
Supplementary Material for Multi-Agent Reinforcement Learning via Double Averaging Primal-Dual Optimization

Hoi-To Wai
The Chinese University of Hong Kong
Shatin, Hong Kong
htwai@se.cuhk.edu.hk

Zhuoran Yang
Princeton University
Princeton, NJ, USA
zy6@princeton.edu

Zhaoran Wang
Northwestern University
Evanston, IL, USA
zhaoranwang@gmail.com

Mingyi Hong
University of Minnesota
Minneapolis, MN, USA
mhong@umn.edu

A Proof of Theorem 1

We repeat the statement of the theorem as follows:

Theorem. Under A1 and A2, we denote by $(\theta^*, \{\mathbf{w}_i^*\}_{i=1}^N)$ the primal-dual optimal solution to the optimization problem in (10). Set the step sizes as $\gamma_2 = \beta\gamma_1$ with $\beta := 8(\rho + \lambda_{\max}(\hat{\mathbf{A}}^\top \hat{\mathbf{C}}^{-1} \hat{\mathbf{A}})) / \lambda_{\min}(\hat{\mathbf{C}})$. Define $\bar{\theta}(t) := \frac{1}{N} \sum_{i=1}^N \theta_i^t$ as the average of parameters. If the primal step size γ_1 is sufficiently small, then there exists a constant $0 < \sigma < 1$ that

$$\|\bar{\theta}(t) - \theta^*\|^2 + (1/\beta N) \sum_{i=1}^N \|\mathbf{w}_i^t - \mathbf{w}_i^*\|^2 = \mathcal{O}(\sigma^t), \quad (1/N) \sum_{i=1}^N \|\theta_i^t - \bar{\theta}(t)\| = \mathcal{O}(\sigma^t).$$

If $N, M \gg 1$ and the graph is geometric with $\lambda = 1 - c/N$ for $c > 0$, a sufficient condition for convergence is to set $\gamma = \mathcal{O}(1/\max\{N^2, M^2\})$ and the resultant rate is $\sigma = 1 - \mathcal{O}(1/\max\{MN^2, M^3\})$.

Notation We first define a set of notations pertaining to the proof. For any $\beta > 0$, observe that the primal-dual optimal solution, $(\theta^*, \{\mathbf{w}_i^*\}_{i=1}^N)$, to the optimization problem (10) can be written as

$$\underbrace{\begin{pmatrix} \rho \mathbf{I} & \sqrt{\frac{\beta}{N}} \hat{\mathbf{A}}^\top & \cdots & \sqrt{\frac{\beta}{N}} \hat{\mathbf{A}}^\top \\ -\sqrt{\frac{\beta}{N}} \hat{\mathbf{A}} & \beta \hat{\mathbf{C}} & \cdots & \cdots \\ \vdots & \mathbf{0} & \ddots & \mathbf{0} \\ -\sqrt{\frac{\beta}{N}} \hat{\mathbf{A}} & \cdots & \cdots & \beta \hat{\mathbf{C}} \end{pmatrix}}_{:=\mathbf{G}} \begin{pmatrix} \theta^* \\ \frac{1}{\sqrt{\beta N}} \mathbf{w}_1^* \\ \vdots \\ \frac{1}{\sqrt{\beta N}} \mathbf{w}_N^* \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ -\sqrt{\frac{\beta}{N}} \mathbf{b}_1 \\ \vdots \\ -\sqrt{\frac{\beta}{N}} \mathbf{b}_N \end{pmatrix}, \quad (\text{A.1})$$

where we denote the matrix on the left hand side as \mathbf{G} . This equation can be obtained by checking the first-order optimality condition. In addition, for any $p \in \{1, \dots, M\}$, we define the \mathbf{G}_p as

$$\mathbf{G}_p := \begin{pmatrix} \rho \mathbf{I} & \sqrt{\frac{\beta}{N}} \mathbf{A}_p^\top & \cdots & \sqrt{\frac{\beta}{N}} \mathbf{A}^\top \\ -\sqrt{\frac{\beta}{N}} \mathbf{A}_p & \beta \mathbf{C}_p & \cdots & \cdots \\ \vdots & \mathbf{0} & \ddots & \mathbf{0} \\ -\sqrt{\frac{\beta}{N}} \mathbf{A}_p & \cdots & \cdots & \beta \mathbf{C}_p \end{pmatrix}. \quad (\text{A.2})$$

By definition, \mathbf{G} is the sample average of $\{\mathbf{G}_p\}_{p=1}^M$. Define $\bar{\boldsymbol{\theta}}(t) := (1/N) \sum_{i=1}^N \boldsymbol{\theta}_i^t$ as the average of the local parameters at iteration t . Furthermore, we define

$$\mathbf{h}_{\boldsymbol{\theta}}(t) := \rho \bar{\boldsymbol{\theta}}(t) + \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{A}}^\top \mathbf{w}_i^t, \quad \mathbf{g}_{\boldsymbol{\theta}}(t) := \frac{1}{NM} \sum_{i=1}^N \sum_{p=1}^M (\rho \boldsymbol{\theta}_i^{\tau_p^t} + \mathbf{A}_p^\top \mathbf{w}_i^{\tau_p^t}), \quad (\text{A.3})$$

$$\mathbf{h}_{\mathbf{w}_i}(t) := \hat{\mathbf{A}} \bar{\boldsymbol{\theta}}(t) - \hat{\mathbf{C}} \mathbf{w}_i^t - \hat{\mathbf{b}}_i, \quad \mathbf{g}_{\mathbf{w}_i}(t) := \frac{1}{M} \sum_{p=1}^M (\mathbf{A}_p \boldsymbol{\theta}_i^{\tau_p^t} - \mathbf{C}_p \mathbf{w}_i^{\tau_p^t} - \mathbf{b}_{p,i}), \quad (\text{A.4})$$

where $\mathbf{h}_{\boldsymbol{\theta}}(t)$ and $\mathbf{h}_{\mathbf{w}}(t) := [\mathbf{h}_{\mathbf{w}_1}(t), \dots, \mathbf{h}_{\mathbf{w}_N}(t)]$ represent the gradients evaluated by a *centralized* and *batch* algorithm. Note that $\mathbf{g}_{\boldsymbol{\theta}}(t)$ defined in (A.3) coincides with that in (16). Using Lemma 1, it can be checked that $\bar{\boldsymbol{\theta}}(t+1) = \bar{\boldsymbol{\theta}}(t) - \gamma_1 \mathbf{g}_{\boldsymbol{\theta}}(t)$ and $\mathbf{w}_i^{t+1} = \mathbf{w}_i^t - \gamma_2 \mathbf{g}_{\mathbf{w}_i}(t)$ for all $t \geq 1$. That is, $\bar{\boldsymbol{\theta}}(t+1)$ and \mathbf{w}_i^{t+1} can be viewed as primal-dual updates using $\mathbf{g}_{\boldsymbol{\theta}}(t)$ and $\mathbf{g}_{\mathbf{w}_i}(t)$, which are decentralized counterparts of gradients $\mathbf{h}_{\boldsymbol{\theta}}(t)$ and $\mathbf{h}_{\mathbf{w}_i}(t)$ defined in (A.3) (A.4).

To simplify the notation, hereafter, we define vectors $\underline{\mathbf{h}}(t)$, $\underline{\mathbf{g}}(t)$, and $\underline{\mathbf{v}}(t)$ by

$$\underline{\mathbf{h}}(t) := \begin{pmatrix} \mathbf{h}_{\boldsymbol{\theta}}(t) \\ -\sqrt{\frac{\beta}{N}} \mathbf{h}_{\mathbf{w}_1}(t) \\ \vdots \\ -\sqrt{\frac{\beta}{N}} \mathbf{h}_{\mathbf{w}_N}(t) \end{pmatrix}, \quad \underline{\mathbf{g}}(t) := \begin{pmatrix} \mathbf{g}_{\boldsymbol{\theta}}(t) \\ -\sqrt{\frac{\beta}{N}} \mathbf{g}_{\mathbf{w}_1}(t) \\ \vdots \\ -\sqrt{\frac{\beta}{N}} \mathbf{g}_{\mathbf{w}_N}(t) \end{pmatrix}, \quad \underline{\mathbf{v}}(t) := \begin{pmatrix} \bar{\boldsymbol{\theta}}(t) - \boldsymbol{\theta}^* \\ \frac{1}{\sqrt{\beta N}} (\mathbf{w}_1^t - \mathbf{w}_1^*) \\ \vdots \\ \frac{1}{\sqrt{\beta N}} (\mathbf{w}_N^t - \mathbf{w}_N^*) \end{pmatrix}. \quad (\text{A.5})$$

Using (A.1), it can be verified that (see the detailed derivation in Section A.2)

$$\underline{\mathbf{h}}(t) = \mathbf{G} \underline{\mathbf{v}}(t). \quad (\text{A.6})$$

By adopting the analysis in [15] and under Assumption 2, it can be shown that with

$$\beta := \frac{8(\rho + \lambda_{\max}(\hat{\mathbf{A}}^\top \hat{\mathbf{C}}^{-1} \hat{\mathbf{A}}))}{\lambda_{\min}(\hat{\mathbf{C}})},$$

then \mathbf{G} is full rank with its eigenvalues satisfying

$$\lambda_{\max}(\mathbf{G}) \leq \left| \frac{\lambda_{\max}(\hat{\mathbf{C}})}{\lambda_{\min}(\hat{\mathbf{C}})} \right| \lambda_{\max}(\rho \mathbf{I} + \hat{\mathbf{A}}^\top \hat{\mathbf{C}}^{-1} \hat{\mathbf{A}}), \quad \lambda_{\min}(\mathbf{G}) \geq \frac{8}{9} \lambda_{\min}(\hat{\mathbf{A}}^\top \hat{\mathbf{C}}^{-1} \hat{\mathbf{A}}) > 0. \quad (\text{A.7})$$

Moreover, let $\mathbf{G} := \mathbf{U} \boldsymbol{\Lambda} \mathbf{U}^{-1}$ be the eigen-decomposition of \mathbf{G} , where $\boldsymbol{\Lambda}$ is a diagonal matrix consists of the eigenvalues of \mathbf{G} , and the columns of \mathbf{U} are the eigenvectors. Then, \mathbf{U} is full rank with

$$\|\mathbf{U}\| \leq 8(\rho + \lambda_{\max}(\hat{\mathbf{A}}^\top \hat{\mathbf{C}}^{-1} \hat{\mathbf{A}})) \left| \frac{\lambda_{\max}(\hat{\mathbf{C}})}{\lambda_{\min}(\hat{\mathbf{C}})} \right|, \quad \|\mathbf{U}^{-1}\| \leq \frac{1}{\rho + \lambda_{\max}(\hat{\mathbf{A}}^\top \hat{\mathbf{C}}^{-1} \hat{\mathbf{A}})}. \quad (\text{A.8})$$

Furthermore, we also define the following upper bounds on the spectral norms

$$\mathbf{G} := \|\mathbf{G}\|, \quad \bar{\mathbf{G}} := \max_{p=1, \dots, M} \|\mathbf{G}_p\|, \quad \bar{\mathbf{A}} := \max_{p=1, \dots, M} \|\mathbf{A}_p\|, \quad \bar{\mathbf{C}} := \max_{p=1, \dots, M} \|\mathbf{C}_p\|. \quad (\text{A.9})$$

Lastly, we define the following two Lyapunov functions

$$\mathcal{E}_c(t) := \frac{1}{N} \sqrt{\sum_{i=1}^N \|\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}(t)\|^2}, \quad \mathcal{E}_g(t) := \frac{1}{N} \sqrt{\sum_{i=1}^N \|\mathbf{s}_i^t - \mathbf{g}_{\boldsymbol{\theta}}(t)\|^2}. \quad (\text{A.10})$$

Note that we have the following inequalities:

$$\mathcal{E}_c(t) \leq \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}(t)\|, \quad \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}(t)\| \leq \sqrt{N} \mathcal{E}_c(t), \quad (\text{A.11})$$

which follows from the norm equivalence $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1 \leq \sqrt{N} \|\mathbf{x}\|_2$ for any $\mathbf{x} \in \mathbb{R}^N$.

Convergence Analysis We denote that $\gamma_1 = \gamma$ and $\gamma_2 = \beta\gamma$. To study the linear convergence of the PD-DistIAG method, our first step is to establish a bound on the difference from the primal-dual optimal solution, $\underline{\mathbf{v}}(t)$. Observe with the choice of our step size ratio,

$$\underline{\mathbf{v}}(t+1) = (\mathbf{I} - \gamma\mathbf{G})\underline{\mathbf{v}}(t) + \gamma(\underline{\mathbf{h}}(t) - \underline{\mathbf{g}}(t)). \quad (\text{A.12})$$

Consider the difference vector $\underline{\mathbf{h}}(t) - \underline{\mathbf{g}}(t)$. Its first block can be evaluated as

$$\begin{aligned} [\underline{\mathbf{h}}(t) - \underline{\mathbf{g}}(t)]_1 &= \frac{1}{NM} \sum_{i=1}^N \sum_{p=1}^M \left(\rho(\bar{\boldsymbol{\theta}}(t) - \boldsymbol{\theta}_i^{\tau_p^t}) + \mathbf{A}_p^\top (\mathbf{w}_i^t - \mathbf{w}_i^{\tau_p^t}) \right) \\ &= \frac{1}{NM} \sum_{i=1}^N \sum_{p=1}^M \left(\rho(\bar{\boldsymbol{\theta}}(t) - \bar{\boldsymbol{\theta}}(\tau_p^t)) + \mathbf{A}_p^\top (\mathbf{w}_i^t - \mathbf{w}_i^{\tau_p^t}) \right) + \frac{\rho}{NM} \sum_{i=1}^N \sum_{p=1}^M (\bar{\boldsymbol{\theta}}(\tau_p^t) - \boldsymbol{\theta}_i^{\tau_p^t}). \end{aligned} \quad (\text{A.13})$$

Meanwhile, for any $i \in \{1, \dots, N\}$, the $(i+1)$ -th block is

$$\begin{aligned} [\underline{\mathbf{h}}(t) - \underline{\mathbf{g}}(t)]_{i+1} &= \sqrt{\frac{\beta}{N}} \frac{1}{M} \sum_{p=1}^M \left(\mathbf{A}_p (\boldsymbol{\theta}_i^{\tau_p^t} - \bar{\boldsymbol{\theta}}(t)) + \mathbf{C}_p (\mathbf{w}_i^t - \mathbf{w}_i^{\tau_p^t}) \right) \\ &= \sqrt{\frac{\beta}{N}} \frac{1}{M} \sum_{p=1}^M \left(\mathbf{A}_p (\bar{\boldsymbol{\theta}}(\tau_p^t) - \bar{\boldsymbol{\theta}}(t)) + \mathbf{C}_p (\mathbf{w}_i^t - \mathbf{w}_i^{\tau_p^t}) \right) + \sqrt{\frac{\beta}{N}} \frac{1}{M} \sum_{p=1}^M \mathbf{A}_p (\boldsymbol{\theta}_i^{\tau_p^t} - \bar{\boldsymbol{\theta}}(\tau_p^t)). \end{aligned} \quad (\text{A.14})$$

For ease of presentation, we stack up and denote the residual terms (related to consensus error) in (A.13) and (A.14) as the vector $\underline{\boldsymbol{\varepsilon}}_c(t)$. That is, the first block of $\underline{\boldsymbol{\varepsilon}}_c(t)$ is $\rho/(NM) \cdot \sum_{i=1}^N \sum_{p=1}^M (\bar{\boldsymbol{\theta}}(\tau_p^t) - \boldsymbol{\theta}_i^{\tau_p^t})$, and the remaining blocks are given by $\sqrt{\beta/N} \cdot 1/M \cdot \sum_{p=1}^M \mathbf{A}_p (\boldsymbol{\theta}_i^{\tau_p^t} - \bar{\boldsymbol{\theta}}(\tau_p^t))$, $\forall i \in \{1, \dots, N\}$. Then by the definition of \mathbf{G}_p in (A.2), we obtain the following simplification:

$$\underline{\mathbf{h}}(t) - \underline{\mathbf{g}}(t) - \underline{\boldsymbol{\varepsilon}}_c(t) = \frac{1}{M} \sum_{p=1}^M \mathbf{G}_p \left(\sum_{j=\tau_p^t}^{t-1} \Delta \underline{\mathbf{v}}(j) \right), \quad (\text{A.15})$$

where we have defined

$$\Delta \underline{\mathbf{v}}(j) := \begin{pmatrix} \bar{\boldsymbol{\theta}}(j+1) - \bar{\boldsymbol{\theta}}(j) \\ \frac{1}{\sqrt{\beta N}} (\mathbf{w}_1^{j+1} - \mathbf{w}_1^j) \\ \vdots \\ \frac{1}{\sqrt{\beta N}} (\mathbf{w}_N^{j+1} - \mathbf{w}_N^j) \end{pmatrix}. \quad (\text{A.16})$$

Clearly, we can express $\Delta \underline{\mathbf{v}}(j)$ as $\Delta \underline{\mathbf{v}}(j) = \underline{\mathbf{v}}(j+1) - \underline{\mathbf{v}}(j)$ with $\underline{\mathbf{v}}(t)$ defined in (A.5). Combining (A.6) and (A.12), we can also write $\Delta \underline{\mathbf{v}}(j)$ in (A.16) as

$$\Delta \underline{\mathbf{v}}(j) = \gamma [\underline{\mathbf{h}}(j) - \underline{\mathbf{g}}(j)] - \gamma \underline{\mathbf{h}}(j). \quad (\text{A.17})$$

Denoting $\widehat{\underline{\mathbf{v}}}(t) := \mathbf{U}^{-1} \underline{\mathbf{v}}(t)$, multiplying \mathbf{U}^{-1} on both sides of (A.12) yields

$$\widehat{\underline{\mathbf{v}}}(t+1) = (\mathbf{I} - \gamma\mathbf{\Lambda})\widehat{\underline{\mathbf{v}}}(t) + \gamma \mathbf{U}^{-1} (\underline{\mathbf{h}}(t) - \underline{\mathbf{g}}(t)). \quad (\text{A.18})$$

Combining (A.15), (A.17), and (A.18), by triangle inequality, we have

$$\begin{aligned} \|\widehat{\underline{\mathbf{v}}}(t+1)\| &\leq \\ \|\mathbf{I} - \gamma\mathbf{\Lambda}\| \|\widehat{\underline{\mathbf{v}}}(t)\| + \gamma \|\mathbf{U}^{-1}\| &\left\{ \|\underline{\boldsymbol{\varepsilon}}_c(t)\| + \frac{\gamma \bar{\mathbf{G}}}{M} \sum_{p=1}^M \sum_{j=\tau_p^t}^{t-1} [\|\underline{\mathbf{h}}(j)\| + \|\underline{\mathbf{h}}(j) - \underline{\mathbf{g}}(j)\|] \right\}, \end{aligned} \quad (\text{A.19})$$

where \bar{G} appears in (A.9) and $\underline{\mathcal{E}}_c(t)$ is the residue term of the consensus. Furthermore, simplifying the right-hand side of (A.19) yields

$$\begin{aligned} \|\widehat{\mathbf{v}}(t+1)\| &\leq \|\mathbf{I} - \gamma\mathbf{\Lambda}\| \|\widehat{\mathbf{v}}(t)\| + \gamma\|\mathbf{U}^{-1}\| \left\{ \|\underline{\mathcal{E}}_c(t)\| + \gamma\bar{G} \sum_{j=(t-M)_+}^{t-1} [\|\mathbf{h}(j)\| + \|\mathbf{h}(j) - \mathbf{g}(j)\|] \right\} \\ &\leq \|\mathbf{I} - \gamma\mathbf{\Lambda}\| \|\widehat{\mathbf{v}}(t)\| + \gamma\|\mathbf{U}^{-1}\| \left(\|\underline{\mathcal{E}}_c(t)\| \right. \\ &\quad \left. + \gamma\bar{G} \sum_{j=(t-M)_+}^{t-1} \left\{ \|\underline{\mathcal{E}}_c(j)\| + G\|\mathbf{U}\| \|\widehat{\mathbf{v}}(j)\| + \bar{G}\|\mathbf{U}\| \sum_{\ell=(j-M)_+}^{j-1} [\|\widehat{\mathbf{v}}(\ell+1)\| + \|\widehat{\mathbf{v}}(\ell)\|] \right\} \right). \end{aligned} \quad (\text{A.20})$$

Moreover, using the definition and (A.11), we can upper bound $\|\underline{\mathcal{E}}_c(t)\|$ by

$$\|\underline{\mathcal{E}}_c(t)\| \leq \frac{1}{M} \sum_{p=1}^M \left[(\rho + \bar{A}\sqrt{\beta N}) \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}(\tau_p^t)\| \right] \leq \sqrt{N}(\rho + \bar{A}\sqrt{\beta N}) \max_{(t-M)_+ \leq q \leq t} \mathcal{E}_c(q). \quad (\text{A.21})$$

Thus, combining (A.20) and (A.21), we bound $\|\widehat{\mathbf{v}}(t+1)\|$ by

$$\|\widehat{\mathbf{v}}(t+1)\| \leq \|\mathbf{I} - \gamma\mathbf{\Lambda}\| \|\widehat{\mathbf{v}}(t)\| + C_1(\gamma) \max_{(t-2M)_+ \leq q \leq t-1} \|\widehat{\mathbf{v}}(q)\| + C_2(\gamma) \max_{(t-2M)_+ \leq q \leq t} \mathcal{E}_c(q), \quad (\text{A.22})$$

where constants $C_1(\gamma)$ and $C_2(\gamma)$ are given by

$$C_1(\gamma) := \gamma^2 \|\mathbf{U}\| \|\mathbf{U}^{-1}\| \bar{G}M(G + 2\bar{G}M), \quad C_2(\gamma) := \gamma\|\mathbf{U}^{-1}\| (1 + \gamma\bar{G}M) \sqrt{N}(\rho + \bar{A}\sqrt{\beta N}).$$

Notice that since \mathbf{U}^{-1} is full rank, the squared norm $\|\widehat{\mathbf{v}}(t)\|^2$ is proportional to $\|\bar{\boldsymbol{\theta}}(t) - \boldsymbol{\theta}^*\|^2 + (1/\beta N) \sum_{i=1}^N \|\mathbf{w}_i^* - \mathbf{w}_i^t\|^2$, i.e., the optimality gap at the t -th iteration.

We next upper bound $\mathcal{E}_c(t+1)$ as defined in (A.10). Notice that $N\mathcal{E}_c(t+1)$ can be written as Frobenius norm of the matrix $\boldsymbol{\Theta}^{t+1} - \mathbf{1}\bar{\boldsymbol{\theta}}(t+1)^\top$, where $\boldsymbol{\Theta}^{t+1} = ((\boldsymbol{\theta}_1^{t+1})^\top; \dots; (\boldsymbol{\theta}_N^{t+1})^\top)$. Also, we denote $\mathbf{S}^t = ((s_1^t)^\top; \dots; (s_N^t)^\top)$. By the update in (15) and using the triangle inequality, we have

$$\begin{aligned} \mathcal{E}_c(t+1) &= \frac{1}{N} \|\boldsymbol{\Theta}^{t+1} - \mathbf{1}\bar{\boldsymbol{\theta}}(t+1)^\top\|_F = \frac{1}{N} \|\mathbf{W}(\boldsymbol{\Theta}^t - \mathbf{1}\bar{\boldsymbol{\theta}}(t)^\top) - \gamma(\mathbf{S}^t - \mathbf{1}\mathbf{g}_\theta(t)^\top)\|_F \\ &\leq \frac{1}{N} (\|\boldsymbol{\Theta}^{t+1} - \mathbf{1}\bar{\boldsymbol{\theta}}(t)^\top\|_F + \gamma\|\mathbf{S}^t - \mathbf{1}\mathbf{g}_\theta(t)^\top\|_F). \end{aligned} \quad (\text{A.23})$$

Notice that we have $\lambda := \lambda_{\max}(\mathbf{W} - (1/N)\mathbf{1}\mathbf{1}^\top) < 1$ as the graph is connected. Using the fact that $N\mathcal{E}_g(t) = \|\mathbf{S}^t - \mathbf{1}\mathbf{g}_\theta(t)^\top\|_F$, the right-hand side of (A.23) can be bounded by

$$\mathcal{E}_c(t+1) \leq \lambda \mathcal{E}_c(t) + \gamma \mathcal{E}_g(t), \quad (\text{A.24})$$

where the Lyapunov function $\mathcal{E}_g(t)$ is defined in (A.10).

To conclude the proof, we need to further upper bound $\mathcal{E}_g(t+1)$. To simplify the notation, let us define $\mathbf{G}_p^t = (\nabla_{\boldsymbol{\theta}} J_{1,p}(\boldsymbol{\theta}_1^t; \mathbf{w}_1^t)^\top; \dots; \nabla_{\boldsymbol{\theta}} J_{N,p}(\boldsymbol{\theta}_N^t; \mathbf{w}_N^t)^\top)$ and observe that

$$\mathcal{E}_g(t+1) = \frac{1}{N} \left\| \mathbf{S}^{t+1} - \mathbf{1}\mathbf{g}_\theta(t+1)^\top \right\|_F = \frac{1}{N} \left\| \mathbf{W}\mathbf{S}^t + \frac{1}{M} (\mathbf{G}_{p_{t+1}}^{t+1} - \mathbf{G}_{p_{t+1}}^{\tau_{p_{t+1}}^t}) - \mathbf{1}\mathbf{g}_\theta(t+1)^\top \right\|_F \quad (\text{A.25})$$

where we have used (13). Furthermore, we observe

$$\begin{aligned} \mathcal{E}_g(t+1) &= \frac{1}{N} \left\| \mathbf{W}(\mathbf{S}^t - \mathbf{1}\mathbf{g}_\theta(t)^\top) + \frac{1}{M} (\mathbf{G}_{p_{t+1}}^{t+1} - \mathbf{G}_{p_{t+1}}^{\tau_{p_{t+1}}^t}) - \mathbf{1}(\mathbf{g}_\theta(t+1) - \mathbf{g}_\theta(t))^\top \right\|_F \\ &\leq \frac{1}{N} \left(\|\mathbf{W}(\mathbf{S}^t - \mathbf{1}\mathbf{g}_\theta(t)^\top)\|_F + \left\| \frac{1}{M} (\mathbf{G}_{p_{t+1}}^{t+1} - \mathbf{G}_{p_{t+1}}^{\tau_{p_{t+1}}^t}) - \mathbf{1}(\mathbf{g}_\theta(t+1) - \mathbf{g}_\theta(t))^\top \right\|_F \right) \\ &\leq \lambda \mathcal{E}_g(t) + \frac{1}{N} \left\| \frac{1}{M} (\mathbf{G}_{p_{t+1}}^{t+1} - \mathbf{G}_{p_{t+1}}^{\tau_{p_{t+1}}^t}) - \mathbf{1}(\mathbf{g}_\theta(t+1) - \mathbf{g}_\theta(t))^\top \right\|_F \end{aligned} \quad (\text{A.26})$$

We observe $\mathbf{G}_p^t = ((\mathbf{w}_1^t)^\top \mathbf{A}_p; \dots; (\mathbf{w}_N^t)^\top \mathbf{A}_p) + \rho \boldsymbol{\Theta}^t$ and $\mathbf{g}_\theta(t+1) - \mathbf{g}_\theta(t) = M^{-1}(\rho \bar{\boldsymbol{\theta}}(t+1) - \rho \bar{\boldsymbol{\theta}}(\tau_{p_{t+1}}^t) + N^{-1} \mathbf{A}_{p_{t+1}}^\top \sum_{i=1}^N (\mathbf{w}_i^{t+1} - \mathbf{w}_i^{\tau_{p_{t+1}}^t}))$. Adopting the notations $\boldsymbol{\Omega}^t = ((\mathbf{w}_1^t)^\top; \dots; (\mathbf{w}_N^t)^\top)$ and $\bar{\mathbf{w}}^t = N^{-1} \sum_{i=1}^N \mathbf{w}_i^t$, we observe that

$$\begin{aligned} & M^{-1}(\mathbf{G}_{p_{t+1}}^{t+1} - \mathbf{G}_{p_{t+1}}^{\tau_{p_{t+1}}^t}) - \mathbf{1}(\mathbf{g}_\theta(t+1) - \mathbf{g}_\theta(t))^\top \\ &= \frac{\rho}{M}(\boldsymbol{\Theta}^{t+1} - \mathbf{1}\bar{\boldsymbol{\theta}}(t+1)^\top - \boldsymbol{\Theta}^{\tau_{p_{t+1}}^t} - \mathbf{1}\bar{\boldsymbol{\theta}}(\tau_{p_{t+1}}^t)^\top) \\ & \quad + \frac{1}{M}(\boldsymbol{\Omega}^{t+1} - \mathbf{1}(\bar{\mathbf{w}}^{t+1})^\top - \boldsymbol{\Omega}^{\tau_{p_{t+1}}^t} + \mathbf{1}(\bar{\mathbf{w}}^{\tau_{p_{t+1}}^t})^\top) \mathbf{A}_{p_{t+1}}. \end{aligned} \quad (\text{A.27})$$

Using the triangular inequality, the norm of the above can be bounded as

$$\begin{aligned} & \frac{\rho}{M} \left(\|\boldsymbol{\Theta}^{t+1} - \mathbf{1}\bar{\boldsymbol{\theta}}(t+1)^\top\|_F + \|\boldsymbol{\Theta}^{\tau_{p_{t+1}}^t} - \mathbf{1}\bar{\boldsymbol{\theta}}(\tau_{p_{t+1}}^t)^\top\|_F \right) \\ & \quad + \frac{\|\mathbf{A}_{p_{t+1}}\|}{M} \left(\|\boldsymbol{\Omega}^{t+1} - \mathbf{1}(\bar{\mathbf{w}}^{t+1})^\top - \boldsymbol{\Omega}^{\tau_{p_{t+1}}^t} + \mathbf{1}(\bar{\mathbf{w}}^{\tau_{p_{t+1}}^t})^\top\|_F \right) \end{aligned} \quad (\text{A.28})$$

From the norm equivalence $\|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_1$, we recognize that $\|\boldsymbol{\Omega}^{t+1} - \mathbf{1}(\bar{\mathbf{w}}^{t+1})^\top - \boldsymbol{\Omega}^{\tau_{p_{t+1}}^t} + \mathbf{1}(\bar{\mathbf{w}}^{\tau_{p_{t+1}}^t})^\top\|_F \leq \sum_{i=1}^N \|\mathbf{w}_i^{t+1} - \bar{\mathbf{w}}^{t+1} - \mathbf{w}_i^{\tau_{p_{t+1}}^t} + \bar{\mathbf{w}}^{\tau_{p_{t+1}}^t}\|$. It holds for all $t' \leq t$ that

$$\mathbf{w}_i^{t+1} - \mathbf{w}_i^{t'} = -\frac{\gamma}{\beta M} \sum_{\ell=t'}^t \sum_{p=1}^M \left[\mathbf{A}_p(\boldsymbol{\theta}_i^{\tau_p^\ell} - \boldsymbol{\theta}^*) - \mathbf{C}_p(\mathbf{w}_i^{\tau_p^\ell} - \mathbf{w}_i^*) \right].$$

We thus obtain

$$\begin{aligned} & \frac{\|\mathbf{A}_{p_{t+1}}\|}{M} \left(\|\boldsymbol{\Omega}^{t+1} - \mathbf{1}(\bar{\mathbf{w}}^{t+1})^\top - \boldsymbol{\Omega}^{\tau_{p_{t+1}}^t} + \mathbf{1}(\bar{\mathbf{w}}^{\tau_{p_{t+1}}^t})^\top\|_F \right) \\ & \leq \frac{\|\mathbf{A}_{p_{t+1}}\|}{M} \sum_{i=1}^N \left\| \mathbf{w}_i^{t+1} - \frac{1}{N} \sum_{j=1}^N \mathbf{w}_j^{t+1} - \mathbf{w}_i^{\tau_{p_{t+1}}^t} + \frac{1}{N} \sum_{j=1}^N \mathbf{w}_j^{\tau_{p_{t+1}}^t} \right\| \\ & \leq \frac{2\gamma\bar{A}}{\beta M^2} \sum_{i=1}^N \sum_{\ell=(t-M)_+}^t \sum_{p=1}^M \left(\|\mathbf{A}_p(\boldsymbol{\theta}_i^{\tau_p^\ell} - \boldsymbol{\theta}^*) - \mathbf{C}_p(\mathbf{w}_i^{\tau_p^\ell} - \mathbf{w}_i^*)\| \right) \\ & \leq \frac{2\gamma\bar{A}}{\beta M} \sum_{i=1}^N \sum_{\ell=(t-M)_+}^t \left[\max_{(\ell-M)_+ \leq q \leq \ell} \left(\bar{A} \|\boldsymbol{\theta}_i^q - \boldsymbol{\theta}^*\| + \bar{C} \|\mathbf{w}_i^q - \mathbf{w}_i^*\| \right) \right]. \end{aligned} \quad (\text{A.29})$$

Thus, combining (A.24), (A.29), and the definition of \mathcal{E}_c in (A.10), we have

$$\begin{aligned} & \frac{1}{N} \left\| \frac{1}{M} (\mathbf{G}_{p_{t+1}}^{t+1} - \mathbf{G}_{p_{t+1}}^{\tau_{p_{t+1}}^t}) - \mathbf{1}(\mathbf{g}_\theta(t+1) - \mathbf{g}_\theta(t))^\top \right\|_F \\ & \leq \frac{\rho}{M} [\mathcal{E}_c(\tau_{p_{t+1}}^t) + \mathcal{E}_c(t+1)] \\ & \quad + \frac{2\gamma\bar{A}(M+1)}{\beta N M} \sum_{i=1}^N \max_{(t-2M)_+ \leq q \leq t} \left(\bar{A} \|\boldsymbol{\theta}_i^q - \boldsymbol{\theta}^*\| + \bar{C} \|\mathbf{w}_i^q - \mathbf{w}_i^*\| \right) \\ & \leq \frac{\rho}{M} [\mathcal{E}_c(\tau_{p_{t+1}}^t) + \lambda \mathcal{E}_c(t) + \gamma \mathcal{E}_g(t)] \\ & \quad + \frac{2\gamma\bar{A}(M+1)}{\beta M} \max_{(t-2M)_+ \leq q \leq t} \left(\bar{A} \sqrt{N} \mathcal{E}_c(q) + \bar{A} \|\bar{\boldsymbol{\theta}}(q) - \boldsymbol{\theta}^*\| + \frac{\bar{C}}{N} \sum_{i=1}^N \|\mathbf{w}_i^q - \mathbf{w}_i^*\| \right). \end{aligned} \quad (\text{A.30})$$

Combining (A.26) and (A.30), we obtain that

$$\begin{aligned} \mathcal{E}_g(t+1) & \leq \left(\lambda + \frac{\gamma\rho}{M} \right) \mathcal{E}_g(t) + \left[\frac{2\gamma\bar{A}^2(M+1)\sqrt{N}}{\beta M} + \frac{2(1+\lambda)}{M} \right] \max_{(t-2M)_+ \leq q \leq t} \mathcal{E}_c(q) \\ & \quad + \frac{2\gamma\bar{A}(M+1)}{\beta M} \max_{(t-2M)_+ \leq q \leq t} \left(\bar{A} \|\bar{\boldsymbol{\theta}}(q) - \boldsymbol{\theta}^*\| + \frac{\bar{C}}{N} \sum_{i=1}^N \|\mathbf{w}_i^q - \mathbf{w}_i^*\| \right). \end{aligned} \quad (\text{A.31})$$

To bound the last term on the right-hand side of (A.31), For all q , we observe that:

$$\begin{aligned} & \left(\bar{A} \|\bar{\boldsymbol{\theta}}(q) - \boldsymbol{\theta}^*\| + \frac{\bar{C}}{N} \sum_{i=1}^N \|\mathbf{w}_i^q - \mathbf{w}_i^*\| \right)^2 \\ & \leq (N+1)(\bar{A})^2 \left[\|\bar{\boldsymbol{\theta}}(q) - \boldsymbol{\theta}^*\|^2 + \beta \left(\frac{\bar{C}}{\bar{A}} \right)^2 \frac{1}{\beta N} \sum_{i=1}^N \|\mathbf{w}_i^q - \mathbf{w}_i^*\|^2 \right] \\ & \leq (N+1) \|\mathbf{U}\| \max \left\{ (\bar{A})^2, \beta(\bar{C})^2 \right\} \|\mathbf{v}(q)\|^2, \end{aligned}$$

which further implies that

$$\begin{aligned} \mathcal{E}_g(t+1) & \leq \left(\lambda + \frac{\gamma\rho}{M} \right) \mathcal{E}_g(t) + \left(\frac{2\gamma\bar{A}^2(M+1)\sqrt{N}}{\beta M} + \frac{2(1+\lambda)}{M} \right) \max_{(t-2M)_+ \leq q \leq t} \mathcal{E}_c(q) \\ & \quad + \frac{2\gamma\bar{A}\sqrt{N+1}(M+1)}{\beta M} \|\mathbf{U}\| \max\{\bar{A}, \sqrt{\beta\bar{C}}\} \max_{(t-2M)_+ \leq q \leq t} \|\widehat{\mathbf{v}}(q)\|. \end{aligned} \quad (\text{A.32})$$

Finally, combining (A.22), (A.24), (A.32) shows:

$$\begin{pmatrix} \|\widehat{\mathbf{v}}(t+1)\| \\ \mathcal{E}_c(t+1) \\ \mathcal{E}_g(t+1) \end{pmatrix} \leq \mathbf{Q} \begin{pmatrix} \max_{(t-2M)_+ \leq q \leq t} \|\widehat{\mathbf{v}}(q)\| \\ \max_{(t-2M)_+ \leq q \leq t} \mathcal{E}_c(q) \\ \max_{(t-2M)_+ \leq q \leq t} \mathcal{E}_g(q) \end{pmatrix}, \quad (\text{A.33})$$

where the inequality sign is applied element-wisely, and \mathbf{Q} is a non-negative 3×3 matrix, defined by:

$$\begin{pmatrix} \theta(\gamma) + \gamma^2 \|\mathbf{U}\| \|\mathbf{U}^{-1}\| \bar{G}M(G + 2\bar{G}M) & \gamma\sqrt{N} \|\mathbf{U}\| (1 + \gamma\bar{G}M)(\rho + \bar{A}\sqrt{\beta N}) & 0 \\ 0 & \lambda & \gamma \\ \frac{2\gamma\bar{A}\sqrt{N+1}(M+1)}{\beta M} \|\mathbf{U}\| \max\{\bar{A}, \sqrt{\beta\bar{C}}\} & \sqrt{N} \frac{2\gamma\bar{A}^2(M+1)}{\beta M} + \frac{2(1+\lambda)}{M} & \lambda + \frac{\gamma\rho}{M} \end{pmatrix}, \quad (\text{A.34})$$

where $\theta(\gamma) := \|\mathbf{I} - \gamma\mathbf{A}\| = \|\mathbf{I} - \gamma\mathbf{G}\|$. Note that the upper bounds for $\|\mathbf{U}\|$ and $\|\mathbf{U}^{-1}\|$ are provided in (A.8). Furthermore, also note that the eigenvalues of \mathbf{G} are bounded in (A.7). We could set the stepsize γ to be sufficiently small such that $\theta(\gamma) := \|\mathbf{I} - \gamma\mathbf{G}\| < 1$.

Finally, we apply Lemmas 2 and 3 presented in Section A.1 to the recursive inequality in (A.32), which shows that each of $\|\mathbf{v}(t)\|$, $\mathcal{E}_c(t)$, $\mathcal{E}_g(t)$ converges linearly with t . Therefore, we conclude the proof of Theorem 1.

A.1 Two Useful Lemmas

In this section, we present two auxiliary lemmas that are used in the proof of Theorem 1. Our first lemma establish the linear convergence of vectors satisfying recursive relations similar to (A.32), provided the spectral radius of \mathbf{Q} is less than one. In addition, the second lemma verifies this condition for \mathbf{Q} defined in in (A.34).

Lemma 2. Consider a sequence of non-negative vectors $\{\mathbf{e}(t)\}_{t \geq 1} \subseteq \mathbb{R}^n$ whose evolution is characterized by $\mathbf{e}(t+1) \leq \mathbf{Q}\mathbf{e}([(t-M+1)_+, t])$ for all $t \geq 1$ and some fixed integer $M > 0$, where $\mathbf{Q} \in \mathbb{R}^{n \times n}$ is a matrix whose entries are nonnegative, and we define

$$\mathbf{e}(\mathcal{S}) := \begin{pmatrix} \max_{q \in \mathcal{S}} e_1(q) \\ \vdots \\ \max_{q \in \mathcal{S}} e_n(q) \end{pmatrix} \in \mathbb{R}^n \quad \text{for any subset } \mathcal{S} \subseteq \mathbb{N}.$$

Moreover, if \mathbf{Q} irreducible in the sense that there exists an integer m such that the entries of \mathbf{Q}^m are all positive, and the spectral radius of \mathbf{Q} , denoted by $\rho(\mathbf{Q})$, is strictly less than one, then for any $t \geq 1$, we have

$$\mathbf{e}(t) \leq \rho(\mathbf{Q})^{\lceil \frac{t-1}{M} \rceil} C_1 \mathbf{u}_1, \quad (\text{A.35})$$

where $\mathbf{u}_1 \in \mathbb{R}_{++}^n$ is the top right eigenvector of \mathbf{Q} and C_1 is a constant that depends on the initialization.

Proof. We shall prove the lemma using induction. By the Perron-Frobenius theorem, the eigenvector \mathbf{u}_1 associated with $\rho(\mathbf{Q})$ is unique and is an all-positive vector. Therefore, there exists C_1 such that

$$\mathbf{e}(1) \leq C_1 \mathbf{u}_1. \quad (\text{A.36})$$

Let us first consider the base case with $t = 2, \dots, M + 1$, i.e., $\lceil (t-1)/M \rceil = 1$. When $t = 2$, by (A.36) we have,

$$\mathbf{e}(2) \leq \mathbf{Q}\mathbf{e}(1) \leq C_1 \mathbf{Q}\mathbf{u}_1 = \rho(\mathbf{Q}) C_1 \mathbf{u}_1, \quad (\text{A.37})$$

which is valid as \mathbf{Q} , $\mathbf{e}(1)$, \mathbf{u}_1 are all non-negative. Furthermore, we observe that $\mathbf{e}(2) \leq C_1 \mathbf{u}_1$. Next when $t = 3$, we have

$$\mathbf{e}(3) \leq \mathbf{Q}\mathbf{e}([1, 2]) \stackrel{(a)}{\leq} C_1 \mathbf{Q}\mathbf{u}_1 = \rho(\mathbf{Q}) C_1 \mathbf{u}_1,$$

where (a) is due to the non-negativity of vectors/matrix and the fact $\mathbf{e}(1), \mathbf{e}(2) \leq C_1 \mathbf{u}_1$ as shown in (A.37). Telescoping using similar steps, one can show $\mathbf{e}(t) \leq \rho(\mathbf{Q}) C_1 \mathbf{u}_1$ for any $t = 2, \dots, M + 1$.

For the induction step, let us assume that (A.35) holds true for any t up to $t = pM + 1$. That is, we assume that the result holds for all t such that $\lceil (t-1)/M \rceil \leq p$. We shall show that it also holds for any $t = pM + 2, \dots, (p+1)M + 1$, i.e., $\lceil (t-1)/M \rceil = p + 1$. Observe that

$$\mathbf{e}(pM + 2) \leq \mathbf{Q}\mathbf{e}([(p-1)M + 2, pM + 1]) \leq C_1 \rho(\mathbf{Q})^p \mathbf{Q}\mathbf{u}_1 = \rho(\mathbf{Q})^{p+1} C_1 \mathbf{u}_1, \quad (\text{A.38})$$

where we have used the induction hypothesis. It is clear that (A.38) is equivalent to (A.35) with $t = pM + 2$. Similar upper bound can be obtained for $\mathbf{e}(pM + 3)$ as well. Repeating the same steps, we show that (A.35) is true for any $t = pM + 2, \dots, (p+1)M + 1$. Therefore, we conclude the proof of this lemma. **Q.E.D.**

The following Lemma shows that \mathbf{Q} defined in (A.34) satisfies the conditions required in the previous lemma. Combining these two lemmas yields the final step of the proof of Theorem 1.

Lemma 3. Consider the matrix \mathbf{Q} defined in (A.34), it can be shown that (a) \mathbf{Q} is an irreducible matrix in $\mathbb{R}^{3 \times 3}$; (b) there exists a sufficiently small γ such that $\rho(\mathbf{Q}) < 1$; and (c) as $N, M \gg 1$ and the graph is geometric, we can set $\gamma = \mathcal{O}(1/\max\{N^2, M^2\})$ and $\rho(\mathbf{Q}) \leq 1 - \mathcal{O}(1/\max\{N^2, M^2\})$.

Proof. Our proof is divided into three parts. The first part shows the straightforward irreducibility of \mathbf{Q} ; the second part gives an upper bound to the spectral radius of \mathbf{Q} ; and the last part derives an asymptotic bound on $\rho(\mathbf{Q})$ when $N, M \gg 1$.

Irreducibility of \mathbf{Q} To see that \mathbf{Q} is irreducible, notice that \mathbf{Q}^2 is a positive matrix, which could be verified by direct computation.

Spectral Radius of \mathbf{Q} In the sequel, we compute an upper bound to the spectral radius of \mathbf{Q} , and show that if γ is sufficiently small, then its spectral radius will be strictly less than one. First we note that $\theta(\gamma) = 1 - \gamma\alpha$ for some $\alpha > 0$ and the network connectivity satisfies $\lambda < 1$. Also note that $\rho > 0$. For notational simplicity let us define the following

$$\begin{aligned} a_1 &= \|\mathbf{U}\| \|\mathbf{U}^{-1}\| \overline{G}M(G + 2\overline{G}M), \quad a_2 = \|\mathbf{U}\| \sqrt{N}(\rho + \overline{A}\sqrt{\beta N}), \quad a_3 = \overline{G}M \|\mathbf{U}\| \sqrt{N}(\rho + \overline{A}\sqrt{\beta N}) \\ a_4 &= \frac{2\overline{A}\sqrt{N+1}(M+1)}{\beta M} \|\mathbf{U}\| \max\{\overline{A}, \sqrt{\beta C}\}, \quad a_5 = \frac{2\overline{A}^2(M+1)\sqrt{N}}{\beta M}, \quad a_6 = \frac{2(1+\lambda)}{M}. \end{aligned}$$

With the above shorthand definitions, the characteristic polynomial for \mathbf{Q} , denoted by $g: \mathbb{R} \rightarrow \mathbb{R}$, is given by

$$g(\sigma) = \det \begin{pmatrix} \sigma - (1 - \gamma\alpha + \gamma^2 a_1) & -\gamma a_2 - \gamma^2 a_3 & 0 \\ 0 & \sigma - \lambda & -\gamma \\ -\gamma a_4 & -\gamma a_5 - a_6 & \sigma - \left(\lambda + \frac{\gamma\rho}{M}\right) \end{pmatrix}.$$

By direct computation, we have

$$g(\sigma) = (\sigma - (1 - \gamma\alpha + \gamma^2 a_1))g_0(\sigma) - \gamma^3(a_2 + \gamma a_3)a_4 \quad (\text{A.39})$$

where

$$g_0(\sigma) := (\sigma - \lambda)^2 - \frac{\gamma\rho}{M}(\sigma - \lambda) - \gamma(\gamma a_5 + a_6). \quad (\text{A.40})$$

Notice that the two roots of the above polynomial can be upper bounded by:

$$\lambda + \frac{\gamma\rho}{2M} \pm \sqrt{\left(\frac{\gamma\rho\sqrt{N}}{2M}\right)^2 + \gamma(\gamma a_5 + a_6)} \leq \bar{\sigma} := \lambda + \frac{\gamma\rho}{M} + \sqrt{\gamma(\gamma a_5 + a_6)} \quad (\text{A.41})$$

In particular, for all $\sigma \geq \bar{\sigma}$, we have

$$g_0(\sigma) \geq (\sigma - \bar{\sigma})^2. \quad (\text{A.42})$$

Now, let us define

$$\sigma^* := \max \left\{ \frac{\gamma\alpha}{4} + 1 - \gamma\alpha + \gamma^2 a_1, \bar{\sigma} + \gamma \sqrt{\frac{4(a_2 + \gamma a_3)a_4}{\alpha}} \right\} \quad (\text{A.43})$$

Observe that for all $\sigma \geq \sigma^*$, it holds that

$$\begin{aligned} g(\sigma) &\geq (\sigma - (1 - \gamma\alpha + \gamma^2 a_1))(\sigma - \bar{\sigma})^2 - \gamma^3(a_2 + \gamma a_3)a_4 \\ &\geq \frac{\gamma\alpha}{4} \gamma^2 \frac{4(a_2 + \gamma a_3)a_4}{\alpha} - \gamma^3(a_2 + \gamma a_3)a_4 = 0. \end{aligned} \quad (\text{A.44})$$

Lastly, observe that $g(\sigma)$ is strictly increasing for all $\sigma \geq \sigma^*$. Combining with the Perron Frobenius theorem shows that $\rho(\mathbf{Q}) \leq \sigma^*$. Moreover, as $\lambda < 1$ and $\alpha > 0$, there exists a sufficiently small γ such that $\sigma^* < 1$. We conclude that $\rho(\mathbf{Q}) < 1$ in the latter case.

Asymptotic Rate when $M, N \gg 1$ We evaluate a sufficient condition on γ for the proposed algorithm to converge, *i.e.*, when $\sigma^* < 1$. Let us consider (A.43) and the first operand in the $\max\{\cdot\}$. The first operand is guaranteed to be less than one if:

$$\gamma \leq \frac{\alpha}{2a_1} \implies \frac{\gamma\alpha}{4} + 1 - \gamma\alpha + \gamma^2 a_1 \leq 1 - \frac{\gamma\alpha}{4}. \quad (\text{A.45})$$

Moreover, from the definition of a_1 , we note that this requires $\gamma = \mathcal{O}(1/M^2)$ if $M \gg 1$.

Next, we notice that for geometric graphs, we have $\lambda = 1 - c/N$ for some positive c . Substituting this into the second operand in (A.43) gives

$$1 - \frac{c}{N} + \frac{\gamma\rho}{M} + \sqrt{\gamma(\gamma a_5 + a_6)} + \gamma \sqrt{\frac{4(a_2 + \gamma a_3)a_4}{\alpha}} < 1. \quad (\text{A.46})$$

Therefore, (A.45) and (A.46) together give a sufficient condition for $\sigma^* < 1$.

To obtain an asymptotic rate when $M, N \gg 1$. Observe that $a_2 = \Theta(N)$, $a_3 = \Theta(NM)$, $a_4 = \Theta(\sqrt{N})$, $a_5 = \Theta(\sqrt{N})$, $a_6 = \Theta(1/M)$. Moreover, the condition (A.45) gives $\gamma = \mathcal{O}(1/M^2)$, therefore the left hand side of Eq. (A.46) can be approximated by

$$1 - \frac{c}{N} + \gamma\Theta\left(N^{\frac{3}{4}}\right) + \sqrt{\gamma}\Theta(1/\sqrt{M}) \quad (\text{A.47})$$

Setting the above to $1 - c/(2N)$ requires one to have $\gamma = \mathcal{O}(1/N^2)$.

Finally, the above discussions show that setting $\gamma = \mathcal{O}(1/\max\{N^2, M^2\})$ guarantees that $\sigma^* < 1$. In particular, we have $\sigma^* \leq \max\{1 - \gamma\frac{\alpha}{4}, 1 - c/(2N)\} = 1 - \mathcal{O}(1/\max\{N^2, M^2\})$ **Q.E.D.**

A.2 Derivation of Equation (A.6)

We establish (A.6) with details. Recall that $\underline{\mathbf{h}}(t)$ and $\underline{\mathbf{v}}(t)$ are defined in (A.5). We verify this equation for each block of $\underline{\mathbf{h}}(t)$. To begin with, for the first block, for $\mathbf{h}_\theta(t)$ defined in (A.3), we have

$$\mathbf{h}_\theta(t) = \rho\bar{\boldsymbol{\theta}}(t) + \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{A}}^\top \mathbf{w}_i^t = \rho(\bar{\boldsymbol{\theta}}(t) - \boldsymbol{\theta}^* + \boldsymbol{\theta}^*) + \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{A}}^\top \mathbf{w}_i^t.$$

Recall from (A.1) that $\rho\theta^* = -\frac{1}{N} \sum_{i=1}^N \hat{\mathbf{A}}^\top \mathbf{w}_i^*$, which implies that

$$\mathbf{h}_\theta(t) = \rho(\bar{\theta}(t) - \theta^*) + \sum_{i=1}^N \sqrt{\frac{\beta}{N}} \hat{\mathbf{A}}^\top \frac{1}{\sqrt{\beta N}} (\mathbf{w}_i^t - \mathbf{w}_i^*) = [\mathbf{G}\mathbf{v}(t)]_1, \quad (\text{A.48})$$

where $[\mathbf{G}\mathbf{v}(t)]_1$ denotes the first block of $\mathbf{G}\mathbf{v}(t)$.

It remains to establish the equation for the remaining blocks. For any $i \in \{1, \dots, N\}$, let us focus on the $i+1$ -th block. By the definition of $\mathbf{h}_{\mathbf{w}_i}(t)$ in (A.4), we have

$$-\sqrt{\frac{\beta}{N}} \mathbf{h}_{\mathbf{w}_i}(t) = -\sqrt{\frac{\beta}{N}} (\hat{\mathbf{A}}\bar{\theta}(t) - \hat{\mathbf{C}}\mathbf{w}_i^t - \hat{\mathbf{b}}_i) = -\sqrt{\frac{\beta}{N}} (\hat{\mathbf{A}}(\bar{\theta}(t) - \theta^*) + \hat{\mathbf{A}}\theta^* - \hat{\mathbf{C}}\mathbf{w}_i^t - \hat{\mathbf{b}}_i).$$

Again from (A.1), it holds that $\hat{\mathbf{A}}\theta^* = \mathbf{b}_i + \hat{\mathbf{C}}\mathbf{w}_i^*$. Therefore,

$$-\sqrt{\frac{\beta}{N}} (\hat{\mathbf{A}}\bar{\theta}(t) - \hat{\mathbf{C}}\mathbf{w}_i^t - \hat{\mathbf{b}}_i) = -\sqrt{\frac{\beta}{N}} \hat{\mathbf{A}}(\bar{\theta}(t) - \theta^*) + \beta \hat{\mathbf{C}} \frac{\mathbf{w}_i^t - \mathbf{w}_i^*}{\sqrt{\beta N}} = [\mathbf{G}\mathbf{v}(t)]_{i+1}, \quad (\text{A.49})$$

where $[\mathbf{G}\mathbf{v}(t)]_{i+1}$ denotes the $i+1$ -th block of $\mathbf{G}\mathbf{v}(t)$. Combining (A.48) and (A.49) gives the desired equality.

B Additional Experiments

An interesting observation from Theorem 1 is that the convergence rate of PD-DistIAG depends on M and the topology of the graph. The following experiments will demonstrate the effects of these on the algorithm, along with the effects of regularization parameter ρ .

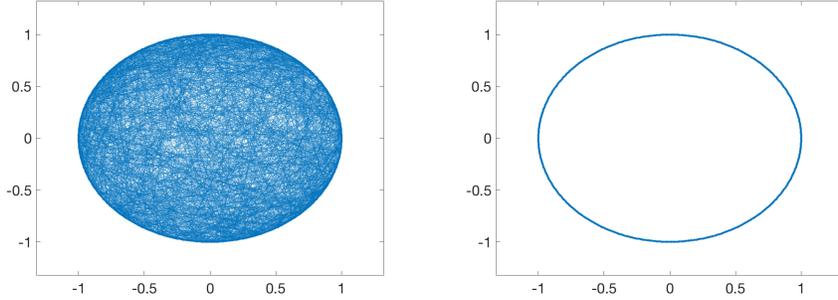


Figure B.1: Illustrating the graph topologies in the additional experiments. (Left) ER graph with connectivity probability of $1.01 \log N/N$. (Right) Ring graph.

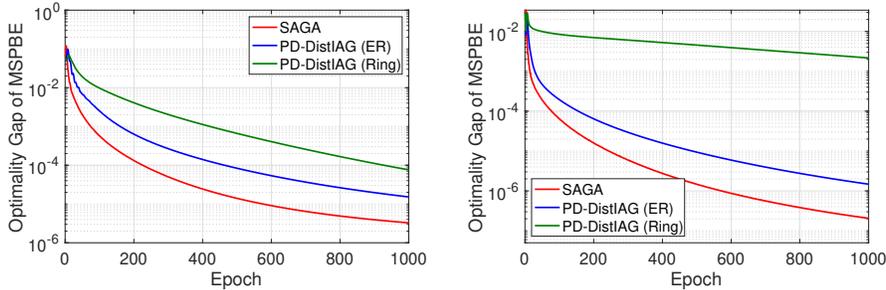


Figure B.2: Experiment with mountaincar dataset. For this problem, we only have $d = 300$, $M = 500$ samples, but yet there are $N = 500$ agents. (Left) We set $\rho = 0.01$. (Right) We set $\rho = 0.1$.

To demonstrate the dependence of PD-DistIAG on the graph topology, we fix the number of agents at $N = 500$ and compare the performances on the ring and the ER graph set with probability of

connection of $p = 1.01 \log N/N$, as illustrated in Fig. B.1. Notice that the ring graph is not a geometric graph and its connectivity parameter, defined as $\lambda := \lambda_{\max}(\mathbf{W} - (1/N)\mathbf{1}\mathbf{1}^\top)$ from the previous section can be much closer to 1 than the ER graph. Therefore, we expect the PD-DistIAG algorithm to converge slower on the ring graph. This is corroborated by Fig. B.2. Furthermore, from the figure, we observe that with a larger regularization ρ , the disadvantage for using the ring graph has exacerbated. We suspect that this is due to the fact that the convergence speed is limited by the graph connectivity, as seen in (A.34); while in the case of ER graph, the algorithm is able to exploit the improved problem's condition number.

Next, we consider the same set of experiment but increase the number of samples to $M = 5000$.

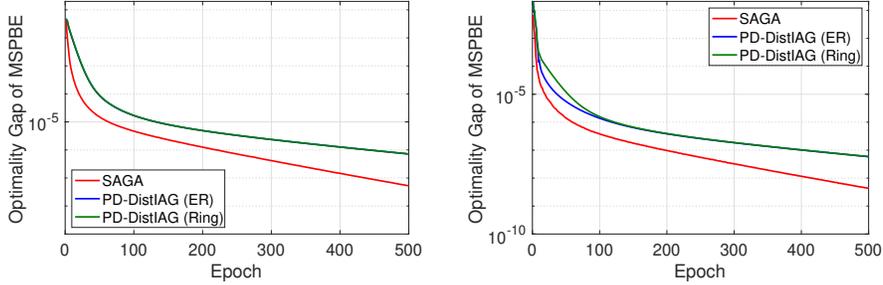


Figure B.3: Experiment with mountaincar dataset. For this problem, we have $d = 300$, $M = 5000$ samples, but yet there are $N = 500$ agents. (Left) We set $\rho = 0.01$. (Right) We set $\rho = 0.1$.

Interestingly, for this example, the performances of the ring graph and the ER graph settings are almost identical in this setting with large sample size M . This is possible as we recall from Theorem 1 that the algorithm converges at a rate of $\mathcal{O}(\sigma^t)$ where $\sigma = 1 - \mathcal{O}(1/\max\{MN^2, M^3\})$. As we have $M \gg N$, the impact from the sample size M becomes dominant, and is thus insensitive to the graph's connectivity.

References

- [1] G. Arslan and S. Yüksel. Decentralized Q-learning for stochastic teams and games. *IEEE Transactions on Automatic Control*, 62(4):1545–1558, 2017.
- [2] V. S. Borkar. *Stochastic approximation: A dynamical systems viewpoint*. Cambridge University Press, 2008.
- [3] D. S. Callaway and I. A. Hiskens. Achieving controllability of electric loads. *Proceedings of the IEEE*, 99(1):184–199, 2011.
- [4] A. Chambolle and T. Pock. On the ergodic convergence rates of a first-order primal–dual algorithm. *Mathematical Programming*, 159(1-2):253–287, 2016.
- [5] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.
- [6] J. Chen and A. H. Sayed. Diffusion adaptation strategies for distributed optimization and learning over networks. *IEEE Transactions on Signal Processing*, 60(8):4289–4305, 2012.
- [7] Y. Chen and M. Wang. Stochastic primal-dual methods and sample complexity of reinforcement learning. *arXiv preprint arXiv:1612.02516*, 2016.
- [8] P. Corke, R. Peterson, and D. Rus. Networked robots: Flying robot navigation using a sensor net. *Robotics Research*, pages 234–243, 2005.
- [9] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.
- [10] B. Dai, N. He, Y. Pan, B. Boots, and L. Song. Learning from conditional distributions via dual embeddings. *arXiv preprint arXiv:1607.04579*, 2016.

- [11] B. Dai, A. Shaw, N. He, L. Li, and L. Song. Boosting the actor with dual critic. *arXiv preprint arXiv:1712.10282*, 2017.
- [12] B. Dai, A. Shaw, L. Li, L. Xiao, N. He, J. Chen, and L. Song. Smoothed dual embedding control. *arXiv preprint arXiv:1712.10285*, 2017.
- [13] E. Dall’Anese, H. Zhu, and G. B. Giannakis. Distributed optimal power flow for smart microgrids. *IEEE Transactions on Smart Grid*, 4(3):1464–1475, 2013.
- [14] A. Defazio, F. Bach, and S. Lacoste-Julien. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in neural information processing systems*, pages 1646–1654, 2014.
- [15] S. S. Du, J. Chen, L. Li, L. Xiao, and D. Zhou. Stochastic variance reduction methods for policy evaluation. *arXiv preprint arXiv:1702.07944*, 2017.
- [16] H. R. Feyzmahdavian, A. Aytekin, and M. Johansson. A delayed proximal gradient method with linear convergence rate. In *Machine Learning for Signal Processing (MLSP), 2014 IEEE International Workshop on*, pages 1–6. IEEE, 2014.
- [17] J. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2137–2145, 2016.
- [18] J. Foerster, N. Nardelli, G. Farquhar, P. Torr, P. Kohli, S. Whiteson, et al. Stabilising experience replay for deep multi-agent reinforcement learning. *arXiv preprint arXiv:1702.08887*, 2017.
- [19] J. K. Gupta, M. Egorov, and M. Kochenderfer. Cooperative multi-agent control using deep reinforcement learning. In *International Conference on Autonomous Agents and Multi-agent Systems*, pages 66–83, 2017.
- [20] M. Gurbuzbalaban, A. Ozdaglar, and P. A. Parrilo. On the convergence rate of incremental aggregated gradient algorithms. *SIAM Journal on Optimization*, 27(2):1035–1048, 2017.
- [21] J. Hu and M. P. Wellman. Nash Q-learning for general-sum stochastic games. *Journal of Machine Learning Research*, 4(Nov):1039–1069, 2003.
- [22] R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.
- [23] J. Kober and J. Peters. Reinforcement learning in robotics: A survey. In *Reinforcement Learning*, pages 579–610. Springer, 2012.
- [24] M. Lauer and M. Riedmiller. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *International Conference on Machine Learning*, 2000.
- [25] D. Lee, H. Yoon, and N. Hovakimyan. Primal-dual algorithm for distributed reinforcement learning: Distributed gtd2. *arXiv preprint arXiv:1803.08031*, 2018.
- [26] X. Lian, M. Wang, and J. Liu. Finite-sum composition optimization via variance reduced gradient descent. *arXiv preprint arXiv:1610.04674*, 2016.
- [27] A. Lin and Q. Ling. Decentralized and privacy-preserving low-rank matrix completion. 2014. Preprint.
- [28] M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 157–163, 1994.
- [29] M. L. Littman. Value-function reinforcement learning in Markov games. *Cognitive Systems Research*, 2(1):55–66, 2001.
- [30] B. Liu, J. Liu, M. Ghavamzadeh, S. Mahadevan, and M. Petrik. Finite-sample analysis of proximal gradient td algorithms. In *UAI*, pages 504–513, 2015.

- [31] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *arXiv preprint arXiv:1706.02275*, 2017.
- [32] S. V. Macua, J. Chen, S. Zazo, and A. H. Sayed. Distributed policy evaluation under multiple behavior strategies. *IEEE Transactions on Automatic Control*, 60(5):1260–1274, 2015.
- [33] S. V. Macua, A. Tukiainen, D. G.-O. Hernández, D. Baldazo, E. M. de Cote, and S. Zazo. Diff-dac: Distributed actor-critic for multitask deep reinforcement learning. *arXiv preprint arXiv:1710.10363*, 2017.
- [34] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [35] A. Nedić and D. P. Bertsekas. Least squares policy evaluation algorithms with linear function approximation. *Discrete Event Dynamic Systems*, 13(1-2):79–110, 2003.
- [36] A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.
- [37] S. Omidshafiei, J. Papis, C. Amato, J. P. How, and J. Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *International Conference on Machine Learning*, pages 2681–2690, 2017.
- [38] B. Palaniappan and F. Bach. Stochastic variance reduction methods for saddle-point problems. In *Advances in Neural Information Processing Systems*, pages 1416–1424, 2016.
- [39] E. Parisotto, J. L. Ba, and R. Salakhutdinov. Actor-mimic: Deep multi-task and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342*, 2015.
- [40] S. Pu and A. Nedić. Distributed stochastic gradient tracking methods. *arXiv preprint arXiv:1805.11454*, 2018.
- [41] G. Qu and N. Li. Harnessing smoothness to accelerate distributed optimization. *IEEE Transactions on Control of Network Systems*, 2017.
- [42] M. Rabbat and R. Nowak. Distributed optimization in sensor networks. In *International Symposium on Information Processing in Sensor Networks*, pages 20–27, 2004.
- [43] M. Schmidt, N. Le Roux, and F. Bach. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1-2):83–112, 2017.
- [44] W. Shi, Q. Ling, G. Wu, and W. Yin. Extra: An exact first-order algorithm for decentralized consensus optimization. *SIAM Journal on Optimization*, 25(2):944–966, 2015.
- [45] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- [46] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [47] R. S. Sutton, H. R. Maei, and C. Szepesvári. A convergent $o(n)$ temporal-difference algorithm for off-policy learning with linear function approximation. In *Advances in neural information processing systems*, pages 1609–1616, 2009.
- [48] Y. W. Teh, V. Bapst, W. M. Czarnecki, J. Quan, J. Kirkpatrick, R. Hadsell, N. Heess, and R. Pascanu. Distral: Robust multi-task reinforcement learning. *arXiv preprint arXiv:1707.04175*, 2017.
- [49] J. Tsitsiklis, D. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.

- [50] M. Wang. Primal-dual π learning: Sample complexity and sublinear run time for ergodic markov decision problems. *arXiv preprint arXiv:1710.06100*, 2017.
- [51] X. Wang and T. Sandholm. Reinforcement learning to play an optimal Nash equilibrium in team Markov games. In *Advances in Neural Information Processing Systems*, pages 1603–1610, 2003.
- [52] A. Wilson, A. Fern, S. Ray, and P. Tadepalli. Multi-task reinforcement learning: A hierarchical Bayesian approach. In *International Conference on Machine Learning*, pages 1015–1022, 2007.
- [53] K. Zhang, Z. Yang, H. Liu, T. Zhang, and T. Başar. Fully decentralized multi-agent reinforcement learning with networked agents. *arXiv preprint arXiv:1802.08757*, 2018.
- [54] M. Zhu and S. Martínez. Discrete-time dynamic average consensus. *Automatica*, 46(2):322–329, 2010.