
Multi-Agent congestion cost minimization with linear function approximations

Prashant Trivedi

IE&OR, IIT Bombay, India

trivedi.prashant15@iitb.ac.in

Nandyala Hemachandra

IE&OR, IIT Bombay, India

nh@iitb.ac.in

Abstract

This work considers multiple agents traversing a network from a source node to the goal node. The cost to an agent for traveling a link has a private as well as a congestion component. The agent's objective is to find a path to the goal node with minimum overall cost in a decentralized way. We model this as a fully decentralized multi-agent reinforcement learning problem and propose a novel multi-agent congestion cost minimization (MACCM) algorithm. Our MACCM algorithm uses linear function approximations of transition probabilities and the global cost function. In the absence of a central controller and to preserve privacy, agents communicate the cost function parameters to their neighbors via a time-varying communication network. Moreover, each agent maintains its estimate of the global state-action value, which is updated via a multi-agent extended value iteration (MAEVI) sub-routine. We show that our MACCM algorithm achieves a sub-linear regret. The proof requires the convergence of cost function parameters, the MAEVI algorithm, and analysis of the regret bounds induced by the MAEVI triggering condition for each agent. We implement our algorithm on a two node network with multiple links to validate it. We first identify the optimal policy, the optimal number of agents going to the goal node in each period. We observe that the average regret is close to zero for 2 and 3 agents. The optimal policy captures the trade-off between the minimum cost of staying at a node and the congestion cost of going to the goal node. Our work is a generalization of learning the stochastic shortest path problem.

Keywords: Multi-Agent Systems; Decentralized Models; Congestion Cost; Private Cost; Network Model; Function Approximations; Value Iteration; Sub-Linear Regret; Stochastic Shortest Path.

1 INTRODUCTION

The shortest path problems are ubiquitous in many domains, such as driving directions on Google maps, automated warehouse systems, fleet management, and communication and computer networks. However, in most theoretical research, it is assumed that a single agent is traversing the network [Bellman \(1958\)](#); [Min et al. \(2022\)](#); [Vial et al. \(2022\)](#). So, the actual traversal cost does not factor in the crucial components such as congestion due to other agents and the agent's private travel efficiency. In this work, we consider a multi-agent setup where a set of agents traverse through a given network from a fixed initial/source node to a pre-specified goal node in a fully decentralized way. The cost to an agent for traveling a network link depends on two components: 1) congestion and 2) its private operational/fuel efficiency factor. Here the congestion is the number of agents using same link. The common objective of the agents is to find a path to the goal node in a completely decentralized way, and minimizing overall congestion cost while maintaining the agents' privacy. The decentralized setup has two major benefits over a centralized setup: 1) it can handle humongous state and action spaces, and 2) the agents can preserve the privacy of their actions and actual rewards. This is a generalization of the well-known learning stochastic shortest path (SSP) problem.

We model it as a fully decentralized multi-agent reinforcement learning (MARL) problem and propose a multi-agent congestion cost minimization (MACCM) algorithm. To incorporate privacy, we first parameterize the global cost function and share its parameters across the neighbors via a consensus matrix. Moreover, an agent's transition probability of going to the next node is also unknown to an agent. To this end, we propose the linear mixture MDP model, where the model parameters are expressed as the linear mixture of a given basis function. Our MACCM algorithm considers the privacy and works in an episodic manner. Each episode begins at a fixed initial node and ends if all

the agents reach the goal node. In the MACCM algorithm, each agent maintains an estimate of the global state-action value function and takes actions accordingly. This estimate is updated according to a multi-agent extended value iteration (MAEVI) sub-routine when a ‘doubling criteria’ is triggered. The intuition of using this updated estimate is that it will suggest a ‘better’ policy. We show that the updated optimistic estimator indeed provides a better policy, and our algorithm achieves a sub-linear regret. Specifically, our main contributions are:

a) In Section 3 we introduce a multi-agent version of SSP and propose a fully decentralized MACCM algorithm and show that it achieves a sub-linear regret (in Section 4). The regret depends on \sqrt{nK} and c_{\min} , where n is the number of agents, K is the number of episodes, and c_{\min} is the minimum cost of staying at any node except the goal node.

b) To prove the regret bound (in Section 4), we first show the convergence of the consensus based cost function parameters via the stochastic approximation method. Moreover, we show the convergence of the MAEVI algorithm. Finally, we separately bound the regret terms induced by the agents for whom MAEVI is triggered or not. The results of Min et al. (2022) and Vial et al. (2022) that consider learning SSP are special cases of our work (Remark 1).

c) To validate the usefulness of our algorithm, we provide some computational evidence on a hard instance in Section 5. In particular, we consider a network with two nodes and multiple links on each node. The average regret is very close to zero for 2 and 3 agents’ cases. The regret computation requires an optimal policy, the optimal number of agents going to the goal node in each period. This optimal policy captures the agents trade-off between the minimum cost of staying at the initial node and the congestion based cost of going to the goal node.

2 PROBLEM SETTING

Let (Z, E) be a given network, where $Z = \{s_{init}, 1, 2, \dots, g, q\}$ is the set of nodes and $E = \{(i, j) \mid i, j \in Z\}$ are the set of edges in the network, where s_{init} and g are fixed initial and the goal nodes respectively. Let $N = \{1, 2, \dots, n\}$ be the set of agents. The common objective of agents is to traverse through the network from initial node s_{init} to a goal node g while minimizing the sum of all agents’ path travel costs. The cost incurred by an agent includes a private efficiency component and another component that is congestion based, as given in Eq. (1) below. To achieve this objective and to preserve privacy, we model this problem as a fully decentralized multi-agent reinforcement learning (MARL) and provide a Multi-Agent Congestion Cost Minimization (MACCM) algorithm that achieves a sub-linear regret.

Formally, an instance of MACCM problem is described as

$(N, \mathcal{S} \cup \{s_{init}\}, \mathbf{g}, \{\mathcal{A}^i\}_{i \in N}, \{c^i\}_{i \in N}, \mathbb{P}, \{\mathcal{G}_t\}_{t \geq 0})$. Here $\mathcal{S} \cup \{s_{init}\}$ is the global state space with a fixed source state s_{init} . Each $s \in \mathcal{S}$ is a vector of size n representing the node at which each agent is present in that order, that is $s = (s^1, s^2, \dots, s^n)$, where $s^i \in Z$ for each agent $i \in N$. Often we write $s = (s^i, s^{-i})$ to denote the state s , where agent i is at node s^i and $s^{-i} = (s^1, \dots, s^{i-1}, s^{i+1}, \dots, s^n)$ is the node of all but agent i .

Let $\mathcal{A}^i(s)$ be the set of links available to agent $i \in N$ when the global state is s . So, the global action at the state s is $\mathcal{A}(s) = \prod_{i \in N} \mathcal{A}^i(s)$. A typical element in $\mathcal{A}(s)$ is a vector of size n , one for each agent as $(\mathbf{a}^1(s), \mathbf{a}^2(s), \dots, \mathbf{a}^n(s))$, where $\mathbf{a}^i(s) \in \mathcal{A}^i(s)$. Here $\mathbf{a}^i(s)$ represents the action taken by agent i when the global state is s . However, the action taken by each agent is private information and not known to other agents. For a global state s , and global action \mathbf{a} , each agent $i \in N$ realizes a local cost $c^i(s, \mathbf{a})$. It is important to note that the cost of agent i depends on the action taken by other agents also. In this work, we assume that the cost $c^i(s, \mathbf{a})$ depends on two components. The first one is the private component that is local to the agent, and the second is the congestion component. The private component might represent the agent’s efficiency. The congestion is defined as the number of agents using the same link. In particular,

$$c^i(s, \mathbf{a}) = K^i(s, \mathbf{a}) \cdot \sum_{j \in N} \mathbb{1}_{\{a^j(s) = a^j(s)\}}, \quad (1)$$

where $K^i(s, \mathbf{a})$ is a bounded private component of the i -th agent cost, not known to other agents, and the summation of indicators is the congestion seen by agent i in the state s . This implies, $c^i(\cdot, \cdot)$ is bounded. We assume all the costs are realized/incurred just before a period ends.

Moreover, we assume that $c^i(\cdot, \cdot) \geq c_{\min} > 0$. This assumption is common in many recent works for single agent Stochastic Shortest Path (SSP) problem (Rosenberg et al., 2020; Tarbouriech et al., 2020). $c_{\min} > 0$ ensures that agents do not have the incentive to wait indefinitely at any node for the clearance of congestion. Suppose $s = (g, s^{-i})$, that is agent i has reached the goal state g and the other agents are somewhere else in the network, then agent i stays at goal node, and $c^i((g, s^{-i}), \mathbf{a}) = 0$. Moreover, once an action \mathbf{a} is taken in state s the state will change to s' with probability $\mathbb{P}(s'|s, \mathbf{a})$ and this continue till $s = (g, \dots, g) = \mathbf{g}$, i.e., all the agents reach goal state.

Let $T^\pi(s)$ be the time to reach the goal state \mathbf{g} starting from the state s and following policy π . A stationary and deterministic policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ is called a proper policy if $T^\pi(s) < \infty$ almost surely. Let Π_p be the set of all stationary, deterministic and proper policies. We make the following assumption about proper policy, common in SSP literature Tarbouriech et al. (2020); Cohen et al. (2021).

Assumption 1 (Proper Policy Existence). *There exists at least a proper policy, meaning $\Pi_p \neq \emptyset$.*

Next, we define the cost-to-go or the value function for a given policy π

$$V^\pi(\mathbf{s}) := \lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\sum_{t=1}^T \bar{c}(\mathbf{s}_t, \pi(\mathbf{s}_t)) \middle| \mathbf{s}_1 = \mathbf{s} \right], \quad (2)$$

where $\bar{c}(\mathbf{s}_t, \pi(\mathbf{s}_t)) := \frac{1}{n} \sum_{i=1}^n c^i(\mathbf{s}_t, \pi(\mathbf{s}_t))$. So, our global objective translates to finding a policy π^* such that $\pi^* = \arg \min_{\pi \in \Pi_p} V^\pi(\mathbf{s}_{init})$. Let $V^* = V^{\pi^*}$ be the value of the optimal policy π^* . We also define the state-action value function $Q^\pi(\mathbf{s}, \mathbf{a})$ of a policy π as

$$\lim_{T \rightarrow \infty} \mathbb{E}_\pi \left[\bar{c}(\mathbf{s}_1, \mathbf{a}_1) + \sum_{t=2}^T \bar{c}(\mathbf{s}_t, \pi(\mathbf{s}_t)) \middle| \mathbf{s} = \mathbf{s}_1, \mathbf{a} = \mathbf{a}_1 \right]. \quad (3)$$

Since $\bar{c}(\cdot, \cdot)$ is bounded, for any proper policy $\pi \in \Pi_p$, both $V^\pi(\cdot)$ and $Q^\pi(\cdot, \cdot)$ are also bounded. This work assumes that the transition probability function \mathbb{P} is written as the linear mixture of given basis functions (Min et al., 2022; Vial et al., 2022). In particular, we make the following assumption about the transition probability function.

Assumption 2 (Transition probability approximation). *Suppose the feature mapping $\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^{nd}$ is known and pre-given. There exists a $\theta^* \in \mathbb{R}^{nd}$ with $\|\theta^*\|_2 \leq \sqrt{nd}$ such that $\mathbb{P}(s'|s, \mathbf{a}) = \langle \phi(s'|s, \mathbf{a}), \theta^* \rangle$ for any triplet $(s', \mathbf{a}, s) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$. Also, for a bounded function $V : \mathcal{S} \mapsto [0, B]$, it holds that $\|\phi_V(s, \mathbf{a})\|_2 \leq B\sqrt{nd}$, where $\phi_V(s, \mathbf{a}) = \sum_{s' \in \mathcal{S}} \phi(s'|s, \mathbf{a})V(s')$.*

For simplicity of notation, given any function $V : \mathcal{S} \rightarrow [0, B]$ we define $\mathbb{P}V(s, \mathbf{a}) = \sum_{s' \in \mathcal{S}} \mathbb{P}(s'|s, \mathbf{a})V(s')$, $\forall (s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$. Then, under the Assumption 2 we have,

$$\mathbb{P}V(s, \mathbf{a}) = \sum_{s' \in \mathcal{S}} \langle \phi(s'|s, \mathbf{a}), \theta^* \rangle V(s') = \langle \phi_V(s, \mathbf{a}), \theta^* \rangle.$$

Moreover, for any function $V : \mathcal{S} \rightarrow [0, B]$, we define the Bellman operator \mathcal{L} as $\mathcal{L}V(s) := \min_{\mathbf{a} \in \mathcal{A}} \{\bar{c}(s, \mathbf{a}) + \mathbb{P}V(s, \mathbf{a})\}$. Throughout, we assume that B_* is the upper bound on the optimal value function V^* , i.e., $B_* = \max_{s \in \mathcal{S}} V^*(s)$. Without loss of generality, we assume that $B_* \geq 1$, and denote the optimal state-action value by $Q^* = Q^{\pi^*}$, which satisfy the following Bellman equation for all $(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$

$$Q^*(s, \mathbf{a}) = \bar{c}(s, \mathbf{a}) + \mathbb{P}V^*(s, \mathbf{a}); V^*(s) = \min_{\mathbf{a} \in \mathcal{A}} Q^*(s, \mathbf{a}).$$

However, in our decentralized model, the global cost $\bar{c}(\cdot, \cdot)$ is unknown to any agent. So, at every decision epoch, each agent shares some parameters of the model using a time-varying communication network \mathcal{G}_t to its neighbors. To this end, we propose to estimate the globally averaged cost function \bar{c} . Let $\bar{c}(\cdot, \cdot; \mathbf{w}) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ be the class of parameterized functions where $\mathbf{w} \in \mathbb{R}^k$ for some $k \ll |\mathcal{S}| |\mathcal{A}|$. To obtain the estimate $\bar{c}(\cdot, \cdot; \mathbf{w})$ we seek to minimize the following least square estimate

$$\min_{\mathbf{w}} \mathbb{E}_{s, \mathbf{a}} [\bar{c}(s, \mathbf{a}) - \bar{c}(s, \mathbf{a}; \mathbf{w})]^2. \quad (\text{OP } 1)$$

A key result that ensures the working of a decentralized algorithm is the following; see also Zhang et al. (2018); Trivedi and Hemachandra (2022)

Proposition 1. *The optimization problem in Eq. (OP 1) is equivalently characterized as (both have the same stationary points)*

$$\min_{\mathbf{w}} \sum_{i=1}^n \mathbb{E}_{s, \mathbf{a}} [c^i(s, \mathbf{a}) - \bar{c}(s, \mathbf{a}; \mathbf{w})]^2. \quad (\text{OP } 2)$$

The proof details are available in the Appendix A.1. Note that the objective function in Eq. (OP 2) has the same form with separable objectives over agents as in the distributed optimization literature Nedic and Ozdaglar (2009); Boyd et al. (2006). This motivates the following updates for parameters of the global cost function estimate by agent i , \mathbf{w}^i to minimize the objective in Eq. (OP 2)

$$\begin{aligned} \tilde{\mathbf{w}}_t^i &\leftarrow \mathbf{w}_t^i + \gamma_t \cdot [c_t^i(\cdot, \cdot) - \bar{c}(\cdot, \cdot; \mathbf{w}_t^i)] \cdot \nabla_{\mathbf{w}} \bar{c}(\cdot, \cdot; \mathbf{w}_t^i) \\ \mathbf{w}_{t+1}^i &= \sum_{j \in N} l_t(i, j) \tilde{\mathbf{w}}_t^j, \end{aligned} \quad (4)$$

where $l_t(i, j)$ is the (i, j) -th entry of the consensus matrix L_t obtained using communication network \mathcal{G}_t at time t . γ_t is the step-size satisfying $\sum_t \gamma_t = \infty$ and $\sum_t \gamma_t^2 < \infty$, and $\bar{c}(\cdot, \cdot; \mathbf{w}_t^i)$ is the estimate of global cost function by agent i at time t . In the above equation, at time t , each agent updates an intermediate cost function parameters $\tilde{\mathbf{w}}_t^i$ using the stochastic gradient descent method to get the minima of the optimization problem given in (OP 2). Each agent shares these intermediate parameters to the neighbors via the communication matrix and update the true parameters \mathbf{w}^i .

The central idea of using the communication/consensus matrix in decentralized RL algorithms is to share some information among the consensus matrix neighbors while preserving privacy. Actions, which are private information, are not shared. However, a global objective cannot be attained without a central controller and without sharing any information or parameters. Thus, the cost function parameters \mathbf{w} 's (not actual costs) are shared with the neighbors as per the communication matrix. They converge to their true parameters *a.s.* (Theorem 2). Such a sharing of the cost function parameters via communication matrix to the neighbors is an intermediate construct that maintains privacy of each actions and costs, but, achieves the global objective, Equation (2). We make following assumption (Zhang et al., 2018; Bianchi et al., 2013) on communication matrix $\{L_t\}_{t \geq 0}$.

Assumption 3 (Consensus matrix). *The consensus matrices $\{L_t\}_{t \geq 0} \subseteq \mathbb{R}^{n \times n}$ satisfy (i) L_t is row stochastic, i.e., $L_t \mathbf{1} = \mathbf{1}$ and $\mathbb{E}(L_t)$ is column stochastic, i.e., $\mathbf{1}^\top \mathbb{E}(L_t) = \mathbf{1}^\top$. Further, there exists a constant $\kappa \in (0, 1)$ such that for*

any $l_t(i, j) > 0$, we have $l_t(i, j) \geq \kappa$; (ii) Consensus matrix L_t respects \mathcal{G}_t , i.e., $l_t(i, j) = 0$, if $(i, j) \notin \mathcal{E}_t$; (iii) The spectral norm of $\mathbb{E}[L_t^\top (I - \mathbb{1}\mathbb{1}^\top/n)L_t]$ is less than one.

We make the following assumption (Zhang et al., 2018; Trivedi and Hemachandra, 2022) on the features associated with the cost function while showing the convergence of the cost function parameters \mathbf{w} (Theorem 2).

Assumption 4 (Full rank). *For each agent $i \in N$, the cost function $\bar{c}(s, \mathbf{a})$ is parameterized as $\bar{c}(s, \mathbf{a}; \mathbf{w}) = \langle \psi(s, \mathbf{a}), \mathbf{w} \rangle$. Here $\psi(s, \mathbf{a}) = [\psi_1(s, \mathbf{a}), \dots, \psi_k(s, \mathbf{a})] \in \mathbb{R}^k$ are the features associated with pair (s, \mathbf{a}) . Further, we assume that these features are uniformly bounded. Moreover, let the feature matrix $\Psi \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}| \times k}$ have $[\psi_m(s, \mathbf{a}), s \in \mathcal{S}, \mathbf{a} \in \mathcal{A}]^\top$ as its m -th column for any $m \in [k]$, then Ψ has full column rank.*

Since the global cost is unknown, each agent uses the parameterized cost and maintains its estimate of $V(\cdot)$ and $Q(\cdot, \cdot)$. Let $V^i(\cdot)$ and $Q^i(\cdot, \cdot)$ be the estimate of these functions by agent i . So, the modified Bellman optimality equation for all (s, \mathbf{a}) and for all agents $i \in N$ is

$$\begin{aligned} Q^{i*}(s, \mathbf{a}; \mathbf{w}^i) &= \bar{c}(s, \mathbf{a}; \mathbf{w}^i) + \mathbb{P}V^{i*}(s, \mathbf{a}; \mathbf{w}^i) \\ V^{i*}(s; \mathbf{w}^i) &= \min_{\mathbf{a} \in \mathcal{A}} Q^{i*}(s, \mathbf{a}; \mathbf{w}^i) \end{aligned} \quad (5)$$

We later show in Theorem 2 that $\mathbf{w}_t^i \rightarrow \mathbf{w}^*$. Hence $\bar{c}(s, \mathbf{a}; \mathbf{w}_t^i) \rightarrow \bar{c}(s, \mathbf{a}; \mathbf{w}^*)$, $Q^{i*}(s, \mathbf{a}; \mathbf{w}_t^i) \rightarrow Q^{i*}(s, \mathbf{a})$ and $V^{i*}(s; \mathbf{w}_t^i) \rightarrow V^{i*}(s)$ as $\bar{c}(s, \mathbf{a}; \mathbf{w}_t^i)$, $Q^{i*}(s, \mathbf{a}; \mathbf{w}_t^i)$ and $V^{i*}(s; \mathbf{w}_t^i)$ are continuous functions of \mathbf{w}^i , where $Q^{i*}(s, \mathbf{a})$ and $V^{i*}(s)$ are defined as

$$Q^{i*}(s, \mathbf{a}) = \bar{c}(s, \mathbf{a}; \mathbf{w}^*) + \mathbb{P}V^{i*}(s, \mathbf{a}); \quad V^{i*}(s) = \min_{\mathbf{a} \in \mathcal{A}} Q^{i*}(s, \mathbf{a}). \quad (6)$$

With the above assumptions, we aim to design an algorithm for the episodic setting where an episode begins from a common initial state s_{init} and ends at \mathbf{g} such that the following regret over K episodes is minimized

$$R_K = \sum_{j=1}^K \sum_{l=1}^{I_j} \frac{1}{n} \sum_{i \in N} (\bar{c}(s_{j,l}, \mathbf{a}_{j,l}; \mathbf{w}_{j,l}^i) - K \cdot V^{i*}(s_{init})), \quad (7)$$

here I_j is the length of the episode $j = 1, \dots, K$, and $\bar{c}(s_{j,l}, \mathbf{a}_{j,l}; \mathbf{w}_{j,l}^i)$ is the estimate of the global cost function by agent i in the l -th step of the j -th episode. Note that in the above regret expression, instead of the global optimal value, we use the average of V^{i*} , averaged over all the agents. This is because (1) V^* is not available to any agent; however as mentioned above, each agent i maintains its estimate V^{i*} . (2) Theorem 2 implies $V^{i*} = V^*$ for all $i \in N$; however, we write V^{i*} in the regret definition to avoid any confusions, in both the Equations (6) and (7). Our proofs will remain the same, with this minor change. (3) We empirically observe that $V^{i*} = V^*$ for all $i \in N$. So, the regret definition in Equation (7) is the same as the true regret in terms of true optimal values. In the next section, we

present the MACCM algorithm that is fully decentralized and achieves a sub-linear regret.

3 MACCM ALGORITHM

We next describe the MACCM algorithm design. It is inspired by the single agent UCLK algorithm for discounted linear mixture MDPs of Zhou et al. (2021). It also uses some structure of the LEVIS algorithm of Min et al. (2022).

Let t be the global time index, and K be the number of episodes. Each episode starts at fixed state $s_{init} = (s_{init}, s_{init}, \dots, s_{init})$ and ends when all the agents reach to the goal state $\mathbf{g} = (g, g, \dots, g)$. An episode k is decomposed into many epochs; let j_i denote the j -th epoch of the agent i . Within this epoch, agent i uses $Q_{j_i}^i$ as an optimistic estimator of the global state-action value function.

Initially, for each agent $i \in N$, the estimate of the global state value function V^i and the global state-action value function Q^i is taken as 1 for all $s \neq \mathbf{g}$, and 0 for $s = \mathbf{g}$. In each episode k , if an agent $i \in N$ is not at the goal node, it takes action using the current optimistic estimator of the global state-action value $Q_{j_i}^i$. In particular, the action taken by agent i is according to the min max criteria that captures its best action against the worst possible action in terms of congestion by other agents. However, if agent i has reached the goal node, it will stay there until the episode gets over, i.e., till all other agents reach the goal node (lines 5-12).

(Lines 16-20) Apart from executing a policy that uses an optimistic estimator, each agent $i \in N$ also updates the Σ_t and b_t . These updates are used to estimate the true model parameters. Σ_t and b_t together are inspired from the ridge regression based minimizer of the model parameters; similar updates were used by Zhou et al. (2021); Abbasi-Yadkori et al. (2011) in single agent model. The doubling criteria are used to update the optimistic estimator of the state-action value function. The determinant doubling criteria reflects the diminishing returns of the underlying transition. However, more than this update is required as it cannot guarantee the finite length of the epoch, so a simple time doubling criteria is used. Moreover, the cost function parameters are also updated as given in Eq. (4).

In lines 21-26 of the algorithm, we maintain a set S_t containing those agents for whom the doubling criteria are satisfied at time t . The determinant doubling criteria is used in many previous works of linear bandits, and RL (Abbasi-Yadkori et al., 2011; Zhou et al., 2021). It is often referred to as lazy policy update. It reflects the diminishing returns of learning the underlying transitions. However, the determinant doubling criteria alone is insufficient and cannot guarantee the finite length of each epoch since feature norm $\|\phi_{V^i}(\cdot, \cdot)\|$ are not bounded from below. So, a simple time doubling criteria is introduced. This criterion possesses many excellent properties, including the easiness of

Algorithm 1 MACCM

```

1: Input: regularization parameter  $\lambda$ , confidence radius
    $\{\beta_t\}$ , an estimate  $B \geq B_*$ .
2: Initialize: set  $t \leftarrow 1$ . For each agent  $i \in N$ , set
    $j_i = 0, t_0^i = 0, \Sigma_0^i = \lambda I, b_0^i = 0, Q_0^i(s, \cdot), V_0^i(s) =$ 
    $1, \forall s \neq \mathbf{g}$ , and 0 otherwise;  $\mathbf{w}_0^i = \mathbf{0}; \gamma_t = \frac{1}{t+1}$ .
3: for  $k = 1, \dots, K$  do
4:   Set  $\mathbf{s}_t = \mathbf{s}_{init} = (s_{init}, s_{init}, \dots, s_{init})$ .
5:   while  $\mathbf{s}_t \neq \mathbf{g}$  do
6:     for  $i \in N$  do
7:       if  $s_t^i \neq \mathbf{g}$  then
8:          $\mathbf{a}_t^i = \arg \min_{\mathbf{a} \in \mathcal{A}_t^i} \max_{\mathbf{a}^- \in \mathcal{A}_t^{-i}} Q_{j_i}^i(\mathbf{s}_t, \mathbf{a}, \mathbf{a}^-)$ 
9:       else
10:         $\mathbf{a}_t^i = \mathbf{g}$ 
11:       end if
12:     end for
13:     Set  $\mathbf{a}_t = (\mathbf{a}_t^1, \mathbf{a}_t^2, \dots, \mathbf{a}_t^n)$ 
14:     receive cost  $\bar{c}_t(\mathbf{s}_t, \mathbf{a}_t; \mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \bar{c}(\mathbf{s}_t, \mathbf{a}_t; \mathbf{w}_t^i)$ 
15:     next state  $\mathbf{s}_{t+1} \sim \mathbb{P}(\cdot | \mathbf{s}_t, \mathbf{a}_t)$ 
16:     for  $i \in N$  do
17:       Set  $\tilde{\mathbf{w}}_t^i \leftarrow \mathbf{w}_t^i + \gamma_t [c^i(\mathbf{s}_t, \mathbf{a}_t) - \bar{c}(\mathbf{s}_t, \mathbf{a}_t; \mathbf{w}_t^i)] \cdot$ 
          $\nabla_{\mathbf{w}} \bar{c}(\mathbf{s}_t, \mathbf{a}_t; \mathbf{w}_t^i)$ 
18:       Set  $\Sigma_t^i \leftarrow \Sigma_{t-1}^i + \phi_{V_{j_i}^i}(\mathbf{s}_t, \mathbf{a}_t) \phi_{V_{j_i}^i}(\mathbf{s}_t, \mathbf{a}_t)^\top$ 
19:       Set  $b_t^i \leftarrow b_{t-1}^i + \phi_{V_{j_i}^i}(\mathbf{s}_t, \mathbf{a}_t) V_{j_i}^i(\mathbf{s}_{t+1})$ 
20:     end for
21:     Set  $S_t = \emptyset, S_t^c = N$ 
22:     for  $i \in N$  do
23:       if  $\det(\Sigma_t^i) \geq 2\det(\Sigma_{t_{j_i}^i}^i)$  or  $t \geq 2t_{j_i}^i$  then
24:         Set  $S_t = S_t \cup \{i\}$ ; and  $S_t^c = S_t^c \setminus \{i\}$ 
25:       end if
26:     end for
27:     for  $i \in S_t$  do
28:       Set  $j_i \leftarrow j_i + 1; t_{j_i}^i \leftarrow t$ , and  $\epsilon_{j_i} \leftarrow \frac{1}{t_{j_i}^i}$ 
29:        $\hat{\theta}_{j_i}^i \leftarrow \Sigma_t^{i-1} b_t^i$ 
30:       Set  $\mathcal{C}_{j_i}^i \leftarrow \left\{ \theta : \|\Sigma_{t_{j_i}^i}^{i/2} (\theta - \hat{\theta}_{j_i}^i)\|_2 \leq \beta_{t_{j_i}^i} \right\}$ 
31:       Set  $Q_{j_i}^i(\cdot, \cdot) \leftarrow \text{MAEVI}(\mathcal{C}_{j_i}^i, \epsilon_{j_i}, \frac{1}{t_{j_i}^i}, \mathbf{w}_t^i)$ 
32:       Set  $V_{j_i}^i(\cdot) = \min_{\mathbf{a}^- \in \mathcal{A}^-} Q_{j_i}^i(\cdot, \mathbf{a}^i, \mathbf{a}^-)$ 
33:     end for
34:     Send parameters to the neighbors
35:     Get  $\tilde{\mathbf{w}}_{t+1}^i = \sum_{j \in N} l_t(i, j) \tilde{\mathbf{w}}_t^j$ 
36:     Set  $t \leftarrow t + 1$ 
37:   end while
38: end for

```

implementation and the low space and time complexity.

(Lines 27-33) We switch the epoch for all agents $i \in S_t$ and update their optimistic estimator of the state-action value function using the multi-agent extended value iteration (MAEVI) subroutine. Moreover, we set the MAEVI error parameter $\epsilon_{j_i} = \frac{1}{t_{j_i}^i}$ to bound the cumulative error from

the value iterations by a constant, i.e., $(2t_{j_i} - t_{j_i}) \cdot \epsilon_{j_i} = 1$.

The optimism in the MACCM algorithm in period t is due to the construction of confidence set $\mathcal{C}_{j_i}^i$ for all the agents $i \in S_t$, which is input to the MAEVI sub-routine. The MAEVI sub-routine requires a confidence ellipsoid $\mathcal{C}_{j_i}^i$ containing true model parameters. This confidence set is constructed in line 30 of the MACCM algorithm, which is obtained via minimization of a suitable ridge regression problem with confidence radius β_t . Moreover, we construct a set \mathcal{B} to ensure that the model parameters form a valid transition probability function. In particular, the model parameters are taken from $\mathcal{C}_{j_i}^i \cap \mathcal{B}$. Here $\mathcal{B} := \{\theta : \forall (s, \mathbf{a}), \langle \phi(\cdot | s, \mathbf{a}), \theta \rangle \text{ is probability distribution and } \langle \phi(s' | s, \mathbf{a}) \rangle = \mathbb{1}_{\{s' \neq \mathbf{g}\}}\}$. In Theorem 3 we prove that the true model parameters θ^* are in the set $\mathcal{C}_{j_i}^i \cap \mathcal{B}$ with high probability.

The MAEVI algorithm uses a discount term q , which is key in the convergence of the MAEVI sub-routine. This is because $\langle \cdot, \phi_{V^i}(\cdot, \cdot) \rangle$ is not a contractive map, so we use an extra discount term $(1 - q)$ that provides the contraction property. This may lead to an additional bias that can be suppressed by suitably choosing a q . Particularly, we choose $q = \frac{1}{t_{j_i}^i}$, and it will yield an additional regret of $\mathcal{O}(\log T)$ in the final regret. The term $(1 - q)$ biases the estimated transition kernel towards the goal state \mathbf{g} , which also encourages further optimism. Similar design is also available in Min et al. (2022); Tarbouriech et al. (2021).

It is important to note that we use the stochastic approximation based rule to update the cost function parameters \mathbf{w}^i . We want to emphasize that MAEVI uses the most recent cost function parameters. These parameters are updated at each time period t via the consensus matrix in line 35 of the MACCM algorithm. The convergence of the cost function parameters ensures the convergence of state-action value function and these are used in regret analysis of the MACCM algorithm (Theorems 1 and 2).

4 MAIN RESULTS AND OUTLINE OF PROOFS

In this Section, we outline the main results. Due to space considerations, we defer all the proof details to the Appendix A. The following Theorem provides the bound on the regret R_K given in Eq. (7) for the MACCM algorithm.

Theorem 1. *Under the Assumptions 1, 2, for any $\delta > 0$, let $\beta_t = B \sqrt{nd \log \left(\frac{4}{\delta} (nt^2 + \frac{nt^3 B^2}{\lambda}) \right)} + \sqrt{\lambda nd}$, for all $t \geq 1$, where $B \geq B_*$ and $\lambda \geq 1$. Then, with a probability of at least $1 - \delta$, the regret of the MACCM algorithm satisfies*

$$R_K = \tilde{\mathcal{O}} \left(B^{1.5} d \sqrt{nk/c_{\min}} \cdot \log^2 \left(\frac{KBnd}{c_{\min} \delta} \right) + \frac{B^2 nd^2}{c_{\min}} \log^2 \left(\frac{KBnd}{c_{\min} \delta} \right) \right). \quad (8)$$

Algorithm 2 Multi-Agent EVI routine

1: **Input:** Set of agents S ; Confidence set $\mathcal{C}^k; \epsilon^k; \mathbf{w}^k, \forall k \in S$; discount term q .

2: **Initialize:** Set $Q^{k,(0)}(\cdot, \cdot), V^{k,(0)}(\cdot) = 0, V^{k,(-1)}(\cdot) = \infty, \forall k \in S$

3: **for** $k \in S$ **do**

4: **if** $\mathcal{C}^k \cap \mathcal{B} \neq \emptyset$ **then**

5: **while** $\|V^{k,(l)} - V^{k,(l-1)}\|_\infty \geq \epsilon^k$ **do**

6: $Q^{k,(l+1)}(\cdot, \cdot) = \bar{c}(\cdot, \cdot; \mathbf{w}^k)$
 $+ (1 - q) \min_{\theta \in \mathcal{C}^k \cap \mathcal{B}} \langle \theta, \phi_{V^{k,(l)}}(\cdot, \cdot) \rangle$

7: $V^{k,(l+1)}(\cdot) = \min_{\mathbf{a} \in \mathcal{A}} Q^{k,(l+1)}(\cdot, \mathbf{a})$

8: **end while**

9: Set $l \leftarrow l + 1$

10: **end if**

11: **end for**

12: Set $Q^k(\cdot, \cdot) \leftarrow Q^{k,(l+1)}(\cdot, \cdot), \forall k \in S$

13: **Output:** $Q^k(\cdot, \cdot), \forall k \in S$

If $B = O(B_\star)$, then the regret is $\tilde{O}(B_\star^{1.5} d \sqrt{nK/c_{\min}})$. The proof of Theorem 1 includes three major steps: 1) convergence of cost function parameters (Theorem 2); 2) MAEVI analysis (Theorem 3); 3) regret decomposition (Theorem 4). The proof is available in Appendix A.2

Convergence of cost function parameters: We next show the convergence of the cost function parameters \mathbf{w}^i . To this end, let $d(\mathbf{s})$ be the probability and stationary distribution of the Markov chain $\{\mathbf{s}_t\}_{t \geq 0}$ under policy π , and $\pi(\mathbf{s}, \mathbf{a})$ be the probability of taking action \mathbf{a} in state \mathbf{s} . Moreover, let $D^{\mathbf{s}, \mathbf{a}} = \text{diag}[d(\mathbf{s}) \cdot \pi(\mathbf{s}, \mathbf{a}), \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}]$ be the diagonal matrix with $d(\mathbf{s}) \cdot \pi(\mathbf{s}, \mathbf{a})$ as diagonal entries.

Theorem 2. *Under assumptions 3 and 4, with sequence $\{\mathbf{w}_t^i\}$, we have $\lim_t \mathbf{w}_t^i = \mathbf{w}^*$ almost surely for each agent $i \in N$, where \mathbf{w}^* is unique solution to*

$$\Psi^\top D^{\mathbf{s}, \mathbf{a}} (\Psi \mathbf{w}^* - \bar{c}) = 0. \quad (9)$$

The proof of this theorem uses the stochastic approximations of the single time-scale algorithms of Borkar (2022). The detailed proof is deferred to the Appendix A.3. The equation in above theorem is obtained by taking the first-order derivative of the least square minimization of the difference between the actual cost function and its linear parameterization.

Multi-Agent EVI analysis: Next, we show that the MAEVI algorithm converges in finite time to the optimistic estimator of the state-action value function. Specifically, we have the following theorem.

Theorem 3. *Let $\beta_t = B \sqrt{nd \log(\frac{4}{\delta}(nt^2 + \frac{nt^3 B^2}{\lambda}))} + \sqrt{\lambda nd}$, for all $t \geq 1$. Then with probability at least $1 - \delta/2$,*

and for each agent $i \in N$, for all $j_i \geq 1$, MAEVI converges in finite time, and the following hold: $\theta^ \in \mathcal{C}_{j_i}^i \cap \mathcal{B}$, $0 \leq Q_{j_i}^i(\cdot, \cdot) \leq Q^{i*}(\cdot, \cdot; \mathbf{w}^i)$, and $0 \leq V_{j_i}^i(\cdot) \leq V^{i*}(\cdot; \mathbf{w}^i)$.*

The proof of this theorem is deferred to Appendix A.4 due to space considerations.

Regret Decomposition: Next, we give the details of the regret decomposition. To this end, we first show that the total number of calls J to the MAEVI algorithm in the entire analysis is bounded. Let J^i be the total number of calls to the MAEVI algorithm made by agent i . Note that $J \leq \sum_{i \in N} J^i$. An agent $i \in N$ makes a call to the MAEVI algorithm if either the determinant doubling criteria or the time doubling is satisfied. Let J_1^i be the number of calls to MAEVI made via determinant doubling criteria, and J_2^i be the calls to the EVI algorithm via the time doubling criteria. Therefore, $J^i \leq J_1^i + J_2^i$.

Lemma 1. *The total number of calls to the MAEVI algorithm in the entire analysis, J , is bounded as*

$$J \leq 2n^2 d \log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + 2n \log(T). \quad (10)$$

The proof of the above lemma is deferred to the Appendix A.5. For the regret decomposition, we divide the time horizon into disjoint intervals $m = 1, 2, \dots, M$. The endpoint of an interval is decided by one of the two conditions: 1) MAEVI is triggered for at least one agent, and 2) all the agents have reached the goal node. This decomposition is explicitly used in the regret analysis only and not in the actual algorithm implementation. It is easy to observe that the interval length (number of periods in the interval) can vary; let H_m denote the length of the interval m . Moreover, at the end of the M th interval, all the K episodes are over. Therefore, the total length of all the intervals is $\sum_{m=1}^M H_m$, which is the same as $\sum_{k=1}^K T_k$, where T_k is the time to finish the episode k . Hence, both representations reflect the total time T to finish all the K episodes. Using the above interval decomposition, we write the regret R_K as

$$R_K = R(M) \leq \sum_{m=1}^M \sum_{h=1}^{H_m} \frac{1}{n} \sum_{i=1}^n \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}, \mathbf{w}^i) + 1 - \sum_{m \in \mathcal{M}(M)} \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{init}), \quad (11)$$

here $\mathcal{M}(M)$ is the set of all intervals that are the first intervals in each episode. In RHS we add 1 because $|V_0^i| \leq 1$. In the Theorem below, we decompose the regret R_K as

Theorem 4. *Assume that the event in Theorem 3 holds, then we have the following upper bound on the regret*

$$R(M) \leq E_1 + E_2 + 2n^2 dB_\star \log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + 2nB_\star \log(T) \quad (12)$$

where E_1 and E_2 are defined as

$$E_1 = \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n \left\{ \bar{c}(s_{m,h}, \mathbf{a}_{m,h}, \mathbf{w}^i) + \mathbb{P}V_{j_i(m)}^i(s_{m,h}, \mathbf{a}_{m,h}) - V_{j_i(m)}^i(s_{m,h}) \right\} \right],$$

$$E_2 = \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(s_{m,h+1}) - \frac{1}{n} \sum_{i=1}^n \mathbb{P}V_{j_i(m)}^i(s_{m,h}, \mathbf{a}_{m,h}) \right].$$

The proof of this theorem is deferred to the Appendix A.6. To complete the proof of Theorem 1, we bound E_1 and E_2 separately (Appendix C.3, C.4). Bounding E_1 uses all the intrinsic properties of MACCM algorithm 1. Unlike the single-agent setup, here, we separately consider the set of all agents for whom the MAEVI is triggered or not. Thus, E_1 is decomposed into $E_1(S_m)$ and $E_1(S_m^c)$ where S_m is the set of agents for whom MAEVI is triggered in m -th interval, and S_m^c are remaining agents. The bounds on $E_1(S_m)$ and $E_1(S_m^c)$ are specifically required in our multi-agent setup and are novel. Moreover, E_2 is the martingale difference sum, so it is bounded using the concentration inequalities.

5 COMPUTATIONAL EXPERIMENTS

In this Section, we provide the details of the computations to validate the usefulness of our MACCM algorithm. Consider a network with two nodes $\{s_{init}, g\}$. Thus, the number of states is 2^n , where n is the number of agents. In each state the actions available to each agent are $\mathcal{A}^i = \{-1, 1\}^{d-1}$, for some given $d \geq 2$. So, the total number of actions is 2^{nd} . Since the number of states and actions is exponentially large, we parameterize the transition probability. For each $(s', \mathbf{a}, s) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, the global transition probability is parameterized as $\mathbb{P}_\theta(s'|s, \mathbf{a}) = \langle \phi(s'|s, \mathbf{a}), \boldsymbol{\theta} \rangle$. The features $\phi(s'|s, \mathbf{a})$ are described below.

$$\begin{cases} (\phi(s'^1 | s^1, \mathbf{a}^1), \dots, \phi(s'^n | s^n, \mathbf{a}^n)), & \text{if } s \neq g, \\ \mathbf{0}_{nd}, & \text{if } s = g, s' \neq g, \\ (\mathbf{0}_{nd-1}, 2^{n-1}), & \text{if } s = g, s' = g, \end{cases}$$

where $\phi(s'^i | s^i, \mathbf{a}^i)$ is defined as

$$\phi(s'^i | s^i, \mathbf{a}^i) = \begin{cases} (-\mathbf{a}^i, \frac{1-\delta}{n})^\top, & \text{if } s^i = s'^i = s_{init} \\ (\mathbf{a}^i, \frac{\delta}{n})^\top, & \text{if } s^i = s_{init}, s'^i = g \\ \mathbf{0}_d^\top, & \text{if } s^i = g, s'^i = s_{init} \\ (\mathbf{0}_{d-1}, \frac{1}{n})^\top, & \text{if } s^i = g, s'^i = g. \end{cases}$$

Here $\mathbf{0}_d^\top = (0, 0, \dots, 0)^\top$ is a vector of d dimension with all zeros. Thus, the features $\phi(s'^i | s^i, \mathbf{a}^i) \in \mathbb{R}^{nd}$.

Moreover, the transition probability parameters are taken as $\boldsymbol{\theta} = (\boldsymbol{\theta}^1, \frac{1}{2^{n-1}}, \boldsymbol{\theta}^2, \frac{1}{2^{n-1}}, \dots, \boldsymbol{\theta}^n, \frac{1}{2^{n-1}})$ where $\boldsymbol{\theta}^i \in \left\{ -\frac{\Delta}{n(d-1)}, \frac{\Delta}{n(d-1)} \right\}^{d-1}$, and $\Delta < \delta$.

Lemma 2. *The features $\phi(s'|s, \mathbf{a})$ satisfy the following: (a) $\sum_{s'} \langle \phi(s'|s, \mathbf{a}), \boldsymbol{\theta} \rangle = 1, \forall s, \mathbf{a}$; (b) $\langle \phi(s' = \mathbf{g} | s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle = 1, \forall \mathbf{a}$; (c) $\langle \phi(s' \neq \mathbf{g} | s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle = 0, \forall \mathbf{a}$.*

The proof is deferred to the Appendix B.1. Recall, from Assumption 4, $\bar{c}(s, \mathbf{a}; \mathbf{w}) = \langle \psi(s, \mathbf{a}), \mathbf{w} \rangle$. We take the features as $\psi(s, \mathbf{a}) = (\psi(s^1, \mathbf{a}^1), \psi(s^2, \mathbf{a}^2), \dots, \psi(s^n, \mathbf{a}^n))$, where $\psi(s_{init}, \mathbf{a}^i) = \sum_{j=1}^n \mathbb{1}_{\{a^i = a^j | s^j = s_{init}\}}$, and $\psi(g, \mathbf{a}^i) = 0$ for any $\mathbf{a}^i \in \mathcal{A}^i$. That is the feature $\psi(s_{init}, \mathbf{a}^i)$ captures the congestion realized by any agent i present at s_{init} . For each agent $i \in N$ and for each s, \mathbf{a} pair the private component $K^i(s, \mathbf{a}) \sim \text{Uniform}(c_{\min}, 1)$. Further, each entry of the consensus matrix is taken as $1/n$, where n is the number of agents.

The intent of this small 2 node network is 3 fold: **(1)** This 2 node network is common in RL literature Min et al. (2022), as it depicts the worst-case performance in the form of a lower bound on the regret. **(2)** Though the network seems very small, the number of states is 2^n , and both the number of actions and model parameters are 2^{nd} , i.e., exponential in the number of agents, and the feature dimension. So, the model complexity increases exponentially in the number of agents offering a computationally challenging model. **(3)** The major computational head in our algorithm is due to the MAEVI sub-routine; in each step, we solve a computationally challenging discrete combinatorial optimization problem in model parameters. Given these constraints, we consider a small network with a low number of agents; however, our algorithm achieves sub-linear regret even in this hard instance. For a general network with any number of agents, a separate feature design and a suitable model parameters choice can make the MAEVI algorithm easy; such a feature design in itself is a complex problem.

Next, we compute the value of the optimal policy for the above model. We require it in the regret computations. First, note that the value of the optimal policy is defined as the sum of expected costs if all agents reach the goal state exactly in 1 time period, in 2 time period, and so on. Let x_j be the number of agents move to the goal node at each period $j = 1, 2, \dots, t-1$, and remaining $x_t = n - \sum_{j=1}^{t-1} x_j$ agents move to the goal node by t -th period. We call this sequence of departures of agents to the goal node as the ‘departure sequence’. For the above departure sequence, the cost incurred is $C_\alpha(x_1, \dots, x_{t-1}, x_t) =$

$$\alpha \sum_{j=1}^{t-1} x_j^2 + \sum_{j=1}^{t-1} \left(n - \sum_{i=1}^j x_i \right) \cdot c_{\min} + \alpha x_t^2, \quad (13)$$

where α is the mean of the uniform distribution $\mathcal{U}(c_{\min}, 1)$, i.e., $\alpha = \frac{c_{\min} + 1}{2}$. We use this α instead of the private cost to compute the optimal value. The first term in the above

equation is because all x_j agents have moved to the goal node in time period j ; hence the congestion is x_j . Moreover, each agent incurs the private cost α in this period. So, the cost incurred to any agent as per Eq. (1) is αx_j , and the number of agents moved are x_j , and hence the total cost incurred is $(\alpha x_j) \times x_j = \alpha x_j^2$. This happens for all the time periods $j = 1, 2, \dots, t-1$, so we sum this for $(t-1)$ periods. The remaining agents at each time period incurred a waiting cost of c_{\min} ; thus, we have a second term. Finally, the third term is because the remaining agents move to the goal node at the last period t . So, the optimal value, V^* is

$$V^* = \sum_{t=1}^{\infty} \mathbb{P}[x_1^*, \dots, x_t^*] \cdot C_{\alpha}(x_1^*, \dots, x_t^*) \quad (14)$$

where x_1^*, \dots, x_t^* are the optimal departure sequence. These x_j^* are obtained by minimizing the cost function $C_{\alpha}(x_1, \dots, x_t)$ in Eq. (13). Moreover, $\mathbb{P}[x_1^*, \dots, x_t^*]$ is the probability of occurrence of this optimal departure sequence. Note that, unlike the regret defined in Eq. (7) that use the cost function parameters, here we take the value of the optimal policy defined above. The following Theorem provides x_1^*, \dots, x_t^* and its cost $C_{\alpha}(x_1^*, \dots, x_t^*)$.

Theorem 5. *The optimal departure sequence is given by*

$$x_j^* = \left\lfloor \frac{n}{t} + \left(\frac{t+1}{2} - j \right) \cdot \frac{c_{\min}}{2\alpha} \right\rfloor, \forall j \in [t-1],$$

$$x_t^* = n - \sum_{j=1}^{t-1} x_j^* \quad (15)$$

The cost $C_{\alpha}(x_1^*, x_2^*, \dots, x_t^*)$ of using the above optimal departure sequence is

$$\alpha t \left(\frac{n}{t} \right)^2 + n(t-1) \cdot \frac{c_{\min}}{2\alpha} - \frac{t(t-1)(t+1)}{12} \cdot \frac{c_{\min}^2}{4\alpha^2}. \quad (16)$$

The proof of above theorem is deferred to the Appendix B.2. The above cost captures the trade-off between the minimum cost to an agent for staying at the initial node and the congestion cost. In particular, the first term captures all agents' total cost of going to the goal node. The remaining terms capture the minimum cost of staying at s_{init} .

Apart from the cost of the optimal departure sequence, we also require the probability of this departure sequence to compute V^* . To this end, we recall the feature design of transition probability that allows an agent to stay or depart from the initial node s_{init} . Note that agent i stays or departs from initial state iff the sign of the action \mathbf{a}^i matches the sign of the transition probability function parameter θ^i , i.e., $\text{sgn}(\mathbf{a}_j^i) = \text{sgn}(\theta_j^i)$ for all $j = 1, 2, \dots, d-1$ (more details are available in SM). Using this sign matching property, we have the following theorem for the transition probability.

Theorem 6. *The transition probability $\mathbb{P}[x_1^*, \dots, x_t^*]$ is given by*

$$\prod_{k=1}^{t-1} (1 - \gamma n + (\gamma - \eta) \sum_{j=1}^k x_j^*) \times (\gamma n - (\gamma - \eta) \sum_{j=1}^{t-1} x_j^*), \quad (17)$$

where $\gamma = \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right)$ and $\eta = \frac{1}{n \cdot 2^{n-1}}$.

The proof of this Theorem is deferred to the Appendix B.3. Using the above transition probability along with cost in Eq. (16) get V^* . However, it is hard to get the closed form expression of V^* , so we use its approximate value for regret computation. The approximate optimal value V_T^* in terms of a given large $T < \infty$ is

$$V_T^* = \sum_{t=1}^T \mathbb{P}[x_1^*, \dots, x_t^*] \cdot C_{\alpha}(x_1^*, \dots, x_t^*), \quad (18)$$

To obtain a better approximation of the optimal value V^* , we tune T . Moreover, note that we are approximating the




Figure 1: Average regret for $n = 2$ agents in green and $n = 3$ agents in red. Here $d = 2$, $\delta = 0.1$, $\Delta = 0.2$, $K = 7000$. $V_T^* = 2.15$ and 4.365 for 2 and 3 agents respectively. All the values are averaged over 15 runs.

private component of the cost by α . Thus, the above V_T^* will have some error; we minimize it in our computations by running the MAACM algorithm for multiple runs and average the regret over these runs. The average regret for 2 and 3 agents are very close to zero, as shown in Figure 1.

Remark 1. *Suppose $n = 1$ and $c(s, a) = c_{\min} = 1$ for all s, a . Then, the optimal departure sequence is $x_1^* = \dots = x_{t-1}^* = 0, x_t^* = 1$. So, from Eq. (13), the optimal cost is t . Moreover, the probability in Eq. (17) reduces to $(1 - \Delta - \delta)^{t-1}(\Delta + \delta)$ and hence $V^* = \sum_{t=1}^{\infty} (1 - \Delta - \delta)^{t-1}(\Delta + \delta) \cdot t = \frac{1}{\delta + \Delta}$. Thus, we recover Min et al. (2022) results for 2 node network with 1 agent (and hence with no congestion cost). Also, the features used in Min et al. (2022) and Vial et al. (2022) are interchangeable. So, we have more general results with extra complexity regarding congestion and privacy.*

6 RELATED WORK

The single agent SSP is well known for decades (Bertsekas, 2012). However, these assume the knowledge of the transition probabilities and the cost of each edge. Recently, there has been much work on the online SSP problem when the transition or the cost is not known or random. In such cases,

the RL based algorithms are proposed (Min et al., 2022; Vial et al., 2022; Tarbouriech et al., 2020, 2021). Many instances of SSP are run over multiple episodes using these algorithm, and the regret over K episodes of the SSP is defined.

The online SSP problem is first described in Tarbouriech et al. (2020) with $\tilde{O}(K^{2/3})$ regret. Later this is improved in Rosenberg et al. (2020). They gave an upper bound of $\tilde{O}(B_*|S|\sqrt{|A|K})$, and a lower bound of $\Omega(B_*\sqrt{|S||A|K})$ where B_* is an upper bound on the expected cost of the optimal policy. However, they assume that the cost functions are known and deterministic. Authors in Cohen et al. (2021) assume the cost function to be i.i.d. and initially unknown. They give an upper and a lower bound proving the optimal regret of $\Theta(\sqrt{(B_*^2 + B_*)|S||A|K})$. The algorithms proposed by Rosenberg et al. (2020) and Tarbouriech et al. (2020) uses the ‘‘optimism in the face of uncertainty (OFU)’’ principle that in-turn uses the ideas from UCRL2 algorithm of Jaksch et al. (2010) for average reward MDPs. Cohen et al. (2021) uses a black-box reduction of SSP to finite horizon MDPs. Similar reduction is used in Chen and Luo (2021); Chen et al. (2021) for SSP with adversarially changing costs. Cohen et al. (2021) gave a new algorithm named ULCVI for regret minimization in finite-horizon MDPs. Tarbouriech et al. (2021) extended the work of Cohen et al. (2021) to obtain a comparable regret bound for SSP without prior knowledge of expected time to reach the goal state.

Often the cost and the transition probabilities are unknown, and the state and action space is humongous. To this end, many researchers use function approximations of the transition probabilities, the per-period cost, or both (Wang et al., 2020; Yang and Wang, 2020). Some recent works in this direction are Jia et al. (2020); Min et al. (2022); Jin et al. (2020); Zhou et al. (2021). In particular, Min et al. (2022) proposes a LEVIS algorithm that uses the optimistic update of the estimated Q function using the extended value iteration (EVI) algorithm. Unlike Min et al. (2022), authors in Vial et al. (2022) uses OFU principle and parameterize the cost function also. Moreover, the feature design of the transition probabilities is also different; however, the features in these works are interchangeable. The regrets obtained in these works depend K , B_* , and c_{\min} . In our work, we simultaneously use the linear function approximations of the transition probabilities and the global cost function. We also provide a decentralized multi-agent algorithm incorporating congestion cost and agents’ privacy of cost.

7 DISCUSSION

This work considers a multi-agent variant of the optimal path-finding problem on a given network with pre-specified initial and goal nodes. The cost of traversing a link of the

network depends on the private cost of the agent (capturing the agent’s travel efficiency) and the congestion on the link. The unknown transition probability is the linear function of a given basis function. Moreover, each agent maintains an estimate of global cost function parameters, that are shared among the agents via a communication matrix.

We propose a fully decentralized multi-agent congestion cost minimization (MACCM) algorithm that achieves a sub-linear regret. In every episode of the MACCM algorithm, each agent maintains an optimistic estimate of the state-action value function; this estimate is updated according to a MAEVI sub-routine. The update happens if the ‘doubling criteria’ is triggered for any agent. To our knowledge, this is the first work that considers the multi-agent version of the congestion cost minimization problem over a network with linear function approximations. It has broader applicability in many real-life scenarios, such as decentralized fleet management. Experiments for 2 and 3 agent cases on a network validate our results.

The work we consider offers many challenges and future directions, and we mention some of these here. The current algorithm is based on the optimistic state-action value function updated according to doubling criteria. A better optimistic estimator with tighter regret bounds can be tried. For example, one can use a parameterized policy space and incorporate the feedback from the policy parameter in the state-action value function. For large networks we can explore a distinct feature design and suitable model parameters to address the scalability of computations. Further, a lower bound on the regret of our multi-agent congestion cost minimization, MACCM, algorithm is desirable. Decentralized learning of a multi-agent congestion cost minimization with capacity constraints on the network links (a generalization of capacitated multi-commodity flow problems) is a further possible research direction.

Acknowledgements

We thank the anonymous Reviewers and Meta Reviewer for their useful comments and encouragement. While working at this publication, Prashant Trivedi was partially supported by the Teaching Assistantship offered by Government of India. Some part of this work was done when Nandyala Hemachandra was visiting IIM Bangalore on a sabbatical leave.

Societal Impact

This work can only have a positive societal impact because it can be a good data-driven (RL) decentralised model for fleet management in transport, etc.

Code Release

Some details of the code for the MACCM algorithm implementations are available at <https://github.com/PRA1TRIVEDI/MACCM>. Refer to the Readme.md file for the description about how to run the code.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24.
- Bellman, R. (1958). On a routing problem. *Quarterly of applied mathematics*, 16(1):87–90.
- Bertsekas, D. (2012). *Dynamic programming and optimal control*, volume 1. Athena scientific.
- Bianchi, P., Fort, G., and Hachem, W. (2013). Performance of a distributed stochastic approximation algorithm. *IEEE Transactions on Information Theory*, 59(11):7405–7418.
- Borkar, V. S. (2022). *Stochastic approximation: a dynamical systems viewpoint. Second Edition*, volume 48. Springer.
- Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. (2006). Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530.
- Chen, L. and Luo, H. (2021). Finding the stochastic shortest path with low regret: The adversarial cost and unknown transition case. In *International Conference on Machine Learning*, pages 1651–1660. PMLR.
- Chen, L., Luo, H., and Wei, C.-Y. (2021). Minimax regret for stochastic shortest path with adversarial costs and known transition. In *Conference on Learning Theory*, pages 1180–1215. PMLR.
- Cohen, A., Efroni, Y., Mansour, Y., and Rosenberg, A. (2021). Minimax regret for stochastic shortest path. *Advances in Neural Information Processing Systems*, 34:28350–28361.
- Jaksch, T., Ortner, R., and Auer, P. (2010). Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4).
- Jia, Z., Yang, L., Szepesvári, C., and Wang, M. (2020). Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*, pages 666–686. PMLR.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. (2020). Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR.
- Kushner, H. and Yin, G. G. (2003). *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media.
- Metivier, M. and Priouret, P. (1984). Applications of a Kushner and Clark lemma to general classes of stochastic algorithms. *IEEE Transactions on Information Theory*, 30(2):140–151.
- Min, Y., He, J., Wang, T., and Gu, Q. (2022). Learning stochastic shortest path with linear function approximation. In *International Conference on Machine Learning*, pages 15584–15629. PMLR.
- Nedic, A. and Ozdaglar, A. (2009). Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61.
- Rosenberg, A., Cohen, A., Mansour, Y., and Kaplan, H. (2020). Near-optimal regret bounds for stochastic shortest path. In *International Conference on Machine Learning*, pages 8210–8219. PMLR.
- Tarbouriech, J., Garcelon, E., Valko, M., Pirotta, M., and Lazaric, A. (2020). No-regret exploration in goal-oriented reinforcement learning. In *International Conference on Machine Learning*, pages 9428–9437. PMLR.
- Tarbouriech, J., Zhou, R., Du, S. S., Pirotta, M., Valko, M., and Lazaric, A. (2021). Stochastic shortest path: Minimax, parameter-free and towards horizon-free regret. *Advances in Neural Information Processing Systems*, 34:6843–6855.
- Trivedi, P. and Hemachandra, N. (2022). Multi-agent natural actor-critic reinforcement learning algorithms. *Dynamic Games and Applications, Special Issue on Multi-agent Dynamic Decision Making and Learning*, edited by Konstantin Avrachenkov, Vivek S. Borkar and U. Jayakrishnan Nair, pages 1–31.
- Vial, D., Parulekar, A., Shakkottai, S., and Srikant, R. (2022). Regret bounds for stochastic shortest path problems with linear function approximation. In *International Conference on Machine Learning*, pages 22203–22233. PMLR.
- Wang, Y., Wang, R., Du, S. S., and Krishnamurthy, A. (2020). Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations*.
- Yang, L. and Wang, M. (2020). Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pages 10746–10756. PMLR.
- Zhang, K., Yang, Z., Liu, H., Zhang, T., and Basar, T. (2018). Fully decentralized multi-agent reinforcement learning with networked agents. In *International Conference on Machine Learning*, pages 5872–5881. PMLR.
- Zhou, D., Gu, Q., and Szepesvári, C. (2021). Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pages 4532–4576. PMLR.

APPENDIX

In this appendix we give the details of proofs omitted in the main results in Section A; further details of the feature design and the computations in Section B; proofs of the intermediate lemmas and propositions in Section C; and some useful results that we use from the existing literature in Section D.

For better readability, we first reiterate the result, and then give its proof.

A PROOFS AND DETAILS OF THE MAIN RESULTS

First, we give proofs of the results we omitted in the main paper.

A.1 Proof of Proposition 1

We first provide the equivalence of the optimization problems (OP 1) and (OP 2) obtain from the least square minimizer of the global cost function. Recall the optimization problem (OP 1) is

$$\min_{\mathbf{w}} \mathbb{E}_{s,\mathbf{a}}[\bar{c}(s,\mathbf{a}) - \bar{c}(s,\mathbf{a};\mathbf{w})]^2. \quad (\text{OP 1})$$

Recall the proposition: The optimization problem in (OP 1) is equivalently characterized as (both have the same stationary points)

$$\min_{\mathbf{w}} \sum_{i=1}^n \mathbb{E}_{s,\mathbf{a}}[c^i(s,\mathbf{a}) - \bar{c}(s,\mathbf{a};\mathbf{w})]^2. \quad (\text{OP 2})$$

Proof. Taking the first order derivative of the objective function in optimization problem (OP 1) w.r.t. \mathbf{w} , we have:

$$\begin{aligned} -2 \times \mathbb{E}_{s,\mathbf{a}}[\bar{c}(s,\mathbf{a}) - \bar{c}(s,\mathbf{a};\mathbf{w})] \times \nabla_{\mathbf{w}} \bar{c}(s,\mathbf{a};\mathbf{w}), &= -2 \times \mathbb{E}_{s,\mathbf{a}} \left[\frac{1}{n} \sum_{i \in N} c^i(s,\mathbf{a}) - \bar{c}(s,\mathbf{a};\mathbf{w}) \right] \times \nabla_{\mathbf{w}} \bar{c}(s,\mathbf{a};\mathbf{w}), \\ &= -\frac{2}{n} \times \mathbb{E}_{s,\mathbf{a}} \left[\sum_{i \in N} c^i(s,\mathbf{a}) - n \cdot \bar{c}(s,\mathbf{a};\mathbf{w}) \right] \times \nabla_{\mathbf{w}} \bar{c}(s,\mathbf{a};\mathbf{w}), \\ &= -\frac{2}{n} \times \mathbb{E}_{s,\mathbf{a}} \left[\sum_{i \in N} (c^i(s,\mathbf{a}) - \bar{c}(s,\mathbf{a};\mathbf{w})) \right] \times \nabla_{\mathbf{w}} \bar{c}(s,\mathbf{a};\mathbf{w}). \end{aligned}$$

Ignoring the factor $\frac{1}{n}$ in the above equation, we exactly have the first order derivative of the objective function in (OP 2). Thus, both optimization problems have the same stationary points. Hence, (OP 1) is an *equivalent characterization* of the optimization problem (OP 2). \square

A.2 Proof of Theorem 1

In this section, we give the proof of our main result (Theorem 1). It provides the upper bound on the regret of our MACCM algorithm. The proof of this theorem relies on an intermediate Lemma 3 (given below).

Recall the theorem: Under the Assumptions 1, 2, for any $\delta > 0$, let $\beta_t = B\sqrt{nd \log \left(\frac{4}{\delta} \left(nt^2 + \frac{nt^3 B^2}{\lambda} \right) \right)} + \sqrt{\lambda nd}$, for all $t \geq 1$, where $B \geq B_*$ and $\lambda \geq 1$. Then, with probability at least $1 - \delta$, the regret of the MACCM algorithm satisfies

$$R_K = \mathcal{O} \left(B^{1.5} d \sqrt{nK/c_{\min}} \cdot \log^2 \left(\frac{KBnd}{c_{\min} \delta} \right) + \frac{B^2 nd^2}{c_{\min}} \log^2 \left(\frac{KBnd}{c_{\min} \delta} \right) \right) \quad (19)$$

Proof. Note that the total cost incurred by each agent in K episodes is upper bounded by $R_K + KB_*$ and it is lower bounded by $c_{\min} \cdot n \cdot T$. This provides the relation between a fixed quantity K and the random quantity T . To complete the proof we use the following lemma.

Lemma 3. Under Assumptions 1 and 2, for any $\delta > 0$, let $\beta_t = B\sqrt{nd \log\left(\frac{4}{\delta}\left(nt^2 + \frac{nt^3 B^2}{\lambda}\right)\right)} + \sqrt{\lambda nd}$ for all $t \geq 0$, where $B \geq B_*$ where $\lambda \geq 1$. Then with the probability of at least $1 - \delta$, the regret of the MACCM algorithm satisfies

$$R_K \leq 10\beta_T \sqrt{Td \log\left(1 + \frac{TB_*^2}{\lambda}\right)} + 16n^2 dB_* \log\left(T + \frac{T^2 B_*^2 nd}{\lambda}\right),$$

where T is the total number of time periods.

The proof of this lemma is given in Appendix C.1. Using the above lemma, with probability at least $1 - \delta$, we have

$$c_{\min} \cdot n \cdot T \leq 10\beta_T \sqrt{Td \log\left(1 + \frac{TB_*^2}{\lambda}\right)} + 16n^2 dB_* \log\left(T + \frac{T^2 B_*^2 nd}{\lambda}\right) + KB_*.$$

Solving the above equation for T , we have

$$T = \mathcal{O}\left(\log^2\left(\frac{n}{\delta}\right) \cdot \left(\frac{KB_*}{nc_{\min}} + \frac{B_*^2 d^2}{c_{\min}^2}\right)\right).$$

Plugging this back into Lemma 3, we have the desired result of Theorem 1. \square

To complete the proof of above theorem, we need to prove the above Lemma 3. To this end, we require the convergence of the cost function parameters as stated in Theorem 2. Let $d(s)$ be the probability and stationary distribution of the Markov chain $\{s_t\}_{t \geq 0}$ under policy π , and $\pi(s, \mathbf{a})$ be the probability of taking action \mathbf{a} is state s . Moreover, let $D^{s, \mathbf{a}} = \text{diag}[d(s) \cdot \pi(s, \mathbf{a})]$ be the diagonal matrix with $d(s) \cdot \pi(s, \mathbf{a})$ as diagonal elements.

A.3 Proof of Theorem 2

Recall the theorem: Under the Assumptions 3 and 4, with sequence $\{\mathbf{w}_t^i\}$, we have $\lim_t \mathbf{w}_t^i = \mathbf{w}^*$ a.s. for each agent $i \in N$, where \mathbf{w}^* is unique solution to

$$\Psi^\top D^{s, \mathbf{a}} (\Psi \mathbf{w}^* - \bar{c}) = 0.$$

Proof. To prove the convergence of the cost function parameters, we use the following proposition to give bounds on \mathbf{w}_t^i for all $i \in N$. For proof, we refer to Zhang et al. (2018).

Proposition 2. Under Assumptions 3, and 4 the sequence $\{\mathbf{w}_t^i\}$ satisfy $\sup_t \|\mathbf{w}_t^i\| < \infty$ a.s., for all $i \in N$.

Let $\mathcal{F}_t = \sigma(c_\tau, \mathbf{w}_\tau, s_\tau, \mathbf{a}_\tau, L_{\tau-1}, \tau \leq t)$ be the filtration which is an increasing σ -algebra over time t . Define the following for notation convenience. Let $c_t = [c_t^1, \dots, c_t^n]^\top \in \mathbb{R}^n$, and $\mathbf{w}_t = [(\mathbf{w}_t^1)^\top, \dots, (\mathbf{w}_t^n)^\top]^\top \in \mathbb{R}^{nk}$. Moreover, let $A \otimes B$ represent the Kronecker product of any two matrices A and B . Let $y_t = [(y_t^1)^\top, \dots, (y_t^n)^\top]^\top$, where $y_{t+1}^i = [(c_{t+1}^i - \psi_t^\top \mathbf{w}_t^i) \psi_t^\top]^\top$. Recall, $\psi_t = \psi(s_t, \mathbf{a}_t)$. Let I be the identity matrix of the dimension $k \times k$. Then update of \mathbf{w}_t can be written as

$$\mathbf{w}_{t+1} = (L_t \otimes I)(\mathbf{w}_t + \gamma_t \cdot y_{t+1}). \quad (20)$$

Let $\mathbf{1} = (1, \dots, 1)$ represents the vector of all 1's. We define the operator $\langle \mathbf{w} \rangle = \frac{1}{n}(\mathbf{1}^\top \otimes I)\mathbf{w} = \frac{1}{n} \sum_{i \in N} \mathbf{w}_t^i$. This $\langle \mathbf{w} \rangle \in \mathbb{R}^k$ represents the average of the vectors in $\{\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^n\}$. Moreover, let $\mathcal{J} = (\frac{1}{n} \mathbf{1} \mathbf{1}^\top) \otimes I \in \mathbb{R}^{nk \times nk}$ is the projection operator that projects a vector into the consensus subspace $\{\mathbf{1} \otimes u : u \in \mathbb{R}^k\}$. Thus $\mathcal{J}\mathbf{w} = \mathbf{1} \otimes \langle \mathbf{w} \rangle$. Now define the disagreement vector $\mathbf{w}_\perp = \mathcal{J}_\perp \mathbf{w} = \mathbf{w} - \mathbf{1} \otimes \langle \mathbf{w} \rangle$, where $\mathcal{J}_\perp = I - \mathcal{J}$. Here I is $nk \times nk$ dimensional identity matrix. The iteration \mathbf{w}_t can be decomposed as the sum of a vector in disagreement space and a vector in consensus space, i.e., $\mathbf{w}_t = \mathbf{w}_{\perp, t} + \mathbf{1} \otimes \langle \mathbf{w}_t \rangle$. The proof of convergence consists of two steps.

Step 01: To show $\lim_t \mathbf{w}_{\perp, t} = 0$ a.s. From Proposition 2 we have $\mathbb{P}[\sup_t \|\mathbf{w}_t\| < \infty] = 1$, i.e., $\mathbb{P}[\cup_{K_1 \in \mathbb{Z}^+} \{\sup_t \|\mathbf{w}_t\| < K_1\}] = 1$. It suffices to show that $\lim_t \mathbf{w}_{\perp, t} \mathbf{1}_{\{\sup_t \|\mathbf{w}_t\| < K_1\}} = 0$ for any $K_1 \in \mathbb{Z}^+$. Lemma 5.5 in Zhang et al. (2018) proves the boundedness of $\mathbb{E}[\|\beta_t^{-1} \mathbf{w}_{\perp, t}\|^2]$ over the set $\{\sup_t \|\mathbf{w}_t\| \leq K_1\}$, for any $K_1 > 0$. We state the lemma here.

Proposition 3 (Lemma 5.5 in Zhang et al. (2018)). Under assumptions 3, and 4 for any $K_1 > 0$, we have

$$\sup_t \mathbb{E}[\|\beta_t^{-1} \mathbf{w}_{\perp, t}\|^2 \mathbf{1}_{\{\sup_t \|\mathbf{w}_t\| \leq K_1\}}] < \infty.$$

From Proposition 3 we obtain that for any $K_1 > 0$, $\exists K_2 < \infty$ such that for any $t \geq 0$, $\mathbb{E}[\|\mathbf{w}_{\perp,t}\|^2] < K_2\gamma_t^2$ over the set $\{\sup_t \|\mathbf{w}_t\| < K_1\}$. Since $\sum_t \gamma_t^2 < \infty$, by Fubini's theorem we have $\sum_t \mathbb{E}(\|\mathbf{w}_{\perp,t}\|^2 \mathbb{1}_{\{\sup_t \|\mathbf{w}_t\| < K_1\}}) < \infty$. Thus, $\sum_t \|\mathbf{w}_{\perp,t}\|^2 \mathbb{1}_{\{\sup_t \|\mathbf{w}_t\| < K_1\}} < \infty$ a.s. Therefore, $\lim_t \mathbf{w}_{\perp,t} \mathbb{1}_{\{\sup_t \|\mathbf{w}_t\| < K_1\}} = 0$ a.s. Since $\{\sup_t \|\mathbf{w}_t\| < \infty\}$ with probability 1, thus $\lim_t \mathbf{w}_{\perp,t} = 0$ a.s. This ends the proof of Step 01.

Step 02: To show the convergence of the consensus vector $\mathbb{1} \otimes \langle \mathbf{w}_t \rangle$, first note that the iteration of $\langle \mathbf{w}_t \rangle$ (Equation (20)) can be written as

$$\begin{aligned} \langle \mathbf{w}_{t+1} \rangle &= \frac{1}{N} (\mathbb{1}^\top \otimes I) (L_t \otimes I) (\mathbb{1} \otimes \langle \mathbf{w}_t \rangle + \mathbf{w}_{\perp,t} + \gamma_t y_{t+1}) \\ &= \langle \mathbf{w}_t \rangle + \gamma_t \langle (L_t \otimes I) (y_{t+1} + \gamma_t^{-1} \mathbf{w}_{\perp,t}) \rangle \\ &= \langle \mathbf{w}_t \rangle + \gamma_t \mathbb{E}(\langle y_{t+1} \rangle | \mathcal{F}_t) + \beta_t \xi_{t+1}, \end{aligned} \quad (21)$$

where

$$\begin{aligned} \xi_{t+1} &= \langle (L_t \otimes I) (y_{t+1} + \gamma_t^{-1} \mathbf{w}_{\perp,t}) \rangle - \mathbb{E}(\langle y_{t+1} \rangle | \mathcal{F}_t), \text{ and} \\ \langle y_{t+1} \rangle &= [(\bar{c}_{t+1} - \psi_t^\top \langle \mathbf{w}_t \rangle) \psi_t^\top]^\top. \end{aligned}$$

Note that $\mathbb{E}(\langle y_{t+1} \rangle | \mathcal{F}_t)$ is Lipschitz continuous in $\langle \mathbf{w}_t \rangle$. Moreover, ξ_{t+1} is a martingale difference sequence and satisfies

$$\mathbb{E}[\|\xi_{t+1}\|^2 | \mathcal{F}_t] \leq \mathbb{E}[\|y_{t+1} + \gamma_t^{-1} \mathbf{w}_{\perp,t}\|_{R_t}^2 | \mathcal{F}_t] + \|\mathbb{E}(\langle y_{t+1} \rangle | \mathcal{F}_t)\|^2, \quad (22)$$

where $R_t = \frac{L_t^\top \mathbb{1} \mathbb{1}^\top L_t \otimes I}{n^2}$ has bounded spectral norm. Bounding first and second terms in RHS of Equation (22), we have, for any $K_1 > 0$

$$\mathbb{E}[\|\xi_{t+1}\|^2 | \mathcal{F}_t] \leq K_3(1 + \|\langle \mathbf{w}_t \rangle\|^2),$$

over the set $\{\sup_t \|\mathbf{w}_t\| \leq K_1\}$ for some $K_3 < \infty$. Thus condition (3) of Assumption 5 is satisfied. The ODE associated with the Equation (21) has the form

$$\langle \dot{\mathbf{w}} \rangle = -\Psi^\top D^{s,a} \Psi \langle \mathbf{w} \rangle + \Psi^\top D^{s,a} \bar{c}. \quad (23)$$

Let the RHS of Equation (23) be $h(\langle \mathbf{w} \rangle)$. Note that $h(\langle \mathbf{w} \rangle)$ is Lipschitz continuous in $\langle \mathbf{w} \rangle$. Also, recall that $D^{s,a} = \text{diag}[d(s) \cdot \pi(s, \mathbf{a}), s \in \mathcal{S}, \mathbf{a} \in \mathcal{A}]$. Hence the ODE given in Equation (23) has unique globally asymptotically stable equilibrium \mathbf{w}^* satisfying

$$\Psi^\top D^{s,a} (\bar{c} - \Psi \mathbf{w}^*) = 0.$$

Moreover, from Propositions 2, and 3, the sequence $\{\mathbf{w}_t\}$ is bounded almost surely, so is the sequence $\{\langle \mathbf{w}_t \rangle\}$. Specializing Corollary 8.1 and Theorem 8.3 on page 114-115 in Borkar (2022) we have $\lim_t \langle \mathbf{w}_t \rangle = \mathbf{w}^*$ a.s. over the set $\{\sup_t \|\mathbf{w}_t\| \leq K_1\}$ for any $K_1 > 0$. This concludes the proof of Step 02.

The proof follows from Proposition 2 and results from Step 01. Thus, we have $\lim_t \mathbf{w}_t^i = \mathbf{w}^*$ a.s. for each $i \in N$. \square

Apart from the convergence of the cost function parameter, we also require the MAEVI analysis; in particular, we now prove the Theorem 3.

A.4 Proof of Theorem 3

Recall the theorem: Let $\beta_t = B \sqrt{nd \log \left(\frac{4}{3} (nt^2 + \frac{nt^3 B^2}{\lambda}) \right)} + \sqrt{\lambda nd}$, for all $t \geq 1$. Then with probability at least $1 - \delta/2$, and for each agent $i \in N$, for all $j_i \geq 1$, MAEVI converges in finite time, and the following holds

$$\theta^* \in C_{j_i}^i \cap \mathcal{B}, 0 \leq Q_{j_i}^i(\cdot, \cdot) \leq Q^{i*}(\cdot, \cdot; \mathbf{w}^i), 0 \leq V_{j_i}^i(\cdot) \leq V^{i*}(\cdot; \mathbf{w}^i)$$

Proof. To prove this Theorem, for each agent $i \in N$, we decompose t into different rounds. Each round $j_i \geq 1$ of agent i corresponds to $t \in [t_{j_i}^i + 1, t_{j_i+1}^i]$, during which the action-value function estimator is the output $Q_{j_i}^i$ of MAEVI sub-routine by agent i . We will apply the induction argument on the rounds to show that optimism holds for all $j_i \geq 1$.

Consider round $j_i = 1$ for agent $i \in N$. In this round, let $t \in [1, t_1^i]$. We have $V_0^i \leq B_*$, from the algorithm's initialization. Define $\eta_t^i = V_0^i(s_{t+1}) - \langle \phi_{V_0^i}(s_t, \mathbf{a}_t), \theta^* \rangle$ for $t \in [1, t_1^i]$. Then $\{\eta_t^i\}_{t=1}^{t_1^i}$ are B_* -sub-Gaussian.

Applying Theorem 7 of Abbasi-Yadkori et al. (2011) in our case, with probability at least $\left(1 - \frac{\delta}{n \cdot t_1^i(t_1^i + 1)}\right)$ we have,

$$\begin{aligned}
 \left\| \sum_t^{i-1/2} \sum_{k=1}^t \phi_{V_0^i}(s_k, \mathbf{a}_k) \eta_k^i \right\|_2 &\leq B_\star \sqrt{2 \log \left(\frac{\det(\Sigma_t^i)^{1/2}}{\delta / (n \cdot t_1^i(t_1^i + 1)) \cdot \det(\Sigma_{t-1}^i)^{1/2}} \right)} \\
 &\stackrel{(i)}{=} B_\star \sqrt{2 \log \left(\frac{\det(\Sigma_t^i)^{1/2}}{\delta / (n \cdot t_1^i(t_1^i + 1)) \cdot \lambda^{nd/2}} \right)} \\
 &\stackrel{(ii)}{\leq} B_\star \sqrt{2 \log \left(\frac{\left(\lambda + \frac{tB_\star^2 nd}{nd} \right)^{nd/2}}{\delta / (n \cdot t_1^i(t_1^i + 1)) \cdot \lambda^{nd/2}} \right)} \\
 &= B_\star \sqrt{2 \log \left(\frac{\lambda^{nd/2} \cdot \left(1 + \frac{tB_\star^2}{\lambda}\right)^{nd/2}}{\delta / (n \cdot t_1^i(t_1^i + 1)) \cdot \lambda^{nd/2}} \times \frac{(\delta / (n \cdot t_1^i(t_1^i + 1)))^{nd/2}}{(\delta / (n \cdot t_1^i(t_1^i + 1)))^{nd/2}} \right)} \\
 &= B_\star \sqrt{2 \log \left(\frac{\left(1 + \frac{tB_\star^2}{\lambda}\right)^{nd/2}}{(\delta / (n \cdot t_1^i(t_1^i + 1)))^{nd/2}} \times \frac{(\delta / (n \cdot t_1^i(t_1^i + 1)))^{nd/2}}{\delta / (n \cdot t_1^i(t_1^i + 1))} \right)} \\
 &= B_\star \sqrt{2 \log \left(\frac{1 + \frac{tB_\star^2}{\lambda}}{\delta / (n \cdot t_1^i(t_1^i + 1))} \right)^{nd/2} + 2 \log \left(\frac{\delta}{n \cdot t_1^i(t_1^i + 1)} \right)^{nd/2-1}} \\
 &\stackrel{(iii)}{\leq} B_\star \sqrt{2 \log \left(\frac{1 + \frac{tB_\star^2}{\lambda}}{\delta / (n \cdot t_1^i(t_1^i + 1))} \right)^{nd/2}} \\
 &= B_\star \sqrt{nd \log \left(\frac{1 + \frac{tB_\star^2}{\lambda}}{\delta / (n \cdot t_1^i(t_1^i + 1))} \right)} \\
 &= B_\star \sqrt{nd \log \left(\frac{n \cdot t_1^i(t_1^i + 1) + \frac{n \cdot t \cdot t_1^i(t_1^i + 1) B_\star^2}{\lambda}}{\delta} \right)}, \tag{24}
 \end{aligned}$$

where (i) follows from the fact that $\det(\Sigma_{t-1}^i) = \lambda^{nd}$; (ii) uses the determinant trace inequality (Theorem 8) along with the assumption that $\|\phi_{V_0^i}\| \leq B_\star \sqrt{nd}$; and (iii) uses the fact that $0 < \frac{\delta}{n \cdot t_1^i(t_1^i + 1)} \leq 1$ hence $\log \left(\frac{\delta}{n \cdot t_1^i(t_1^i + 1)} \right) < 0$.

Next, we consider the LHS of Equation (24) and give the lower bound for the same.

$$\begin{aligned}
 \left\| \sum_t^{i-1/2} \sum_{k=1}^t \phi_{V_0^i}(s_k, \mathbf{a}_k) \eta_k^i \right\|_2 &\stackrel{(i)}{=} \left\| \sum_t^{i-1/2} \sum_{k=1}^t \phi_{V_0^i}(s_k, \mathbf{a}_k) \cdot (V_0^i(s_{k+1}) - \langle \phi_{V_0^i}(s_k, \mathbf{a}_k), \boldsymbol{\theta}^\star \rangle) \right\|_2 \\
 &= \left\| \Sigma_t^{i/2} \Sigma_t^{i-1} \sum_{k=1}^t \phi_{V_0^i}(s_k, \mathbf{a}_k) \cdot V_0^i(s_{k+1}) - \Sigma_t^{i/2} \Sigma_t^{i-1} (\Sigma_t^i - \lambda \mathbf{I}) \boldsymbol{\theta}^\star \right\|_2 \\
 &\stackrel{(ii)}{=} \left\| \Sigma_t^{i/2} \hat{\boldsymbol{\theta}}_t^i - \Sigma_t^{i/2} \boldsymbol{\theta}^\star + \lambda \Sigma_t^{i-1/2} \boldsymbol{\theta}^\star \right\|_2 \\
 &\stackrel{(iii)}{\geq} \left\| \Sigma_t^{i/2} (\hat{\boldsymbol{\theta}}_t^i - \boldsymbol{\theta}^\star) \right\|_2 - \left\| \lambda \Sigma_t^{i-1/2} \boldsymbol{\theta}^\star \right\|_2 \\
 &\stackrel{(iv)}{\geq} \left\| \Sigma_t^{i/2} (\hat{\boldsymbol{\theta}}_t^i - \boldsymbol{\theta}^\star) \right\|_2 - \sqrt{\lambda nd}, \tag{25}
 \end{aligned}$$

where (i) follows by the definition of η_k^i ; (ii) uses the updates of Σ_t^i and definition of $\hat{\boldsymbol{\theta}}_t^i$ in the algorithm; (iii) is the consequence of the triangle inequality, i.e., $\|a\| - \|b\| \leq \|a + b\|$; and finally (iv) follows from the fact that $\|\boldsymbol{\theta}^\star\| \leq \sqrt{nd}$.

From Equations (24) and (25) we have the following:

$$B_\star \sqrt{nd \log \left(\frac{n \cdot t_1^i (t_1^i + 1) + \frac{n \cdot t_1^i (t_1^i + 1) B_\star^2}{\lambda}}{\delta} \right)} \geq \left\| \Sigma_t^{i/2} (\hat{\theta}_t^i - \theta^\star) \right\|_2 - \sqrt{\lambda nd}.$$

From the definition of β_t in this Theorem, we have

$$\left\| \Sigma_t^{i/2} (\hat{\theta}_t^i - \theta^\star) \right\|_2 \leq B_\star \sqrt{nd \log \left(\frac{n \cdot t_1^i (t_1^i + 1) + \frac{n \cdot t_1^i (t_1^i + 1) B_\star^2}{\lambda}}{\delta} \right)} + \sqrt{\lambda nd} \leq \beta_{t_1^i}.$$

Since above holds for all $t \in [1, t_1^i]$ with probability at least $1 - \frac{\delta}{n \cdot t_1^i (t_1^i + 1)}$, the true parameters $\theta^\star \in \mathcal{C}_1^i \cap \mathcal{B}$. This implies in round 1, for each agent $i \in N$, the true parameters are in set $\mathcal{C}_1^i \cap \mathcal{B}$. To complete round 1, we need to show that the output Q_1^i and V_1^i of MAEVI are optimistic for each agent $i \in N$. This will be done using the second induction argument on the loop of MAEVI. For the base step, it follows from the non-negativity of the $Q^{i^\star}(\cdot, \cdot; \mathbf{w}^i)$ and $V^{i^\star}(\cdot; \mathbf{w}^i)$ that $Q^{i,(0)}(\cdot, \cdot) \leq Q^{i^\star}(\cdot, \cdot; \mathbf{w}^i)$ and $V^{i,(0)}(\cdot) \leq V^{i^\star}(\cdot; \mathbf{w}^i)$.

Now, assume that $Q^{i,(l)}(\cdot, \cdot)$ and $V^{i,(l)}(\cdot, \cdot)$ are optimistic. For the $(l+1)$ -th iterate in MAEVI algorithm, we have

$$\begin{aligned} Q^{i,(l+1)}(\cdot, \cdot) &= \bar{c}(\cdot, \cdot; \mathbf{w}^i) + (1-q) \min_{\theta \in \mathcal{C}_1^i \cap \mathcal{B}} \langle \theta, \phi_{V^{i,(l)}}(\cdot, \cdot) \rangle \\ &\stackrel{(i)}{=} \bar{c}(\cdot, \cdot; \mathbf{w}^i) + (1-q) \langle \theta^\star, \phi_{V^{i,(l)}}(\cdot, \cdot) \rangle \\ &\stackrel{(ii)}{=} \bar{c}(\cdot, \cdot; \mathbf{w}^i) + (1-q) \mathbb{P} V^{i,(l)}(\cdot, \cdot) \\ &\stackrel{(iii)}{\leq} \bar{c}(\cdot, \cdot; \mathbf{w}^i) + \mathbb{P} V^{i,(l)}(\cdot, \cdot) \\ &\stackrel{(iv)}{\leq} Q^{i^\star}(\cdot, \cdot; \mathbf{w}^i), \end{aligned}$$

where (i) holds because θ^\star is the minimizer; (ii) is by definition of the linear function approximation of the cost-to-go function. (iii) is because $(1-q)$ is a positive fraction. The last inequality uses the definition of $Q^{i^\star}(\cdot, \cdot; \mathbf{w}^i)$ and the induction hypothesis on l , $V^{i,(l)}(\cdot)$ is optimistic, i.e., $V^{i,(l)}(\cdot) \leq V^{i^\star}(\cdot; \mathbf{w}^i)$. Therefore, by induction $Q^{i,(l+1)}(\cdot, \cdot)$ is also optimistic for all l , and hence the final output $Q_1^i(\cdot, \cdot)$ and $V_1^i(\cdot)$ are optimistic. This finishes the proof of round 1.

For the outer induction, suppose that the event in Theorem 3 holds for round 1 to $j_i - 1$ for each agent $i \in N$ with high probability. Therefore, $\theta^\star \in \mathcal{C}_k^i \cap \mathcal{B}$ for rounds $k = 1, 2, \dots, j_i - 1$. And in each round $k = 1, 2, \dots, j_i - 1$, the output of the MAEVI algorithms are optimistic, i.e., $Q_k^i(\cdot, \cdot) \leq Q_k^{i^\star}(\cdot, \cdot; \mathbf{w}^i)$ and $V_k^i(\cdot) \leq V_k^{i^\star}(\cdot; \mathbf{w}^i)$, $\forall i \in N$. So, for each agent $i \in N$, we define the following event

$$\mathcal{E}_{j_i-1}^i = \{ \theta^\star \in \mathcal{C}_k^i \cap \mathcal{B}; Q_k^i(\cdot, \cdot) \leq Q_k^{i^\star}(\cdot, \cdot; \mathbf{w}^i), V_k^i(\cdot) \leq V_k^{i^\star}(\cdot; \mathbf{w}^i) \text{ for each } k = 1, 2, \dots, j_i - 1 \}.$$

and assume that the $\mathbb{P}(\mathcal{E}_{j_i-1}^i) \geq 1 - \frac{\delta'}{n}$ for some $\delta' > 0$. We now show that the event $\mathcal{E}_{j_i}^i$ also holds with high probability. To do this, we will construct the auxiliary sequence of functions for each agent $i \in N$ as follows:

$$\tilde{V}_k^i(\cdot) := \min\{V_k^i(\cdot), B_\star\}, \forall k = 1, 2, \dots, j_i - 1.$$

Moreover, for any $k = 1, 2, \dots, j_i$ and for any $t \in [t_{k-1}^i + 1, t_k^i]$, define the following:

$$\begin{aligned} \tilde{\eta}_t^i &= V_{k-1}^i(\mathbf{s}_{t+1}) - \langle \phi_{\tilde{V}_{k-1}^i}(\mathbf{s}_t, \mathbf{a}_t), \theta^\star \rangle \\ \tilde{\Sigma}_t^i &= \lambda \mathbf{I} + \sum_{r=1}^t \phi_{\tilde{V}_{k(r)-1}^i}(\mathbf{s}_r, \mathbf{a}_r) \phi_{\tilde{V}_{k(r)-1}^i}(\mathbf{s}_r, \mathbf{a}_r)^\top \\ \tilde{\theta}_t^i &= \tilde{\Sigma}_t^{i-1} \sum_{r=1}^t \phi_{\tilde{V}_{k(r)-1}^i}(\mathbf{s}_r, \mathbf{a}_r) \tilde{V}_{k(r)-1}^i(\mathbf{s}_{r+1}) \\ \tilde{\mathcal{C}}_k^i &= \{ \theta \in \mathbb{R}^{nd} : \left\| \tilde{\Sigma}_{t_k^i}^{i/2} (\tilde{\theta}_{t_k^i}^i - \theta) \right\|_2 \leq \beta_{t_k^i} \}, \end{aligned}$$

where $k(r)$ is the round that contains the time period r , i.e., $r \in [t_{k-1}^i + 1, t_k^i]$.

By construction $\{\tilde{\eta}_t^i\}_{t=1}^{t_{j_i}^i}$ are almost surely B_* sub-Gaussian. Moreover, using the similar computations as above and Abbasi-Yadkori et al. (2011) result (Theorem 7), we can say that event $\tilde{\mathcal{E}}_{j_i}^i$ will hold with high probability where

$$\tilde{\mathcal{E}}_{j_i}^i = \{\boldsymbol{\theta}^* \in \tilde{\mathcal{C}}_{j_i}^i \cap \mathcal{B}; Q_{j_i}^i(\cdot, \cdot) \leq Q^{i*}(\cdot, \cdot; \mathbf{w}^i); V_{j_i}^i(\cdot) \leq V^{i*}(\cdot; \mathbf{w}^i)\},$$

and $Q_{j_i}^i$ is the output of $MAEVI(\tilde{\mathcal{C}}_{j_i}^i, \epsilon_{j_i}^i, \frac{1}{t_{j_i}^i}, \mathbf{w}^i)$.

Moreover, under the event $\mathcal{E}_{j_i-1}^i$, the optimism implies that $\tilde{V}_k^i = V_k^i$ for all $k = 1, 2, \dots, j_i - 1$ and for all $i \in N$. Also, under the event $\mathcal{E}_{j_i-1}^i$, we have $\tilde{\eta}_t^i = \eta_t^i$, $\tilde{\Sigma}_t^i = \Sigma_t^i$, $\tilde{\boldsymbol{\theta}}_t^i = \hat{\boldsymbol{\theta}}_t^i$ for all $t \leq t_{j_i}^i$, and thus $\tilde{\mathcal{C}}_{j_i}^i = \mathcal{C}_{j_i}^i$. So, for each agent $i \in N$, we have

$$\mathcal{E}_{j_i}^i = \mathcal{E}_{j_i-1}^i \cap \tilde{\mathcal{E}}_{j_i}^i, \quad \forall i \in N.$$

Therefore, using the union bound, we have $\mathbb{P}(\mathcal{E}_{j_i}^i) \geq 1 - \frac{\delta'}{n} - \frac{\delta}{n \cdot t_{j_i}^i (t_{j_i}^i + 1)}$. Now, by induction argument and taking the union bound, we have

$$\sum_{j_i=1}^{J^i} \frac{\delta}{n \cdot t_{j_i}^i (t_{j_i}^i + 1)} = \sum_{j_i=1}^{J^i} \frac{\delta}{n} \cdot \left(\frac{1}{t_{j_i}^i} - \frac{1}{(t_{j_i}^i + 1)} \right) \leq \frac{\delta}{n},$$

from here we conclude that with a probability of at least $1 - \frac{\delta}{n}$, the good event holds for all the $j_i \leq J^i$, where J^i is the number of calls to the MAEVI algorithm by agent $i \in N$.

Next, it remains to show that the MAEVI will converge in finite time with ϵ^i tolerances for any agent $i \in N$. To do so, it suffices to show that $\|V^{i,(l)} - V^{i,(l-1)}\|_\infty$ shrinks exponentially. We now claim that $\|Q^{i,(l)} - Q^{i,(l-1)}\|_\infty$ shrinks exponentially, which together with Q^i update in algorithm gives the desired result since $\|V^{i,(l)} - V^{i,(l-1)}\|_\infty \leq \|Q^{i,(l)} - Q^{i,(l-1)}\|_\infty$. To show this, first note that for any (s, \mathbf{a}) pair we have,

$$\begin{aligned} |Q^{i,(l)}(s, \mathbf{a}) - Q^{i,(l-1)}(s, \mathbf{a})| &= (1-q) \cdot \left| \min_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, \phi_{V^{i,(l-1)}}(s, \mathbf{a}) \rangle - \min_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, \phi_{V^{i,(l-2)}}(s, \mathbf{a}) \rangle \right| \\ &= (1-q) \cdot \left| \min_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, \phi_{V^{i,(l-1)}}(s, \mathbf{a}) \rangle + \max_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, -\phi_{V^{i,(l-2)}}(s, \mathbf{a}) \rangle \right| \\ &= (1-q) \cdot \left| -\max_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, -\phi_{V^{i,(l-1)}}(s, \mathbf{a}) \rangle + \max_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, -\phi_{V^{i,(l-2)}}(s, \mathbf{a}) \rangle \right| \\ &\stackrel{(i)}{\leq} (1-q) \cdot \max_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} |-\langle \boldsymbol{\theta}, -\phi_{V^{i,(l-1)}}(s, \mathbf{a}) \rangle + \langle \boldsymbol{\theta}, -\phi_{V^{i,(l-2)}}(s, \mathbf{a}) \rangle| \\ &= (1-q) \cdot \max_{\boldsymbol{\theta} \in \mathcal{C}^i \cap \mathcal{B}} |\langle \boldsymbol{\theta}, \phi_{V^{i,(l-1)}}(s, \mathbf{a}) - \phi_{V^{i,(l-2)}}(s, \mathbf{a}) \rangle| \\ &\stackrel{(ii)}{=} (1-q) \cdot |\langle \bar{\boldsymbol{\theta}}, \phi_{V^{i,(l-1)}}(s, \mathbf{a}) - \phi_{V^{i,(l-2)}}(s, \mathbf{a}) \rangle| \\ &= (1-q) \cdot |\mathbb{P}_{\bar{\boldsymbol{\theta}}}(V^{i,(l-1)} - V^{i,(l-2)})(s, \mathbf{a})| \\ &\stackrel{(iii)}{\leq} (1-q) \cdot \max_{s' \in \mathcal{S}} |V^{i,(l-1)}(s') - V^{i,(l-2)}(s')| \\ &= (1-q) \cdot \max_{s' \in \mathcal{S}} \left| \min_{\mathbf{a}'} Q^{i,(l-1)}(s', \mathbf{a}') - \min_{\mathbf{a}'} Q^{i,(l-2)}(s', \mathbf{a}') \right| \\ &\stackrel{(iv)}{\leq} (1-q) \cdot \max_{s' \in \mathcal{S}, \mathbf{a}' \in \mathcal{A}} |Q^{i,(l-1)}(s', \mathbf{a}') - Q^{i,(l-2)}(s', \mathbf{a}')| \\ &= (1-q) \cdot \|Q^{i,(l-1)} - Q^{i,(l-2)}\|_\infty, \end{aligned}$$

where (i) holds because max is a contraction. (ii) holds because $\bar{\boldsymbol{\theta}}$ is the $\boldsymbol{\theta}$ in the non-empty set $\mathcal{C}^i \cap \mathcal{B}$ that achieves the maximum. Since, $\mathbb{P}_{\bar{\boldsymbol{\theta}}}(\cdot | s, \mathbf{a})$ is a probability distribution we have (iii). Finally, (iv) holds again because of the contraction property of the maximum function. Since (s, \mathbf{a}) are arbitrary in the above, we conclude that $\|Q^{i,(l)} - Q^{i,(l-1)}\|_\infty \leq (1-q) \|Q^{i,(l-1)} - Q^{i,(l-2)}\|_\infty$. Applying this recursively, we have a term that is exponentially decaying and hence $\|Q^{i,(l)} - Q^{i,(l-1)}\|_\infty$ shrinks exponentially implying $\|V^{i,(l)} - V^{i,(l-1)}\|_\infty$ also shrinks exponentially. \square

This concludes the MAEVI analysis. Next, we show that the total number of calls, J to the MAEVI algorithm in the entire analysis is bounded. Note that an agent $i \in N$ calls the MAEVI algorithm if at least one of the doubling criteria is satisfied, i.e., either the determinant doubling criteria or the time doubling criteria is satisfied. Let J^i be the total number of calls to the MAEVI algorithm made by agent i . We will show that J^i is bounded. For agent i , let J_1^i be the number of calls to MAEVI made via determinant doubling criteria, and J_2^i be the calls to the MAEVI algorithm via the time doubling criteria. Note that $J^i \leq J_1^i + J_2^i$. The inequality is taken because there are cases when both the doubling criteria are satisfied. In particular, we use Lemma 1 that ensures the number of calls, J to the MAEVI algorithm is bounded.

A.5 Proof of Lemma 1

Recall the lemma: The total number of calls to the MAEVI algorithm in the entire analysis J , is bounded as

$$J \leq 2n^2d \log \left(1 + \frac{TB_*^2nd}{\lambda} \right) + 2n \log(T).$$

Proof. To prove this result, let us consider an agent $i \in N$. We bound the number of calls to MAEVI algorithm J^i , for each agent $i \in N$ and use the fact that $J \leq \sum_{i \in N} J^i$.

Consider agent $i \in N$, since $V_0^i \leq B_*$. It is easy to see that

$$\begin{aligned} \|\Sigma_T^i\|_2 &= \left\| \lambda \mathbf{I} + \sum_{j=0}^{J^i} \sum_{t=t_j^i+1}^{t_{j+1}^i} \phi_{V_j^i}(\mathbf{s}_t, \mathbf{a}_t) \phi_{V_j^i}(\mathbf{s}_t, \mathbf{a}_t)^\top \right\| \\ &\leq \lambda + \sum_{j=0}^{J^i} \sum_{t=t_j^i+1}^{t_{j+1}^i} \|\phi_{V_j^i}(\mathbf{s}_t, \mathbf{a}_t)\|^2 \\ &\leq \lambda + TB_*^2nd. \end{aligned}$$

The first inequality is because of the triangle inequality. The second one uses the fact that $\|\phi_V\| \leq B_*\sqrt{nd}$ and $V_j^i \leq B_*$ for all $j \geq 0$ (from Assumption 2).

For the determinant doubling criteria, we have $\det(\Sigma_T) \leq (\lambda + TB_*^2nd)^{nd}$. This implies

$$(\lambda + TB_*^2nd)^{nd} \geq 2^{J_1^i} \cdot \det(\Sigma_0) = 2^{J_1^i} \cdot \lambda^{nd},$$

taking log on both sides, we have

$$J_1^i \leq nd \log_2 \left(1 + \frac{TB_*^2nd}{\lambda} \right) \leq 2nd \log \left(1 + \frac{TB_*^2nd}{\lambda} \right).$$

This bounds the number of calls to the MAEVI algorithm when the determinant doubling criteria is satisfied. Next, we consider the number of calls to the MAEVI algorithm when the time doubling criteria is satisfied, i.e., we will bound J_2^i . Note that $t_0 = 1$, so we have $2^{J_2^i} \leq T$, this implies $J_2^i \leq \log_2(T) \leq 2 \log(T)$. Now, summing up the bounds for J_1^i and J_2^i we have the bounds for J^i , i.e.,

$$J^i \leq J_1^i + J_2^i \leq 2nd \log \left(1 + \frac{TB_*^2nd}{\lambda} \right) + 2 \log(T).$$

This proves that the number of calls to the MAEVI algorithm made by an agent $i \in N$ is bounded. So, the total number of calls to MAEVI in the entire analysis is bounded as

$$J \leq \sum_{i \in N} J^i = \sum_{i \in N} 2nd \log \left(1 + \frac{TB_*^2nd}{\lambda} \right) + 2 \log(T) = 2n^2d \log \left(1 + \frac{TB_*^2nd}{\lambda} \right) + 2n \log(T).$$

This ends the proof. \square

The next step is to decompose the regret and bound each term. To this end, we prove Theorem 4. First, we divide the time horizon into disjoint intervals $m = 1, 2, \dots, M$, where the endpoint of each interval is decided by one of the two conditions: 1) at least for one agent, the MAEVI is triggered, and 2) all the agents have reached the goal state. Note that this decomposition is explicitly used in the regret analysis only and not in the actual algorithm implementation. Let H_m denote the length of the interval m . Moreover, at the end of the M -th interval, all the K episodes are over. Therefore, we have the total length of all the intervals as $\sum_{m=1}^M H_m$, which is the same as $\sum_{k=1}^K T_k$, where T_k is the time to finish the episode k . Within an interval, the optimistic estimator of the state-value function for each agent remains the same. An agent for whom the MAEVI is not triggered in the m -th interval will continue using the same optimistic estimator of the state-action value function in the $(m+1)$ -th interval. Later we separate the set of agents depending on whether the MAEVI is triggered for them or not. Let S_m be the set of agents for whom the MAEVI is triggered in the m -th interval. Using the above interval decomposition, the regret R_K can be written as

$$R_K = R(M) \leq \sum_{m=1}^M \sum_{h=1}^{H_m} \frac{1}{n} \sum_{i=1}^n \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + 1 - \sum_{m \in \mathcal{M}(M)} \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{init}), \quad (26)$$

here $\mathcal{M}(M)$ is the set of all intervals that are the first intervals in each episode. In RHS we add 1 because $|V_0^i| \leq 1$. Recall, the above equation is the same as Equation (11) in the main paper.

A.6 Proof of Theorem 4

Recall the theorem: Assume that the event in Theorem 3 holds, then we have the following upper bound on the regret given in Equation (26),

$$\begin{aligned} R(M) &\leq \underbrace{\sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i=1}^n \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h}) \right]}_{E_1} \\ &+ \underbrace{\sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h+1}) - \frac{1}{n} \sum_{i=1}^n \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right]}_{E_2} \\ &+ 2n^2 dB_* \log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2nB_* \log(T) + 2. \end{aligned} \quad (27)$$

Proof. The proof relies on the following proposition that is key result for the regret decomposition.

Proposition 4. *Conditioned on the event given in Theorem 3, for the above mentioned interval decomposition, we have the following:*

$$\begin{aligned} &\sum_{m=1}^M \left(\sum_{h=1}^{H_m} \left\{ \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h+1}) \right\} \right) - \sum_{m \in \mathcal{M}(M)} \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{init}) \\ &\leq 1 + 2n^2 dB_* \log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2nB_* \log(T). \end{aligned}$$

The proof this proposition is given in Section C.2 of this appendix. Using this Proposition in the regret expression given in Equation (26), we get the regret decomposition as desired. \square

In the next Section, we provide the details of feature design of the transition probabilities and the details of the optimal policy value used in the computations.

B DETAILS OF THE FEATURE DESIGN AND THE OPTIMAL POLICY FOR COMPUTATIONAL EXPERIMENTS

In this section, we will provide the details of the optimal policy and related results. Recall, for each $(s', \mathbf{a}, s) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$, the global transition probability is parameterized as $\mathbb{P}_\theta(s'|s, \mathbf{a}) = \langle \phi(s'|s, \mathbf{a}), \boldsymbol{\theta} \rangle$. The features $\phi(s'|s, \mathbf{a})$ are

$$\phi(s'|s, \mathbf{a}) = \begin{cases} (\phi(s'^1|s^1, \mathbf{a}^1), \dots, \phi(s'^n|s^n, \mathbf{a}^n)), & \text{if } s \neq \mathbf{g}, \\ \mathbf{0}_{nd}, & \text{if } s = \mathbf{g}, s' \neq \mathbf{g}, \\ (\mathbf{0}_{nd-1}, 2^{n-1}), & \text{if } s = \mathbf{g}, s' = \mathbf{g}, \end{cases}$$

where $\phi(s'^i|s^i, \mathbf{a}^i)$ are defined as

$$\phi(s'^i|s^i, \mathbf{a}^i) = \begin{cases} (-\mathbf{a}^i, \frac{1-\delta}{n})^\top, & \text{if } s^i = s'^i = s_{init} \\ (\mathbf{a}^i, \frac{\delta}{n})^\top, & \text{if } s^i = s_{init}, s'^i = \mathbf{g} \\ \mathbf{0}_d^\top, & \text{if } s^i = \mathbf{g}, s'^i = s_{init} \\ (\mathbf{0}_{d-1}, \frac{1}{n})^\top, & \text{if } s^i = s'^i = \mathbf{g}. \end{cases}$$

Here $\mathbf{0}_d^\top = (0, 0, \dots, 0)^\top$ is a vector of d dimension with all zeros. Thus, the features $\phi(s'^i|s^i, \mathbf{a}^i) \in \mathbb{R}^{nd}$. Moreover, the transition probability parameters are taken as $\boldsymbol{\theta} = (\boldsymbol{\theta}^1, \frac{1}{2^{n-1}}, \boldsymbol{\theta}^2, \frac{1}{2^{n-1}}, \dots, \boldsymbol{\theta}^n, \frac{1}{2^{n-1}})$, where $\boldsymbol{\theta}^i \in \left\{ -\frac{\Delta}{n(d-1)}, \frac{\Delta}{n(d-1)} \right\}^{d-1}$, and $\Delta < \delta$.

We first proof Lemma 2 to show that these transition probability function features satisfy some basic properties.

B.1 Proof of Lemma 2

Recall the lemma: The features $\phi(s'|s, \mathbf{a})$ satisfy the following: (a) $\sum_{s'} \langle \phi(s'|s, \mathbf{a}), \boldsymbol{\theta} \rangle = 1, \forall s, \mathbf{a}$; (b) $\langle \phi(s' = \mathbf{g}|s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle = 1, \forall \mathbf{a}$; (c) $\langle \phi(s' \neq \mathbf{g}|s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle = 0, \forall \mathbf{a}$.

Proof. To prove this lemma, we consider two cases. In case 1, $s \neq \mathbf{g}$ and case 2, $s = \mathbf{g}$.

Case 01: ($s \neq \mathbf{g}$). Without loss of generality we consider the following state $\mathbf{s} = (\underbrace{s_{init}, s_{init}, \dots, s_{init}}_{k \text{ times}}, \underbrace{\mathbf{g}, \mathbf{g}, \dots, \mathbf{g}}_{(n-k) \text{ times}})$, i.e.,

$k \neq 0$ agents are at s_{init} and remaining $n - k$ are at \mathbf{g} . Consider an agent i , who is at s_{init} node. Out of total 2^n next possible states, there are exactly $\frac{2^n}{2}$ states in which agent i will remain at s_{init} , and in $\frac{2^n}{2}$ states in which the agents move to goal node. The probability that the next node of agent i is s_{init} given that the current node of agent i is s_{init} is given by $-\langle \mathbf{a}^i, \boldsymbol{\theta}^i \rangle + \frac{1-\delta}{n} \times \frac{1}{2^{n-1}}$. And the probability that the next node of agent i is \mathbf{g} given that the current node of agent i is s_{init} is $\langle \mathbf{a}^i, \boldsymbol{\theta}^i \rangle + \frac{\delta}{n} \times \frac{1}{2^{n-1}}$. These probabilities are obtained using the features defined for our two agent model. Since, this is true for all the agents $1, 2, \dots, k$ which are at s_{init} . So, the contribution to the probability term from these k agents who are at s_{init} is

$$\sum_{i=1}^k \left\{ \left(-\langle \mathbf{a}^i, \boldsymbol{\theta}^i \rangle + \frac{1-\delta}{n} \times \frac{1}{2^{n-1}} \right) \times \frac{2^n}{2} \right\} + \sum_{i=1}^k \left\{ \left(\langle \mathbf{a}^i, \boldsymbol{\theta}^i \rangle + \frac{\delta}{n} \times \frac{1}{2^{n-1}} \right) \times \frac{2^n}{2} \right\} = \frac{k}{n}.$$

Next, consider an agent whose state is \mathbf{g} ; for this agent, there are two possibilities, stay at \mathbf{g} or go to s_{init} . The probability of going to s_{init} is zero, however if the agent stays at \mathbf{g} then the probability corresponding to that is $\frac{1}{n} \times \frac{1}{2^{n-1}}$. So the total probability of going to the next state for the agent whose state is \mathbf{g} is

$$0 \times \frac{2^n}{2} + \left(\frac{1}{n} \times \frac{1}{2^{n-1}} \right) \frac{2^n}{2}.$$

Again this is true for the agents $k+1, k+2, \dots, n$ who are at \mathbf{g} , so the overall probability is

$$\sum_{i=k+1}^n \left\{ \left(0 \times \frac{2^n}{2} + \left(\frac{1}{n} \times \frac{1}{2^{n-1}} \right) \frac{2^n}{2} \right) \right\} = \frac{n-k}{n}.$$

So, the sum of these linearly approximated probabilities is

$$\sum_{s'} \langle \phi(s'|s, \mathbf{a}), \boldsymbol{\theta} \rangle = \frac{k}{n} + \frac{n-k}{n} = 1.$$

This ends the proof of the first case.

Case 02: ($s = \mathbf{g}$). For this case, the probability is

$$\begin{aligned} \sum_{s'} \langle \phi(s'|s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle &= \sum_{s' \neq \mathbf{g}} \langle \phi(s'|s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle + \langle \phi(s' = \mathbf{g}|s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle \\ &= \langle \mathbf{0}, \boldsymbol{\theta} \rangle + \langle (\mathbf{0}_{nd-1}, 2^{n-1}), \boldsymbol{\theta} \rangle = 1. \end{aligned}$$

Therefore, in both cases, we have

$$\sum_{s'} \langle \phi(s'|s = \mathbf{g}, \mathbf{a}), \boldsymbol{\theta} \rangle = 1, \quad \forall s, \mathbf{a}.$$

The other two statements of the lemma follow by feature design and model parameter space. \square

Next, we provide the details of the value of the optimal policy for the above model. We require it in the regret computations. First, note that the value of the optimal policy can be defined in terms of expected cost if all agents reach the goal state exactly in a 1 time period, in a 2 time period, and so on. Suppose the agents reach to the goal state exactly by t steps such that x_j agents move to the goal node at each step $j = 1, 2, \dots, t-1$ and remaining $x_t = n - \sum_{j=1}^{t-1} x_j$ agents move to the goal node by t -th step. We call this sequence of departures of agents to the goal node as the ‘departure sequence’. For the above departure sequence, the cost incurred is

$$\begin{aligned} C_\alpha(x_1, \dots, x_t) &= \alpha \times \sum_{j=1}^{t-1} x_j^2 + \sum_{j=1}^{t-1} \left(n - \sum_{i=1}^j x_i \right) \cdot c_{\min} + \alpha \times x_t^2 \\ &= \alpha \times \sum_{j=1}^{t-1} x_j^2 + \sum_{j=1}^{t-1} \left(n - \sum_{i=1}^j x_i \right) \cdot c_{\min} + \alpha \times \left(n - \sum_{j=1}^{t-1} x_j \right)^2. \end{aligned} \quad (28)$$

where α is the mean of the uniform distribution $\mathcal{U}(c_{\min}, 1)$, i.e., $\alpha = \frac{c_{\min}+1}{2}$. We use this α instead of the private cost to compute the optimal value. The first term in the above equation is because all x_j agents have moved to the goal node in time period j ; hence the congestion is x_j . Moreover, each agent incurs the private cost α in this period. So, the cost incurred to any agent is αx_j , and the number of agents moved are x_j , and hence the total cost incurred is $(\alpha x_j) \times x_j = \alpha x_j^2$. This happens for all the time periods $j = 1, 2, \dots, t-1$, so we sum this for $(t-1)$ periods. The remaining agents at each time period incurred a waiting cost of c_{\min} ; thus, we have a second term. Finally, the third term is because the remaining agents, x_t will move to the goal node at the last time period t . The optimal value of a policy π^* , $V^* := V^{\pi^*}$, can be written as

$$V^* = \sum_{t=1}^{\infty} \mathbb{P}[x_1^*, \dots, x_t^*] \cdot C_\alpha(x_1^*, \dots, x_t^*), \quad (29)$$

where $x_1^*, x_2^*, \dots, x_{t-1}^*, x_t^*$ is the optimal departure sequence. These x_j^* are obtained by minimizing the cost of function $C_\alpha(x_1, \dots, x_t)$. Moreover, $\mathbb{P}[x_1^*, \dots, x_t^*]$ is the probability of occurrence of this optimal departure sequence. In Theorem 5 we provide the optimal departure sequence x_1^*, \dots, x_t^* , and the corresponding value $C_\alpha(x_1^*, \dots, x_t^*)$.

B.2 Proof of Theorem 5

Recall the theorem: The optimal departure sequence is given by

$$x_j^* = \left\lfloor \frac{n}{t} + \left(\frac{t+1}{2} - j \right) \cdot \frac{c_{\min}}{2\alpha} \right\rfloor, \quad \forall j = 1, 2, \dots, t-1; \quad x_t^* = n - \sum_{j=1}^{t-1} x_j^*.$$

Moreover, the optimal cost $C_\alpha(x_1^*, \dots, x_t^*)$ of using the above optimal departure sequence is

$$\alpha t \left(\frac{n}{t} \right)^2 + n(t-1) \cdot \frac{c_{\min}}{2\alpha} - \frac{t(t-1)(t+1)}{12} \cdot \frac{c_{\min}^2}{4\alpha^2}.$$

Proof. First, recall for the departure sequence x_1, \dots, x_t , the cost function is given by

$$\begin{aligned} C_\alpha(x_1, \dots, x_t) &= \alpha \times \sum_{j=1}^{t-1} x_j^2 + \sum_{j=1}^{t-1} \left(n - \sum_{i=1}^j x_i \right) \cdot c_{\min} + \alpha \times x_t^2 \\ &= \alpha \times \sum_{j=1}^{t-1} x_j^2 + \sum_{j=1}^{t-1} \left(n - \sum_{i=1}^j x_i \right) \cdot c_{\min} + \alpha \times \left(n - \sum_{j=1}^{t-1} x_j \right)^2. \end{aligned}$$

The proof of this theorem follows by taking the partial derivative of the cost function with respect to the x_j 's and equating them to zero. The first order conditions are necessary and sufficient because the Hessian of the above cost function is positive definite (shown below), and hence the minima exist. The partial derivative of the cost function with respect to x_j is given by

$$\frac{\partial C}{\partial x_j} = 2\alpha x_j - 2\alpha \left(n - \sum_{i=1}^{t-1} x_i \right) - (t-j)c_{\min} = 0, \quad \forall j = 1, 2, \dots, t-1. \quad (30)$$

From the above equations, we have

$$x_j - x_{j+1} = \frac{c_{\min}}{2\alpha}, \quad \forall j = 1, 2, \dots, t-2.$$

Converting all the variables x_2, x_3, \dots, x_{t-1} in terms of x_1 , we have

$$x_j = x_1 - (j-1) \cdot \frac{c_{\min}}{2\alpha}, \quad \forall j = 2, 3, \dots, t-1. \quad (31)$$

Also, from Equation (30), for $j = 1$, we have

$$\frac{\partial C}{\partial x_1} = 2\alpha x_1 - 2\alpha \left(n - \sum_{i=1}^{t-1} x_i \right) - (t-1)c_{\min} = 0.$$

Substituting the variables x_2, x_3, \dots, x_{t-1} in terms of x_1 in above equation, and solving for x_1 using Equation (31), we have

$$2\alpha x_1 - 2\alpha \left(n - \sum_{i=1}^{t-1} x_i \right) - (t-1)c_{\min} = 2\alpha x_1 - 2\alpha \left(n - \sum_{i=1}^{t-1} \left(x_1 - (i-1) \cdot \frac{c_{\min}}{2\alpha} \right) \right) - (t-1)c_{\min} = 0.$$

Solving for x_1 , we have

$$x_1^* = \frac{n}{t} + \frac{t-1}{2} \cdot \frac{c_{\min}}{2\alpha}. \quad (32)$$

Putting them in Equation (31), we have

$$x_j^* = \frac{n}{t} + \left(\frac{t-1}{2} - (j-1) \right) \cdot \frac{c_{\min}}{2\alpha} = \left[\frac{n}{t} + \left(\frac{t+1}{2} - j \right) \cdot \frac{c_{\min}}{2\alpha} \right], \quad \forall j = 1, 2, \dots, t-1.$$

$$x_t^* = n - \sum_{j=1}^{t-1} x_j^*$$

Next, we show that the Hessian of the cost function is positive definite. It is easy to see that the Hessian is given by

$$\nabla^2 C_\alpha(x_1, \dots, x_{t-1}) = \begin{bmatrix} 4\alpha & 2\alpha & 2\alpha & \dots & 2\alpha \\ 2\alpha & 4\alpha & 2\alpha & \dots & 2\alpha \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 2\alpha & 2\alpha & 2\alpha & \dots & 4\alpha \end{bmatrix},$$

i.e., all the diagonal elements are 4α , and all off-diagonal elements are 2α . For such a matrix with the $n \times n$ order one eigenvalue is $2(n+1)\alpha$, and the remaining eigenvalues are 2α . So, all the eigenvalues are positive; hence Hessian is positive definite. So, the function is convex.

To complete the proof, we put back all x_j^* 's in the cost function and simplify it further

$$C(x_j^*; j = 1, 2, \dots, t-1) = \alpha t \left(\frac{n}{t} \right)^2 + n(t-1) \cdot \frac{c_{\min}}{2\alpha} - \frac{t(t-1)(t+1)}{12} \cdot \frac{c_{\min}^2}{4\alpha^2}.$$

This ends the proof. \square

Apart from the cost of the optimal departure sequence, we also require the probability of this departure sequence to compute V^* . To this end, we recall the feature design of transition probability that allows an agent to stay or depart from the initial node s_{init} to the goal node g iff the sign of the action \mathbf{a}^i taken by agent i matches the sign of the transition probability function parameter $\boldsymbol{\theta}^i$, i.e. $\text{sgn}(\mathbf{a}_j^i) = \text{sgn}(\boldsymbol{\theta}_j^i)$ for all $j = 1, 2, \dots, d-1$. This is because for each agent $i \in N$, we have $\langle \phi(s'^i | s^i, \mathbf{a}^i), (\boldsymbol{\theta}^i, \frac{1}{2^{n-1}}) \rangle$ as follows:

$$\left\langle \phi(s'^i | s^i, \mathbf{a}^i), \left(\boldsymbol{\theta}^i, \frac{1}{2^{n-1}} \right) \right\rangle = \begin{cases} -\langle \mathbf{a}^i, \boldsymbol{\theta}^i \rangle + \frac{1-\delta}{n} \times \frac{1}{2^{n-1}}, & \text{if } s^i = s'^i = s_{init} \\ \langle \mathbf{a}^i, \boldsymbol{\theta}^i \rangle + \frac{\delta}{n} \times \frac{1}{2^{n-1}}, & \text{if } s^i = s_{init}, s'^i = g \\ 0, & \text{if } s^i = g, s'^i = s_{init} \\ \frac{1}{n} \times \frac{1}{2^{n-1}} & \text{if } s^i = s'^i = g. \end{cases} \quad (33)$$

Hence the transition probability is as follows

$$\mathbb{P}(s' | s, \mathbf{a}) = \sum_{i=1}^n \left\langle \phi(s'^i | s^i, \mathbf{a}^i), \left(\boldsymbol{\theta}^i, \frac{1}{2^{n-1}} \right) \right\rangle. \quad (34)$$

Using the sign matching property, we have the following theorem for the transition probability.

B.3 Proof of Theorem 6

Recall the theorem: The transition probability $\mathbb{P}[x_1 = x_1^*, \dots, x_t = x_t^*]$ is given by

$$\mathbb{P}[x_1 = x_1^*, \dots, x_t = x_t^*] = \prod_{k=1}^{t-1} \left(1 - \left(\gamma n - (\gamma - \eta) \sum_{j=1}^k x_j^* \right) \right) \times \left(\gamma n - (\gamma - \eta) \sum_{j=1}^{t-1} x_j^* \right),$$

where $\gamma = \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right)$ and $\eta = \frac{1}{n \cdot 2^{n-1}}$.

Proof. To find the probability of optimal departure sequence, we first find the probability of all agents reaching the goal state g exactly by t time period starting from initial state s_{init} . Formally, we find the following probability

$$\mathbb{P}[x_1 = x_1^*, \dots, x_{t-1} = x_{t-1}^*, x_t = x_t^*] = \mathbb{P}(s_{t+1} = g, s_t \neq g, s_{t-1} \neq g, \dots, s_2 \neq g | s_1 = s_{init}, \mathbf{a}_1).$$

This is because, the agents will reach to the goal state in exactly t time periods while using $x_1^*, \dots, x_{t-1}^*, x_t^*$ is the optimal departure sequence. So, t time periods will end if $s_{t+1} = g$.

The above probability can be written as

$$\begin{aligned} \mathbb{P}(s_{t+1} = g, s_t \neq g, \dots, s_2 \neq g | s_1 = s_{init}, \mathbf{a}_1) &= \mathbb{P}(s_2 \neq g | s_1 = s_{init}, \mathbf{a}_1) \times \mathbb{P}(s_3 \neq g | s_1 = s_{init}, s_2, \mathbf{a}_1, \mathbf{a}_2) \\ &\times \dots \times \mathbb{P}(s_t \neq g | s_1 = s_{init}, s_2, \dots, s_{t-1}, \mathbf{a}_1, \dots, \mathbf{a}_{t-1}) \times \mathbb{P}(s_{t+1} = g | s_1 = s_{init}, s_2, \dots, s_t, \mathbf{a}_1, \dots, \mathbf{a}_t). \end{aligned}$$

This can be written as follows

$$\begin{aligned} \mathbb{P}(s_{t+1} = g, s_t \neq g, \dots, s_2 \neq g | s_1 = s_{init}, \mathbf{a}_1) &= \prod_{k=1}^{t-1} \mathbb{P}(s_{k+1} \neq g | s_1 = s_{init}, s_2, \dots, s_k, \mathbf{a}_1, \dots, \mathbf{a}_k) \\ &\times \mathbb{P}(s_{t+1} = g | s_1 = s_{init}, s_2, \dots, s_t, \mathbf{a}_1, \dots, \mathbf{a}_t) \\ &= \prod_{k=1}^{t-1} \left(1 - \mathbb{P}(s_{k+1} = g | s_1 = s_{init}, s_2, \dots, s_k, \mathbf{a}_1, \dots, \mathbf{a}_k) \right) \\ &\times \mathbb{P}(s_{t+1} = g | s_1 = s_{init}, s_2, \dots, s_t, \mathbf{a}_1, \dots, \mathbf{a}_t). \end{aligned} \quad (35)$$

To find this probability first consider the following probability for any $k = 1, 2, \dots, t-1$

$$\mathbb{P}(s_{k+1} = g | s_1 = s_{init}, s_2, \dots, s_k, \mathbf{a}_1, \dots, \mathbf{a}_k) = \sum_{\{i: s_k^i = s_{init}\}} \left(\langle \mathbf{a}_k^i, \boldsymbol{\theta}_k^i \rangle + \frac{\delta}{n \cdot 2^{n-1}} \right) + \sum_{\{i: s_k^i = g\}} \left(\frac{1}{n \cdot 2^{n-1}} \right)$$

$$= \sum_{\{i:s_k^i=s_{init}\}} \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right) + \sum_{\{i:s_k^i=g\}} \left(\frac{1}{n \cdot 2^{n-1}} \right). \quad (36)$$

The above equation uses the definition of transition probability, depending on the number of agents that have moved from the initial node to the goal node and the number of agents already at the goal node. Moreover, we also note that the agent i will move to goal node from the initial node with highest probability if the sign of \mathbf{a}_k^i matches the sign of $\boldsymbol{\theta}_k^i$ for each component and for all $k = 1, 2, \dots, t-1$. Thus, $\langle \mathbf{a}_k^i, \boldsymbol{\theta}_k^i \rangle = \frac{\Delta}{n}$, for all $k = 1, 2, \dots, t-1$.

Substituting Eq. (36) in Eq. (35), we have the following

$$\begin{aligned} & \mathbb{P}(s_{t+1} = \mathbf{g}, s_t \neq \mathbf{g}, \dots, s_2 \neq \mathbf{g} | s_1 = s_{init}, \mathbf{a}_1) \\ &= \prod_{k=1}^{t-1} (1 - \mathbb{P}(s_{k+1} = \mathbf{g} | s_1 = s_{init}, s_2, \dots, s_k, \mathbf{a}_1, \dots, \mathbf{a}_k)) \\ & \quad \times \mathbb{P}(s_{t+1} = \mathbf{g} | s_1 = s_{init}, s_2, \dots, s_t, \mathbf{a}_1, \dots, \mathbf{a}_t) \\ &= \prod_{k=1}^{t-1} \left(1 - \left\{ \sum_{\{i:s_k^i=s_{init}\}} \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right) + \sum_{\{i:s_k^i=g\}} \left(\frac{1}{n \cdot 2^{n-1}} \right) \right\} \right) \\ & \quad \times \sum_{\{i:s_{t-1}^i=s_{init}\}} \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right) + \sum_{\{i:s_{t-1}^i=g\}} \left(\frac{1}{n \cdot 2^{n-1}} \right) \\ &= \prod_{k=1}^{t-1} \left(1 - \left\{ \left(n - \sum_{j=1}^k x_j \right) \cdot \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right) + \left(\sum_{j=1}^k x_j \right) \cdot \left(\frac{1}{n \cdot 2^{n-1}} \right) \right\} \right) \\ & \quad \times \left\{ \left(n - \sum_{j=1}^{t-1} x_j \right) \cdot \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right) + \left(\sum_{j=1}^k x_j \right) \cdot \left(\frac{1}{n \cdot 2^{n-1}} \right) \right\}, \end{aligned}$$

where in the last inequality, we use the fact that the departure sequence is x_1, x_2, \dots, x_{t-1} .

The above probability is the same as the probability of optimal departure sequence x_j^* , i.e.,

$$\mathbb{P}[x_1 = x_1^*, \dots, x_{t-1} = x_{t-1}^*, x_t = x_t^*] = \prod_{k=1}^{t-1} \left(1 - \left(\gamma n - (\gamma - \eta) \sum_{j=1}^k x_j^* \right) \right) \times \left(\gamma n - (\gamma - \eta) \sum_{j=1}^{t-1} x_j^* \right),$$

where $\gamma = \left(\frac{\Delta}{n} + \frac{\delta}{n \cdot 2^{n-1}} \right)$ and $\eta = \frac{1}{n \cdot 2^{n-1}}$. □

Since it is hard to get the closed form expression of the above optimal value we obtain an approximation of the optimal value function for regret computation. So, we define the approximate optimal value in terms of a given large $T < \infty$, as

$$V_T^* = \sum_{t=1}^T \mathbb{P}[x_1^*, \dots, x_t^*] \cdot C_\alpha(x_1^*, \dots, x_t^*).$$

We use this approximate value in the computations, where T is tuned suitably.

C PROOF OF THE INTERMEDIATE LEMMAS AND PROPOSITIONS

In this section, we will provide the proof details of Lemmas and Propositions stated in this appendix.

C.1 Proof of Lemma 3

Proof. The proof of this lemma involves three major steps: 1) the convergence of the cost function parameters (Theorem 2); 2) convergence of the optimistic estimator of the state-action value function whenever MAEVI is triggered (Theorem 3); 3) the regret given in Equation (19) will be decomposed into two components (Theorem 4), and each component will be bounded (Propositions 5, 6). This decomposition further depends on the agents for whom the MAEVI is triggered or

not. The decomposition of regret in this manner is novel to our multi-agent congestion cost minimization. To best of our knowledge this is not available in literature.

To complete the proof, we need to bound both E_1 and E_2 . The following proposition gives the bound to E_1 , which uses *all the intrinsic properties* of our algorithm design.

Proposition 5. *The term E_1 is upper bounded as follows*

$$E_1 \leq 8\beta_T \sqrt{2Td \log \left(1 + \frac{TB_\star^2}{\lambda} \right)} + 8nd(B_\star + 1) \left[\log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + 2 \log(T) \right] + 4(n+1). \quad (37)$$

Proof of this proposition is available in Section C.3 of this appendix. Next, we bound the other term E_2 , which uses the Azuma-Hoeffding inequality. The following proposition provides the bound to E_2 .

Proposition 6. *The term E_2 is upper bounded as follows*

$$E_2 = \frac{1}{n} \sum_{i=1}^n 2B_\star \sqrt{2T \log \left(\frac{2nT}{\delta} \right)} = 2B_\star \sqrt{2T \log \left(\frac{2nT}{\delta} \right)}. \quad (38)$$

The proof of this lemma is present in Section C.4. The proof of Lemma 3 follows by substituting bounds of E_1 and E_2 from Equations (37) and (38) in Theorem 4,

$$\begin{aligned} R(M) &\leq 8\beta_T \sqrt{2Td \log \left(1 + \frac{TB_\star^2}{\lambda} \right)} + 8nd(B_\star + 1) \left[\log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + 2 \log(T) \right] + 4(n+1) \\ &\quad + 2B_\star \sqrt{2T \log \left(\frac{2nT}{\delta} \right)} + 2n^2 d B_\star \log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + 2nB_\star \log(T) + 2. \end{aligned}$$

Combining the lower order terms, we have

$$R_K \leq 10\beta_T \sqrt{Td \log \left(1 + \frac{TB_\star^2}{\lambda} \right)} + 16n^2 d B_\star \log \left(T + \frac{T^2 B_\star^2 nd}{\lambda} \right).$$

This ends the proof of Lemma 3. □

C.2 Proof of Proposition 4

Proof. The proof uses the ideas for the single agent SSP in Tarbouriech et al. (2020) and Min et al. (2022); however, for the multi-agent setup, we have some extra challenges. Consider the following

$$\begin{aligned} &\sum_{m=1}^M \left(\sum_{h=1}^{H_m} \left\{ \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h+1}) \right\} \right) \\ \stackrel{(i)}{=} &\sum_{m=1}^M \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_m+1}) \right) \\ \stackrel{(ii)}{=} &\sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_m+1}) \right) \\ &+ \sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) \right) \\ &+ \frac{1}{n} \sum_{i=1}^n V_{j_i(M)}^i(\mathbf{s}_{M,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(M)}^i(\mathbf{s}_{M,H_M+1}) \\ \stackrel{(iii)}{=} &\sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_m+1}) \right) \end{aligned}$$

$$\begin{aligned}
 & + \frac{1}{n} \sum_{i=1}^n V_{j_i(1)}^i(\mathbf{s}_{1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(M)}^i(\mathbf{s}_{M,1}) \\
 & + \frac{1}{n} \sum_{i=1}^n V_{j_i(M)}^i(\mathbf{s}_{M,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(M)}^i(\mathbf{s}_{M,H_{M+1}}) \\
 = & \sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_{m+1}}) \right) \\
 & + \frac{1}{n} \sum_{i=1}^n V_{j_i(1)}^i(\mathbf{s}_{1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(M)}^i(\mathbf{s}_{M,H_{M+1}}) \\
 \stackrel{(iv)}{\leq} & \sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_{m+1}}) \right) + \frac{1}{n} \sum_{i=1}^n V_{j_i(1)}^i(\mathbf{s}_{1,1}), \tag{39}
 \end{aligned}$$

where (i) follows from the telescopic summation over h ; in (ii) we add and subtract $\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1})$ inside the summation. (iii) again uses the telescopic summation. Finally, (iv) follows by dropping a non-negative term with a negative sign.

Now consider the first term of the RHS of the Equation (39). Note that the interval ends if either of the two conditions are satisfied: 1) MAEVI is triggered by at least one agent $i \in N$; 2) all the agents reach the goal state. Let us suppose, all the agents reach the goal state hence $\mathbf{s}_{m+1,1} = \mathbf{s}_{init}$, and $\mathbf{s}_{m,H_{m+1}} = \mathbf{g}$. This implies

$$\begin{aligned}
 \frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_{m+1}}) & = \frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{init}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{g}) \\
 & = \frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{init}). \tag{40}
 \end{aligned}$$

Next, suppose the interval ended because MAEVI is triggered for some agent $i \in N$. Then we apply a trivial upper bound, i.e.,

$$\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_{m+1}}) \leq \frac{1}{n} \sum_{i=1}^n \max_{j_i} \|V_{j_i}^i\|_{\infty}. \tag{41}$$

Also recall the total number of calls to the MAEVI algorithm from Lemma 1 is given by

$$J \leq 2n^2 d \log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2n \log(T).$$

Thus, from Equation (39) we have,

$$\begin{aligned}
 & \sum_{m=1}^M \left(\sum_{h=1}^{H_m} \left\{ \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h+1}) \right\} \right) \\
 \leq & \sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{m+1,1}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,H_{m+1}}) \right) + \frac{1}{n} \sum_{i=1}^n V_{j_i(1)}^i(\mathbf{s}_{1,1}) \\
 \stackrel{(i)}{\leq} & \sum_{m=1}^{M-1} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m+1)}^i(\mathbf{s}_{init}) \cdot \mathbb{1}_{\{m+1 \in \mathcal{M}(M)\}} \right) + \frac{1}{n} \sum_{i=1}^n V_{j_i(1)}^i(\mathbf{s}_{1,1}) \\
 & + \left[2n^2 d \log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2n \log(T) \right] \cdot \frac{1}{n} \sum_{i=1}^n \max_{j_i} \|V_{j_i}^i\|_{\infty} \\
 \stackrel{(ii)}{\leq} & \sum_{m \in \mathcal{M}(M)} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{init}) \right) + \frac{1}{n} \sum_{i=1}^n V_0^i(\mathbf{s}_{init}) \\
 & + \left[2n^2 d \log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2n \log(T) \right] \cdot \frac{1}{n} \sum_{i=1}^n B_*
 \end{aligned}$$

$$\stackrel{(iii)}{=} \sum_{m \in \mathcal{M}(M)} \left(\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{init}) \right) + 1 + 2n^2 dB_* \log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2nB_* \log(T), \quad (42)$$

where the first inequality is the same as Equation (39). (i) uses the combination of Equations (40) and (41) along with the fact that if the goal state is reached by all the agents, then $m+1 \in \mathcal{M}(M)$. In (ii) we use the fact that $V_{j_i(1)}^i(\mathbf{s}_{1,1}) = V_0^i(\mathbf{s}_{init})$ for all agents $i \in N$. Also, $\|V_{j_i}^i\|_\infty \leq B_*$ for all $i \in N$. Finally (iii) follows from the fact that $V_0^i(\mathbf{s}_{init}) = 1$ for all $i \in N$. The proof follows by arranging the terms of the above Equation (42). \square

C.3 Proof of Proposition 5 (Bounding E_1)

Proof. Recall, from Equation (4) we have

$$E_1 = \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i=1}^n \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h}) \right].$$

First note that $V_{j_i(m)}^i(\mathbf{s}_{m,h}) = \min_a Q_{j_i(m)}^i(\mathbf{s}_{m,h}, a) = Q_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})$, therefore E_1 can be written as

$$E_1 = \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i=1}^n \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i=1}^n Q_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right].$$

Since the interval, m ends if the MAEVI is triggered for at least one agent or all the agents reach the goal state. Let S_m be the set of agents for whom the MAEVI is triggered in the interval m . We can decompose the above summation over S_m and S_m^c . In particular, we have

$$\begin{aligned} E_1 &= \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\ &\quad + \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m^c} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m^c} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m^c} Q_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\ &= \sum_{m=1}^M \sum_{h=1}^{H_m} [E_1(S_m) + E_1(S_m^c)]. \end{aligned}$$

To bound E_1 , we need to bound both $E_1(S_m)$ and $E_1(S_m^c)$ for each m . As opposed to the single agent SSP of [Min et al. \(2022\)](#), we need separate bounds for $E_1(S_m)$ and $E_1(S_m^c)$. First consider $E_1(S_m)$.

Bounding $E_1(S_m)$

First consider $E_1(S_m)$. To bound this, we will use the MAEVI update for each agent $i \in S_m$. Recall that the MAEVI update is given by

$$\begin{aligned} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) &= \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + (1-q) \min_{\boldsymbol{\theta} \in \mathcal{C}_{j_i(m)}^i \cap \mathcal{B}} \langle \boldsymbol{\theta}, \phi_{V^{i,(l-1)}}(\mathbf{s}_{m,h}) \rangle \\ &= \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + (1-q) \cdot \langle \boldsymbol{\theta}_{m,h}, \phi_{V^{i,(l-1)}}(\mathbf{s}_{m,h}) \rangle \\ &= \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + (1-q) \cdot \langle \boldsymbol{\theta}_{m,h}, \phi_{V^{i,(l)}}(\mathbf{s}_{m,h}) \rangle + (1-q) \cdot \langle \boldsymbol{\theta}_{m,h}, [\phi_{V^{i,(l-1)}} - \phi_{V^{i,(l)}}](\mathbf{s}_{m,h}) \rangle \\ &= \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + (1-q) \cdot \langle \boldsymbol{\theta}_{m,h}, \phi_{V^{i,(l)}}(\mathbf{s}_{m,h}) \rangle + (1-q) \cdot \mathbb{P}_{m,h}[V^{i,(l-1)} - V^{i,(l)}](\mathbf{s}_{m,h}) \\ &= \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + (1-q) \cdot \mathbb{P}_{m,h}V^{i,(l)}(\mathbf{s}_{m,h}) + (1-q) \cdot \mathbb{P}_{m,h}[V^{i,(l-1)} - V^{i,(l)}](\mathbf{s}_{m,h}) \\ &\geq \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + (1-q) \cdot \mathbb{P}_{m,h}V^{i,(l)}(\mathbf{s}_{m,h}) - (1-q) \cdot \frac{1}{t_{j_i(m)}}, \end{aligned}$$

where the last inequality follows from the fact that the stopping criteria for the MAEVI for agent $i \in S_m$ is $\|V^{i,(l)} - V^{i,(l-1)}\|_\infty \leq \epsilon_{j_i}^i = \frac{1}{t_{j_i(m)}}$. So, the above inequality implies that

$$\bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) - Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \leq (1-q) \cdot \frac{1}{t_{j_i(m)}} - (1-q) \cdot \mathbb{P}_{m,h}V^{i,(l)}(\mathbf{s}_{m,h}).$$

Adding $\mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})$ on both sides in the above equation, we have

$$\begin{aligned}
 & \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \\
 & \leq \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) + (1-q) \cdot \frac{1}{t_{j_i(m)}} - (1-q) \cdot \mathbb{P}_{m,h} V^{i,(l)}(\mathbf{s}_{m,h}) \\
 & = [\mathbb{P} - \mathbb{P}_{m,h}]V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) + q \cdot \mathbb{P}_{m,h} V^{i,(l)}(\mathbf{s}_{m,h}) + (1-q) \cdot \frac{1}{t_{j_i(m)}} \\
 & \stackrel{(i)}{\leq} [\mathbb{P} - \mathbb{P}_{m,h}]V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) + qB_\star + (1-q) \cdot \frac{1}{t_{j_i(m)}} \\
 & \stackrel{(ii)}{=} [\mathbb{P} - \mathbb{P}_{m,h}]V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) + \frac{B_\star}{t_{j_i(m)}} + (1-q) \cdot \frac{1}{t_{j_i(m)}} \\
 & \stackrel{(iii)}{\leq} \left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle + \frac{B_\star + 1}{t_{j_i(m)}}, \tag{43}
 \end{aligned}$$

where (i) uses the fact that $V_{j_i(m)}^i \leq V^{i\star} \leq B_\star$, (ii) is consequence of the fact that $q = \frac{1}{t_{j_i(m)}}$, and (iii) follows by dropping a negative term. Now taking summation over $i \in S_m$, we have the following:

$$\begin{aligned}
 & \frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \\
 & \leq \left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \frac{1}{n} \sum_{i \in S_m} \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle + \frac{1}{n} \sum_{i \in S_m} \frac{B_\star + 1}{t_{j_i(m)}}.
 \end{aligned}$$

Let $\mathcal{M}_0(M) = \{m \leq M : j_i(m) \geq 1, \forall i \in N\}$ be the set of all intervals for which the output of MAEVI algorithm is $Q_{j_i(m)}^i$ for all agents $i \in N$ rather than the output Q_i^i . We first consider the intervals from $\mathcal{M}_0(M)$. So, from the above equation, we have

$$\begin{aligned}
 & \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\
 & \leq \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \frac{1}{n} \sum_{i \in S_m} \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle + \frac{1}{n} \sum_{i \in S_m} \frac{B_\star + 1}{t_{j_i(m)}} \right] \\
 & = \underbrace{\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \frac{1}{n} \sum_{i \in S_m} \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle \right]}_{A_1} + \underbrace{\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \frac{B_\star + 1}{t_{j_i(m)}} \right]}_{A_2}. \tag{44}
 \end{aligned}$$

To bound the above, we need to bound A_1 and A_2 . First consider A_1 .

$$\begin{aligned}
 A_1 & = \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \frac{1}{n} \sum_{i \in S_m} \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle \right] \\
 & = \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle \right]. \tag{45}
 \end{aligned}$$

To bound the above term, consider the inner term in the above equation for agent $i \in S_m$,

$$\begin{aligned}
 & \left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}, \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle \\
 & \stackrel{(i)}{\leq} \|\boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m,h}\|_{\Sigma_{t(m,h)}^i} \cdot \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \\
 & = \|\boldsymbol{\theta}^\star + \hat{\boldsymbol{\theta}}_{j_i(m)}^i - \hat{\boldsymbol{\theta}}_{j_i(m)}^i - \boldsymbol{\theta}_{m,h}\|_{\Sigma_{t(m,h)}^i} \cdot \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(ii)}{\leq} \left(\|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{j_i(m)}^i\|_{\Sigma_{t(m,h)}^i} + \|\hat{\boldsymbol{\theta}}_{j_i(m)}^i - \boldsymbol{\theta}_{m,h}\|_{\Sigma_{t(m,h)}^i} \right) \cdot \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \\
 &\stackrel{(iii)}{\leq} 2 \left(\|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{j_i(m)}^i\|_{\Sigma_{t(m,h)}^i} + \|\hat{\boldsymbol{\theta}}_{j_i(m)}^i - \boldsymbol{\theta}_{m,h}\|_{\Sigma_{t(m,h)}^i} \right) \cdot \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \\
 &\stackrel{(iv)}{\leq} 4\beta_T \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}.
 \end{aligned} \tag{46}$$

where (i) uses the Cauchy-Schwartz inequality. In (ii) we apply the triangle inequality; (iii) uses the following: recall $t_{j_i(m)}$ is the time at which the $j_i(m)$ -th MAEVI is triggered by agent i , and $t(m, h)$ is the time period corresponding to the h -th step in the m -th interval. Therefore, $t(m, h) \geq t_{j_i(m)}$. Therefore, by determinant doubling criteria we must have $\det(\Sigma_{t(m,h)}) \leq 2 \det(\Sigma_{t_{j_i(m)}})$ otherwise, $t(m, h)$ and $t_{j_i(m)}$ would not belong to the same interval m . The inequality follows from $\lambda_k(\Sigma_{t(m,h)}) \leq 2\lambda_k(\Sigma_{t_{j_i(m)}})$ for all $k \in [nd]$, where $\lambda_k(\cdot)$ is the k -th eigenvalue. Finally, (iv) follows from the fact that $\boldsymbol{\theta}^*$ and $\boldsymbol{\theta}_{m,h}$ belong to the confidence ellipsoid $\mathcal{C}_{j_i(m)}^i$.

Moreover, we also have the following:

$$\begin{aligned}
 \left\langle \boldsymbol{\theta}^* - \boldsymbol{\theta}_{m,h}, \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle &\leq \left\langle \boldsymbol{\theta}^*, \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle \\
 &= \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \\
 &\leq B_\star.
 \end{aligned} \tag{47}$$

From Equations (46) and (47), we have

$$\begin{aligned}
 \left\langle \boldsymbol{\theta}^* - \boldsymbol{\theta}_{m,h}, \phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle &\leq \min\{B_\star, 4\beta_T \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\} \\
 &\leq 4\beta_T \min\{1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\},
 \end{aligned} \tag{48}$$

where the second inequality is because of the fact that $B_\star \leq \beta_T \leq 4\beta_T$. This implies from Equation (45) we have,

$$\begin{aligned}
 A_1 &\leq \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{4\beta_T}{n} \sum_{i \in S_m} \min\{1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\} \right] \\
 &\leq \frac{4\beta_T}{n} \sqrt{\left(\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} 1 \right) \cdot \left(\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\sum_{i \in S_m} \min\{1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\}^2 \right] \right)} \\
 &= \frac{4\beta_T}{n} \sqrt{T \cdot \left(\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\sum_{i \in S_m} \min\{1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\}^2 \right] \right)},
 \end{aligned} \tag{49}$$

the above inequality follows from the Cauchy-Schwartz inequality (product of inner term with 1. Hence summation will become the inner product). To bound the other term in the above square root, we will use the Lemma 4 given in this SM as follows:

$$\begin{aligned}
 &\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\sum_{i \in S_m} \min\{1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\}^2 \right] \\
 &\stackrel{(i)}{\leq} \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\min \left\{ \sum_{i \in S_m} 1, \sum_{i \in S_m} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}^2 \right\} \right] \\
 &\stackrel{(ii)}{\leq} \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\min \left\{ \sum_{i \in N} 1, \sum_{i \in N} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}^2 \right\} \right] \\
 &= n \cdot \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\min \left\{ 1, \frac{1}{n} \sum_{i \in N} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}^2 \right\} \right]
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(iii)}{\leq} 2n \left[nd \log \left(\frac{\text{trace}(\lambda \mathbf{I}) + T \cdot \frac{1}{n} \sum_{i \in N} B_\star^2 nd}{nd} \right) - \log(\det(\lambda \mathbf{I})) \right] \\
 &= 2n \left[nd \log \left(\frac{\lambda nd + TB_\star^2 nd}{nd} \right) - \log(\lambda nd) \right] \\
 &\leq 2n^2 d \log \left(\frac{\lambda nd + TB_\star^2 nd}{\lambda nd} \right) \\
 &= 2n^2 d \log \left(1 + \frac{TB_\star^2}{\lambda} \right),
 \end{aligned}$$

where (i) follows by interchanging the summation and min operator. In (ii) we replace the sum over $i \in S_m$ by $i \in N$; and (iii) follows from Lemma 4 of [Abbasi-Yadkori et al. \(2011\)](#), and the fact that $\max_{m \in \mathcal{M}_0(M)} \|\phi_{V_{j_i(m)}^i}(\cdot, \cdot)\| \leq B_\star \sqrt{nd}$. Combining this with the Equation (49), we have

$$\begin{aligned}
 A_1 &\leq \frac{4\beta_T}{n} \sqrt{T \cdot 2n^2 d \log \left(1 + \frac{TB_\star^2}{\lambda} \right)} \\
 &= 4\beta_T \sqrt{2Td \log \left(1 + \frac{TB_\star^2}{\lambda} \right)}.
 \end{aligned} \tag{50}$$

Next consider A_2 , recall

$$\begin{aligned}
 A_2 &= \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \frac{B_\star + 1}{t_{j_i(m)}} \right] \stackrel{(i)}{\leq} \sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in N} \frac{B_\star + 1}{t_{j_i(m)}} \right] \\
 &\stackrel{(ii)}{=} (B_\star + 1) \cdot \frac{1}{n} \sum_{i \in N} \sum_{j_i=1}^J \sum_{t=t_{j_i}+1}^{t_{j_i+1}} \frac{1}{t_{j_i}} \\
 &\stackrel{(iii)}{\leq} (B_\star + 1) \cdot \frac{1}{n} \sum_{i \in N} \sum_{j_i=1}^J \frac{2t_{j_i}}{t_{j_i}} \\
 &= 2(B_\star + 1)J \\
 &\stackrel{(iv)}{\leq} 2(B_\star + 1) \left[2n^2 d \log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + 2n \log(T) \right] \\
 &= 4(B_\star + 1) \left[n^2 d \log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + n \log(T) \right] \\
 &\stackrel{(v)}{\leq} 4n^2 d (B_\star + 1) \left[\log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + \log(T) \right],
 \end{aligned} \tag{51}$$

where (i) follows by replacing the summation over $i \in S_m$ by summation over $i \in N$. In (ii) we use the fact that the total number of time steps can be represented either as $\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} 1$ or $\sum_{j_i=1}^J \sum_{t=t_{j_i}+1}^{t_{j_i+1}} 1$. Since the time doubling condition $t \geq 2t_{j_i}$ in the algorithm implies $t_{j_i+1} \leq 2t_{j_i}$ for all j_i , we have (iii). In (iv) we use the bound on J given in Lemma 1. Finally, (v) follows from the fact that $n \log(T) \leq n^2 d \log(T)$.

Plugging Equations (50) and (51) in Equation (44) we have

$$\begin{aligned}
 &\sum_{m \in \mathcal{M}_0(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\
 &\leq 4\beta_T \sqrt{2Td \log \left(1 + \frac{TB_\star^2}{\lambda} \right)} + 4n^2 d (B_\star + 1) \left[\log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + \log(T) \right].
 \end{aligned} \tag{52}$$

Finally, to bound $E_1(S_m)$, we need to bound

$$\sum_{m \in \mathcal{M}_0^c(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right].$$

Recall that $\mathcal{M}_0^c(M)$ is the set of all intervals m such that $j_i(m) = 0$ for all $i \in N$, i.e., the intervals before the first call of MAEVI sub-routine. Since $t_0 = 1$ by triggering condition $t \geq 2t_0$ for all agent, so the first MAEVI will be called at $t = 2$ by all the agents. Therefore, we have

$$\begin{aligned}
 & \sum_{m \in \mathcal{M}_0^c(M)} \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\
 &= \sum_{h=1}^2 \left[\frac{1}{n} \sum_{i \in S_m} \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \frac{1}{n} \sum_{i \in S_m} \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - \frac{1}{n} \sum_{i \in S_m} Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\
 &\stackrel{(i)}{\leq} \sum_{h=1}^2 \left[\frac{1}{n} \sum_{i \in N} \bar{c}(\mathbf{s}_{1,h}, \mathbf{a}_{1,h}, \mathbf{w}^i) + \frac{1}{n} \sum_{i \in N} \mathbb{P}V_{j_i(1)}^i(\mathbf{s}_{1,h}, \mathbf{a}_{1,h}) \right] \\
 &= \sum_{h=1}^2 \left[\frac{1}{n} \sum_{i \in N} \bar{c}(\mathbf{s}_{1,h}, \mathbf{a}_{1,h}, \mathbf{w}^i) + \frac{1}{n} \sum_{i \in N} \mathbb{P}V_0^i(\mathbf{s}_{1,h}, \mathbf{a}_{1,h}) \right] \\
 &\stackrel{(ii)}{\leq} 2(n+1), \tag{53}
 \end{aligned}$$

where (i) follows by dropping a negative term and (ii) is because $\bar{c}(\mathbf{s}_{1,h}, \mathbf{a}_{1,h}; \mathbf{w}^i) \leq n$ as the maximum congestion observed by any agent is n and the private component of the cost to each agent is $K^i \leq 1$ moreover, $|V_0^i(\mathbf{s})| \leq 1$.

Combining Equations (52) and (53), we have

$$\sum_{m=1}^M \sum_{h=1}^{H_m} E_1(S_m) \leq 4\beta_T \sqrt{2Td \log \left(1 + \frac{TB_\star^2}{\lambda} \right)} + 4nd(B_\star + 1) \left[\log \left(1 + \frac{TB_\star^2 nd}{\lambda} \right) + \log(T) \right] + 2(n+1). \tag{54}$$

Now, to bound E_1 , we need to bound $\sum_{m=1}^M \sum_{h=1}^{H_m} E_1(S_m^c)$.

Bounding $E_1(S_m^c)$

To bound $E_1(S_m^c)$ we first observe the following: Since $i \in S_m^c$, this implies that the MAEVI is not called for agent i in m -th interval; therefore, agent i will not update its optimistic estimator. So, in the m -th interval, agent i uses the same optimistic estimator used in the $(m-1)$ -th interval. Thus, there is some $1 < k_i < m$ for agent i such that the last MAEVI call made by agent i was at $(m-k_i)$ -th interval, i.e., $Q_{j_i(m)}^i = Q_{j_i(m-1)}^i = Q_{j_i(m-2)}^i = \dots = Q_{j_i(m-k_i)}^i$. And this $Q_{j_i(m-k_i)}^i$ is obtained from the MAEVI update, hence $i \in S_{m-k_i}$, whereas $i \in S_{m-k_i+1}^c, S_{m-k_i+2}^c, \dots, S_m^c$. So, from the analysis done for those agents for whom the MAEVI is called at $(m-k_i)$ -th interval, and Equation (48) we have for all agents $i \in S_{m-k_i}$

$$\left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m-k_i,h}, \phi_{V_{j_i(m-k_i)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle \leq 4\beta_T \min \left\{ 1, \|\phi_{V_{j_i(m-k_i)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m-k_i,h)}^{i-1}} \right\}.$$

Moreover, from equation (43), we have

$$\begin{aligned}
 & \bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \mathbb{P}V_{j_i(m-k_i)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \\
 &\leq \left\langle \boldsymbol{\theta}^\star - \boldsymbol{\theta}_{m-k_i,h}, \phi_{V_{j_i(m-k_i)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right\rangle + \frac{B_\star + 1}{t_{j_i(m-k_i)}} \\
 &\leq 4\beta_T \min \left\{ 1, \|\phi_{V_{j_i(m-k_i)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m-k_i,h)}^{i-1}} \right\} + \frac{B_\star + 1}{t_{j_i(m-k_i)}}.
 \end{aligned}$$

Also, note that $V_{j_i(m)}^i = V_{j_i(m-1)}^i = V_{j_i(m-2)}^i = \dots = V_{j_i(m-k_i)}^i$, this is because MAEVI is not called in the intermediate intervals by agent i , as $V_{j_i(m)}^i(\mathbf{s}_{m,h}) = \min_{\mathbf{a}} Q_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a})$. So for each agent, $i \in S_m^c$, the above equation can be written as

$$\bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \leq 4\beta_T \min \left\{ 1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \right\} + \frac{B_\star + 1}{t_{j_i(m)}}.$$

Furthermore, this is true for all the agents $i \in S_m^c$. Taking summation over $i \in S_m^c$, we have

$$\begin{aligned}
 & \frac{1}{n} \sum_{i \in S_m^c} \left[\bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\
 & \leq \frac{4\beta_T}{n} \sum_{i \in S_m^c} \min\{1, \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}}\} + \frac{1}{n} \sum_{i \in S_m^c} \frac{B_* + 1}{t_{j_i(m)}} \\
 & \stackrel{(i)}{\leq} \frac{4\beta_T}{n} \min \left\{ \sum_{i \in S_m^c} 1, \sum_{i \in S_m^c} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \right\} + \frac{1}{n} \sum_{i \in S_m^c} \frac{B_* + 1}{t_{j_i(m)}} \\
 & \stackrel{(ii)}{\leq} \frac{4\beta_T}{n} \min \left\{ \sum_{i \in N} 1, \sum_{i \in N} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \right\} + \frac{1}{n} \sum_{i \in N} \frac{B_* + 1}{t_{j_i(m)}} \\
 & \leq 4\beta_T \min \left\{ 1, \frac{1}{n} \sum_{i \in N} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \right\} + \frac{1}{n} \sum_{i \in N} \frac{B_* + 1}{t_{j_i(m)}},
 \end{aligned}$$

where (i) follows by interchanging min and summation. In (ii) we replace summation over $i \in S_m^c$ by summation over $i \in N$.

Since, $Q_{j_i(m)}^i$ is same as $Q_{j_i(m-k_i)}^i$, taking summation over m , and h in the above we will have the same bound as we have for $E_1(S_m)$. This implies,

$$\begin{aligned}
 E_1(S_m^c) &= \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i \in S_m^c} \left[\bar{c}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}; \mathbf{w}^i) + \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) - Q^{i,(l)}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \right] \\
 &\leq \sum_{m=1}^M \sum_{h=1}^{H_m} \left[4\beta_T \min \left\{ 1, \frac{1}{n} \sum_{i \in N} \|\phi_{V_{j_i(m)}^i}(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})\|_{\Sigma_{t(m,h)}^{i-1}} \right\} + \frac{1}{n} \sum_{i \in N} \frac{B_* + 1}{t_{j_i(m)}} \right] \\
 &\leq 4\beta_T \sqrt{2Td \log \left(1 + \frac{TB_*^2}{\lambda} \right)} + 4nd(B_* + 1) \left[\log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + \log(T) \right] + 2(n+1), \quad (55)
 \end{aligned}$$

the last inequality uses the same ideas as used for bounding $\sum_{m=1}^M \sum_{h=1}^{H_m} E_1(S_m)$. Combining Equations (54) and (55), we have the bound for E_1 as follows:

$$\begin{aligned}
 E_1 &= \sum_{m=1}^M \sum_{h=1}^{H_m} [E_1(S_m) + E_1(S_m^c)] \\
 &\leq 8\beta_T \sqrt{2Td \log \left(1 + \frac{TB_*^2}{\lambda} \right)} + 8nd(B_* + 1) \left[\log \left(1 + \frac{TB_*^2 nd}{\lambda} \right) + 2 \log(T) \right] + 4(n+1).
 \end{aligned}$$

This ends the proof. \square

C.4 Proof of Proposition 6 (Bounding E_2)

Proof. Recall E_2 is given by

$$E_2 = \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n V_{j_i(m)}^i(\mathbf{s}_{m,h+1}) - \frac{1}{n} \sum_{i=1}^n \mathbb{P}V_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right].$$

The first thing to note is that E_2 is the sum of the martingale differences. However, the function $V_{j_i(m)}^i$ is random, not necessarily bounded. So we can apply Azuma-Hoeffding inequality. Let us define an filtration $\{\mathcal{F}_{m,h}\}_{m,h}$ such that $\mathcal{F}_{m,h}$ is the σ -field of all the history up until $(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})$ but doesn't contain $\mathbf{s}_{m,h+1}$. Thus, $(\mathbf{s}_{m,h}, \mathbf{a}_{m,h})$ is $\mathcal{F}_{m,h}$ -measurable. Moreover, $V_{j_i(m)}^i$ is $\mathcal{F}_{m,h}$ measurable. By definition of operator \mathbb{P} , we have $\mathbb{E}[V_{j_i(m)}^i | \mathcal{F}_{m,h}] = \mathbb{P}V_{j_i(m)}^i$ for all $i \in N$, which shows that E_2 is martingale difference sequence. To deal with the problem that $V_{j_i(m)}^i$ might not be bounded, define an auxiliary sequence

$$\tilde{V}_{j_i(m)}^i := \min\{B_*, V_{j_i(m)}^i\},$$

It follows that $\tilde{V}_{j_i(m)}^i$ is $\mathcal{F}_{m,h}$ -measurable.

$$\begin{aligned} E_2 &= \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n \tilde{V}_{j_i(m)}^i(\mathbf{s}_{m,h+1}) - \frac{1}{n} \sum_{i=1}^n \mathbb{P} \tilde{V}_{j_i(m)}^i(\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right] \\ &\quad + \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n [V_{j_i(m)}^i - \tilde{V}_{j_i(m)}^i](\mathbf{s}_{m,h+1}) - \frac{1}{n} \sum_{i=1}^n \mathbb{P}[V_{j_i(m)}^i - \tilde{V}_{j_i(m)}^i](\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right]. \end{aligned}$$

Since $\tilde{V}_{j_i(m)}^i$ is bounded, we can apply the Azuma-Hoeffding inequality Lemma 5 and get with probability at least $1 - \delta/2n$ that

$$\begin{aligned} E_2 &= \frac{1}{n} \sum_{i=1}^n 2B_* \sqrt{2T \log \left(\frac{T}{\delta/2n} \right)} + \\ &\quad \sum_{m=1}^M \sum_{h=1}^{H_m} \left[\frac{1}{n} \sum_{i=1}^n [V_{j_i(m)}^i - \tilde{V}_{j_i(m)}^i](\mathbf{s}_{m,h+1}) - \frac{1}{n} \sum_{i=1}^n \mathbb{P}[V_{j_i(m)}^i - \tilde{V}_{j_i(m)}^i](\mathbf{s}_{m,h}, \mathbf{a}_{m,h}) \right]. \end{aligned}$$

Note that under the MAEVI analysis and optimism, we have $\tilde{V}_{j_i(m)}^i = V_{j_i(m)}^i$ for all $j_i(m) \geq 1$ for all $i \in N$. Moreover, initialization $\tilde{V}_0^i = V_0^i$ for all $i \in N$ implies the second term in RHS is zero. Thus, with probability at least $1 - \delta$ we have

$$E_2 = \frac{1}{n} \sum_{i=1}^n 2B_* \sqrt{2T \log \left(\frac{2nT}{\delta} \right)} = 2B_* \sqrt{2T \log \left(\frac{2nT}{\delta} \right)},$$

This ends the proof. \square

D SOME USEFUL RESULTS

Theorem 7 (Theorem 1, Abbasi-Yadkori et al. (2011)). *Let $\{\mathcal{F}_t\}_{t=0}^\infty$ be a filtration. Suppose $\{\eta_t\}_{t=1}^\infty$ is a \mathbb{R} -valued stochastic process such that η_t is \mathcal{F}_t -measurable and $\eta_t | \mathcal{F}_{t-1}$ is B -sub-Gaussian. Let $\{\phi_t\}_{t=1}^\infty$ be an \mathbb{R}^d -valued stochastic process such that ϕ_t is \mathcal{F}_{t-1} -measurable. Assume that Σ is an $d \times d$ positive definite matrix. For any $t \geq 1$, define*

$$\Sigma_t = \Sigma + \sum_{k=1}^t \phi_k \phi_k^\top \text{ and } \mathbf{a}_t = \sum_{k=1}^t \eta_k \phi_k.$$

Then, for any $\delta > 0$, with probability at least δ , for all t , we have

$$\|\Sigma_t^{-1/2} \mathbf{a}_t\|_2 \leq B \sqrt{2 \log \left(\frac{\det(\Sigma_t)^{1/2}}{\delta \cdot \det(\Sigma)^{1/2}} \right)}.$$

Theorem 8 (Determinant-trace inequality; Lemma 11 Abbasi-Yadkori et al. (2011)). *Assume $\phi_1, \phi_2, \dots, \phi_t \in \mathbb{R}^d$ and for any $s \leq t$, $\|\phi_s\|_2 \leq L$. Let $\lambda > 0$ and $\Sigma_t = \lambda \mathbf{I} + \sum_{s=1}^t \phi_s \phi_s^\top$. Then*

$$\det(\Sigma_t) \leq (\lambda + tL^2/d)^d.$$

Lemma 4 (Lemma 11 in Abbasi-Yadkori et al. (2011)). *Let $\{\phi_t\}_{t=1}^\infty$ be in \mathbb{R}^d such that $\|\phi_t\| \leq L$ for all t . Assume Σ_0 is psd matrix in $\mathbb{R}^{d \times d}$, and let $\Sigma_t = \Sigma_0 + \sum_{s=1}^t \phi_s \phi_s^\top$. Then we have*

$$\sum_{s=1}^t \min\{1, \|\phi_s\|_{\Sigma_{s-1}^{-1}}\} \leq 2 \left[d \log \left(\frac{\text{trace}(\Sigma_0) + tL^2}{d} \right) - \log \det(\Sigma_0) \right].$$

Lemma 5 (Azuma-Hoeffding inequality, anytime version). *Let $\{X_t\}_{t=0}^\infty$ be a real-valued martingale such that for every $t \geq 1$, it holds that $|X_t - X_{t-1}| \leq B$ for some $B \geq 0$. Then for any $0 < \delta \leq 1/2$, with probability at least $1 - \delta$, the following holds for all $t \geq 0$*

$$|X_t - X_0| \leq 2B \sqrt{2t \log \left(\frac{t}{\delta} \right)}.$$

Lemma 6 (Kushner-Clark Lemma [Kushner and Yin \(2003\)](#); [Metivier and Priouret \(1984\)](#)). Let $\mathcal{X} \subseteq \mathbb{R}^p$ be a compact set and let $h : \mathcal{X} \rightarrow \mathbb{R}^p$ be a continuous function. Consider the following recursion in p -dimensions

$$x_{t+1} = \Gamma\{x_t + \gamma_t[h(x_t) + \zeta_t + \beta_t]\}. \quad (56)$$

Let $\hat{\Gamma}(\cdot)$ be transformed projection operator defined for any $x \in \mathcal{X} \subseteq \mathbb{R}^p$ as

$$\hat{\Gamma}(h(x)) = \lim_{0 < \eta \rightarrow 0} \left\{ \frac{\Gamma(x + \eta h(x)) - x}{\eta} \right\},$$

then the ODE associated with Equation (56) is $\dot{x} = \hat{\Gamma}(h(x))$.

Assumption 5. Kushner-Clark lemma requires the following assumptions

1. Stepsize $\{\gamma_t\}_{t \geq 0}$ satisfy $\sum_t \gamma_t = \infty$, and $\gamma_t \rightarrow 0$ as $t \rightarrow \infty$.
2. The sequence $\{\beta_t\}_{t \geq 0}$ is a bounded random sequence with $\beta_t \rightarrow 0$ almost surely as $t \rightarrow \infty$.
3. For any $\epsilon > 0$, the sequence $\{\zeta_t\}_{t \geq 0}$ satisfy

$$\lim_t \mathbb{P} \left(\sup_{p \geq t} \left\| \sum_{\tau=t}^p \gamma_\tau \zeta_\tau \right\| \geq \epsilon \right) = 0.$$

Kushner-Clark lemma is as follows: suppose that ODE $\dot{x} = \hat{\Gamma}(h(x))$ has a compact set \mathcal{K}^* as its asymptotically stable equilibria, then under Assumption 5, x_t in Equation (56) converges almost surely to \mathcal{K}^* as $t \rightarrow \infty$.