Define the counter process $v_{t,i}$ as before, and let μ^n denote the stationary distribution of the n -th Markov process. Then

$$\frac{v_{t,i}}{t} \rightarrow \sum_{n \in [N]} \mu_i^n.$$

Hence once again Assumption (U2) holds.

Now we establish convergence rates under each of the above updating methods (and indeed, any method such that Assumption (U2) is satisfied). The proof of Theorem 4.10 gives us a hint on how this can be done. Specifically, each of the entities $\bar{A}_{t+1,i}, L_{t+1,i}, C_{t+1,i}$ satisfies a stochastic recursion, whose rate of convergence can be established using Theorems 4.2 and 4.3. These theorems apply to *scalar-valued* stochastic processes with intermittent updating. In principle, when updating θ_t , we could use a mixture of global and local clocks for different components. However, in our view, this would be quite unnatural. Instead, it is assumed that for *every* component, either a global clock or a local clock is used. Recall also the bounds (4.33) and (4.34) on the error ξ_{t+1} .

Theorem 4.11. *Suppose a local clock is used, so that $\alpha_{t,i} = \beta_{v_{t,i}}$ for each i that is updated at time t . Suppose that $\{\mu_t\}$ is nonincreasing; that is, $\mu_{t+1} \leq \mu_t$, $\forall t$, and M_t is uniformly bounded, say by M . Suppose in addition that $\beta_t = O(t^{-(1-\phi)})$, for some $\phi > 0$, and $\beta_t = \Omega(t^{-(1-C)})$ for some $C \in (0, \phi]$. Suppose that $\mu_t = O(t^{-\varepsilon})$ for some $\varepsilon > 0$. Then $\theta_\tau \rightarrow 0$ as $\tau \rightarrow \infty$ for all $\phi < \min\{0.5, \varepsilon\}$. Further, $\theta_\tau = o(\tau^{-\lambda})$ for all $\lambda < \varepsilon - \phi$. In particular, if $\mu_t = 0$ for all t , then $\theta_\tau = o(\tau^{-\lambda})$ for all $\lambda < 1$.*

The proof of the rate of convergence uses Item (3) of Theorem 4.1. In the proof, let us ignore the index i wherever possible, because the subsequent analysis applies to each index i . Recall that $\bar{A}_{t+1,i}$ is defined in (4.64). Since $\ln(1-x) \leq -x$ for all $x \in (0, 1)$, it follows that

$$\ln \prod_{k=0}^t (1 - \alpha_{k,i}) \leq - \sum_{k=0}^t \alpha_{k,i},$$

where $\alpha_{k,i} = 0$ unless there is an update at time k . Now, since a local clock is used, we have that $\alpha_{k,i} = \beta_{v_{k,i}}$ whenever there is an update at time k . Therefore

$$\sum_{k=0}^t \alpha_{k,i} = \sum_{s=0}^{v_{t,i}} \beta_s$$

Now, if Assumption (U2) holds (which it does for each of the three types of updating considered), it follows that $v_{t,i} \approx t/r$ for large t . Thus, if $\beta_\tau = \Omega(\tau^{-(1-C)})$, then we can reason as follows:

$$\sum_{s=0}^{v_t} \beta_s \approx \sum_{s=0}^{t/r} s^{-(1-C)} \approx (t/r)^C.$$

Therefore, for large enough t , we have that

$$\prod_{k=0}^t (1 - \alpha_k) \leq \exp(-(t/r)^C).$$

It follows from (4.64) that $\bar{A}_{t+1,i} \rightarrow 0$ *geometrically fast*.

Next we come to $\bar{B}_{t+1,i}$, which is bounded by $L_{t+1,i}$, as defined in (4.67). Recall the definitions (4.68) and (4.69) for the sequences $\{f_\tau^L\}$ and $\{g_\tau^L\}$. Then (4.12) and (4.13) will hold whenever $C > 0$. Since Assumption (U2) holds, we have that

$$\mu_{v^{-1}(\tau)}^L = C_1 \gamma^{\lfloor v^{-1}(\tau)/\Delta \rfloor} \leq C_2 \gamma'^{\tau}$$

for suitable constants C_2 and γ' . The point to note is that the sequence $\{C_2 \gamma'^{\tau}\}$ is a *geometrically convergent* sequence because $\gamma' < 1$. Therefore (4.14) holds for *every* $\lambda > 0$. Also, (4.15) holds

for all $C > 0$. Hence it follows from Item (3) of Theorem 4.1 that $L_{t+1,i} = o(t^{-\lambda})$ for every $\lambda > 0$.

This leaves only $C_{t+1,i}$. We already know that $C_{t+1,i}$ satisfies the recursion (4.70). Moreover, the modified error sequence $\{\Lambda_t \zeta_{t,i}\}$ satisfies (4.71). The estimates for the rate of convergence now follow from Item (3) of Theorem 4.1, and need not be discussed again.

Theorem 4.12. *Suppose a global clock is used, so that $\alpha_{t,i} = \beta_{t,i}$ whenever the i -th component of θ_t is updated. Suppose that β_t is nonincreasing, so that $\beta_{t+1} \leq \beta_t$ for all t . Suppose in addition that $\beta_t = O(t^{-(1-\phi)})$, for some $\phi > 0$, and $\beta_t = \Omega(t^{-(1-C)})$ for some $C \in (0, \phi]$. Suppose that $\mu_t = O(t^{-\varepsilon})$ for some $\varepsilon > 0$, and $M_t = O(t^\delta)$ for some $\delta \geq 0$. Then $\theta_t \rightarrow 0$ as $t \rightarrow \infty$ whenever*

$$\phi < \min\{0.5 - \delta, \varepsilon\}.$$

Moreover, $\theta_t = o(t^{-\lambda})$ for all $\lambda < \varepsilon - \phi$. In particular, if $\mu_t = 0$ for all t , then $\theta_t = o(t^{-\lambda})$ for all $\lambda < 1$.

The proof is omitted as it is very similar to that of Theorem 4.11.

5. CONCLUSIONS AND PROBLEMS FOR FUTURE RESEARCH

In this paper, we reviewed some results on the convergence of the Stochastic Gradient method from [18]. Then we analyzed the convergence of “intermittently updated” processes of the form (4.1). For this formulation, we derived sufficient conditions for convergence, as well as bounds on the *rate* of convergence. Building on this, we derived both sufficient conditions for convergence, and bounds on the *rate* of convergence, for the full BASA formulation of (1.2). Next, we applied these results to derive sufficient conditions for the convergence of a fixed point iteration with noisy measurements.

There are several interesting problems thrown up by the analysis here. To our knowledge, our paper is the first to provide explicit estimates of the rates of convergence for BASA. A related issue is that of “Markovian” stochastic approximation, in which the update process is the sample path of an irreducible Markov process. It would be worthwhile to examine whether the present approach can handle Markovian SA as well.

Acknowledgments

The research of MV was supported by the Science and Engineering Research Board, India. The authors thank the anonymous reviewer for a very thorough reading of the previous version and for many helpful comments.

REFERENCES

- [1] Y. Arjevani, Y. Carmon, J.C. Duchi, D.J. Foster, N. Srebro, B. Woodworth, Lower bounds for non-convex stochastic optimization, *Mathematical Programming*, 199 (2023) 165-214.
- [2] A. Benveniste, M. Métivier, P. Priouret, *Adaptive Algorithms and Stochastic Approximation*, Springer-Verlag, 1990.
- [3] J. Bhandari, D. Russo, R. Singal, A finite time analysis of temporal difference learning with linear function approximation, *Proceedings of Machine Learning Research*, 75 (2018) 1691-1692.
- [4] J.R. Blum, Multivariable stochastic approximation methods, *Annals of Mathematical Statistics*, 25 (1954) 737-744.
- [5] V.S. Borkar, Asynchronous stochastic approximations, *SIAM Journal on Control and Optimization*, 36 (1998) 840-851.

- [6] V.S. Borkar, S.P. Meyn, The O.D.E. method for convergence of stochastic approximation and reinforcement learning, *SIAM Journal on Control and Optimization*, 38 (2000) 447-469.
- [7] V.S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint* (Second Edition), Cambridge University Press, Cambridge, 2022.
- [8] Z. Chen, S.T. Maguluri, S. Shakkottai, K. Shanmugam, A lyapunov theory for finite-sample guarantees of asynchronous q-learning and td-learning variants, *arxiv:2102.01567v3*, 2021.
- [9] D.P. Derevitskii, A.L. Fradkov, Two models for analyzing the dynamics of adaptation algorithms, *Automation and Remote Control*, 35 (1974) 59-67.
- [10] C. Derman, J. Sacks, On Dvoretzky's stochastic approximation theorem, *Annals of Mathematical Statistics*, 30 (1959) 601-606.
- [11] A. Dvoretzky, On stochastic approximation, In: *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, pp. 39-56. University of California Press, 1956.
- [12] E. Even-Dar, Y. Mansour, Learning rates for q-learning, *Journal of Machine Learning Research*, 5 (2003) 1-25.
- [13] B. Franci, S. Grammatico, Convergence of sequences: A survey, *Annual Reviews in Control*, 53 (2022) 1-26.
- [14] E.G. Gladyshev, On stochastic approximation, *Theory of Probability and Its Applications*, 5 (1965) 275-278.
- [15] L. Grüne, C.M. Kellett, ISS-Lyapunov functions for discontinuous discrete-time systems, *IEEE Transactions on Automatic Control*, 59 (2014) 3098-3103.
- [16] S. Ilandarideva, A. Juditsky, G. Lan, T. Li, Accelerated stochastic approximation with state-dependent noise, *arxiv:2307.01497*, 2023.
- [17] R.L. Karandikar, M. Vidyasagar, Convergence of batch asynchronous stochastic approximation with applications to reinforcement learning, *arXiv:2109.03445*, 2024.
- [18] R.L. Karandikar, M. Vidyasagar, Convergence rates for stochastic approximation: Biased noise with unbounded variance, and applications, *arXiv:2312.02828*, 2024.
- [19] H. Karimi, J. Nutini, M. Schmidt, Linear convergence of gradient and proximal-gradient methods under the polyak- lojasiewicz condition, *Lecture Notes in Computer Science*, 9851 (2016) 795-811.
- [20] J. Kiefer, J. Wolfowitz, Stochastic estimation of the maximum of a regression function, *Annals of Mathematical Statistics*, 23 (1952) 462-466.
- [21] H.J. Kushner, General convergence results for stochastic approximations via weak convergence theory, *Journal of Mathematical Analysis and Applications*, 61 (1977) 490-503.
- [22] H.J. Kushner, D.S. Clark, *Stochastic approximation methods for constrained and unconstrained systems*, Springer, 2012.
- [23] H.J. Kushner, G.G. Yin, *Stochastic Approximation Algorithms and Applications* (Second Edition), Springer-Verlag, 2003.
- [24] T.L. Lai, Stochastic approximation, *The Annals of Statistics*, 31 (2003) 391-406.
- [25] J. Liu, Y. Yuan, On almost sure convergence rates of stochastic gradient methods, In: Po-Ling Loh and Maxim Raginsky, (ed.), *Proceedings of Thirty Fifth Conference on Learning Theory, Proceedings of Machine Learning Research*, vol. 178, pp. 2963-2983, 2022.
- [26] L. Ljung, Analysis of recursive stochastic algorithms, *IEEE Transactions on Automatic Control*, 22 (1977) 551-575.
- [27] L. Ljung, Strong convergence of a stochastic approximation algorithm, *Annals of Statistics*, 6 (1978) 680-696.
- [28] G. Qu, A. Wierman, Finite-time analysis of asynchronous stochastic approximation and q-learning, *Proceedings of Machine Learning Research*, 125 (2020) 1-21.
- [29] H. Robbins, D. Siegmund, A convergence theorem for non negative almost supermartingales and some applications, pp. 233-257, Elsevier, 1971.
- [30] H. Robbins, S. Monro, A stochastic approximation method, *Annals of Mathematical Statistics*, 22 (1951) 400-407.
- [31] R. Srikant, L. Ying, Finite-time error bounds for linear stochastic approximation and td learning, *arxiv:1902.00923v3*, 2019.
- [32] V. Tadić, A. Doucet, Asymptotic bias of stochastic gradient search, *The Annals of Applied Probability*, 27 (2017) 3255-3304.

- [33] J.N. Tsitsiklis, Asynchronous stochastic approximation and q-learning, *Machine Learning*, 16 (1994) 185-202.
- [34] M. Vidyasagar, Convergence of stochastic approximation via martingale and converse Lyapunov methods, *Mathematics of Controls Signals and Systems*, 35 (2023) 351-374.
- [35] M.J. Wainwright, Stochastic approximation with cone-contractive operators: Sharp ℓ_∞ -bounds for q-learning, *arXiv:1905.06265*, 2019.
- [36] C.J.C.H. Watkins, P. Dayan, Q-learning, *Machine Learning*, 8 (1992) 279-292.