Distributed Coordination Maximization over Networks: A Stochastic Approximation Approach

Hyeryung Jang
Department of Electrical Engineering, KAIST
Daejeon, Republic of Korea
hrjang@lanada.kaist.ac.kr

Se-Young Yun
Los Alamos National Laboratory
Los Alamos, New Mexico, USA
syun@lanl.gov

Jinwoo Shin, Yung Yi
Department of Electrical Engineering, KAIST
Daejeon, Republic of Korea
jinwoos@kaist.ac.kr, yiyung@kaist.edu

ABSTRACT

In various online/offline networked environments, it is very popular that the system can benefit from coordinating actions of two interacting nodes, but incur some cost due to such coordination. Examples include a wireless sensor networks with duty cycling, where a sensor node consumes a certain amount of energy when it is awake, but a coordinated operation of sensors enables some meaningful tasks, *e.g.*, sensed data forwarding, collaborative sensing of a phenomenon, or efficient decision of further sensing actions. In this paper, we formulate an optimization problem that captures the amount of coordination gain at the cost of node activation over networks. This problem is challenging since the target utility is a function of the long-term time portion of the inter-coupled activations of two adjacent nodes, and thus a standard Lagrange duality theory is hard to apply to obtain a distributed decomposition as in the standard NUM (Network Utility Maximization). We propose a fully-distributed algorithm that requires only one-hop message passing. Our approach is inspired by a control of Ising model in statistical physics, and the proposed algorithm is motivated by a stochastic approximation method that runs a Markov chain incompletely over time, but provably guarantees its convergence to the optimal solution. We validate our theoretical findings on convergence and optimality through extensive simulations under various scenarios.

CCS Concepts

•**Theory of computation** → **Stochastic approximation;** *Design and analysis of algorithms; Distributed algorithms;* •**Mathematics of computing** → *Network optimization;*

[0]This work was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No.B0717-16-0034, Versatile Network System Architecture for Multi-dimensional Diversity).

MobiHoc'16, July 04-08, 2016, Paderborn, Germany

DOI: http://dx.doi.org/10.1145/2942358.2942366

Keywords

Coordination Maximization, Stochastic Approximation, Distributed Algorithm

1. INTRODUCTION

In many online/offline networks, a variety of gains among nodes and users are generated when they are in a similar state. Examples are diverse, as summarized in what follows:

- *Wireless sensor networks.* In wireless sensor networks with duty cycled node activations for energy saving, each sensor node decides to be awake or not over time, which further depends on its neighbors' wake-up state and distance to the node. When two nearby nodes communicate, they are equipped with a robust wireless channel for mutual communication, and thus their coordination (*e.g.*, message exchange) can become more powerful at the cost of energy consumption while they are awake [5, 6]. Thus, to achieve the desired coordination gain while turning off redundant sensors, each sensor node should smartly make decisions of waking up or not, which should often be done in a distributed manner.
- *Social networks.* In online/offline social networks, social relationships and interactions are of critical interests, since strength of such interactions often determines how the network evolves, *e.g.*, adoption of a technology or spread of information. For example, when a new technology becomes available, using the social relationships, more coordination gain due to compatibility of the technology between two individuals is generated, whereas a certain cost due to technology adoption is incurred, *e.g.*, buying a new OS software [1, 12, 18].
- *Wireless cellular networks.* Multiple base stations coordinate to serve a bandwidth-thirsty user such as one in the cell-edge, to provide more frequency and time resource in a collaborative manner, often referred to as CoMP (Coordinated Multi-Point communications) [2, 13]. In this application, a logical choice for coordination is to appropriately choose a signal power to serve the user. Thus, this clearly generates the coordination gain to the user, but at the cost of energy consumption.

In this paper, we formulate an optimization problem, which we call *coordination maximization* that captures the gain due to peer-to-peer coordination between two nodes, but also the cost due to individual node activations, as in the following form:

$$\max_{\text{all } \lambda_i \text{ and all } \lambda_{ij}} \sum_{\substack{\text{all connected} \\ \text{node pairs } (i,j)}} U_{ij}(\lambda_{ij}) - \sum_{\text{all node } i} C_i(\lambda_i), \quad (1)$$

where $U_{ij}(\cdot)$ and $C_i(\cdot)$ are the coordination utility and the node activation cost functions, respectively. Intuitively, λ_{ij} is the long-term time portion when both nodes i and j are simultaneously activated and thus coordinated, and λ_i is just the long-term time portion when node i is activated. This optimization appears to be a simple variant of a standard NUM (Network Utility Maximization) [11, 15, 21], where it is allowed to easily develop a distributed algorithm converging to the optimal solution. However, the problem in (1) significantly differs from a standard NUM problem, thus developing a distributed algorithm is far from being trivial. The main challenges lie in the fact that the standard Lagrange duality theory for a distributed decomposition is not possible since the objective function includes the term which is a function of the long-term inter-coupling of the states of a pair of connected nodes, and thus, separability is not permitted.

1.1 Main Contribution

The main contribution of this paper lies in developing a fully distributed algorithm that requires only one-hop message passing, yet provably converges to the optimal solution of the coordination maximization problem in (1), which we call **CoordMax** algorithm. In many engineering systems, we often observe the trade-off between efficiency and complexity. Optimal algorithms often require extensive message passing or heavy computational challenges, but lightweight approximate algorithms incur optimality degradation, whereas our algorithm achieves optimality only with locally-limited message passing.

The key technique of our algorithm is sketched in what follows: we first construct a Markov chain whose states are the set of all vectors representing the node activations. Then, each node runs a *distributed* algorithm corresponding to running and sampling the Markov chain in a decentralized manner, *i.e.*, MCMC (Markov Chain Monte Carlo), so that the marginal long-term coordination and node activation rates are the solution of (1). We appropriately choose a step-size to provably ensure its convergence to the optimal solution. The proof technique is based on a stochastic approximation technique where we show that the stochastic dynamics asymptotically converge to a differential equation whose stationary point corresponds to the target optimum.

CoordMax can be interpreted as a way of controlling the parameters of the Glauber dynamics over an *Ising model* [8, 17] in statistical physics. In particular, to capture pairwise coordinations in the network and node-wise costs, we take a graphical model framework, which is Ising model, that represents a singleton and pair-wise interactions among random variables by a parametrized undirected graph structure. **CoordMax** corresponds to an efficient, distributed control mechanism that suitably chooses the parameters of Ising model, leading the corresponding marginals of nodes and edges to the optimal solution of a coordination maximization problem (1). Despite an extensive array of research on Ising model in statistical physics [7, 8, 24, 25], most of them study (i) phase transitions and critical phenomena for given Ising parameters, or (ii) configuration decision strategies again for given Ising parameters towards the fast propagation of a certain state.

1.2 Related Work

A large array of works about network utility maximization (NUM) problem have been studied, see [11, 15, 21] for surveys and tutorial. The objective of NUM problem is to maximize a sum of all nodes' utilities, while not considering any edge-wise status, thus the separability provides a useful dual-based decomposition for an easy development of a distributed algorithm. In recent years, the researches on achieving optimality in both throughput and utility in wireless scheduling (in a decentralized manner) have been studied [9, 16, 19, 22, 23]. The CSMA setting of the multi-hop wireless networks can be understood by a *hard-core* graphical model, which is a special case of Ising model in the way that it only captures node-wise dependency. The Intuitive idea of these works is that wireless links dynamically adjust access intensities by using local information such as queue-length so as to converge to the utility-optimal point or stabilize the network whenever possible.

A variety of gains from coordinating actions of wireless terminals or users have been widely studied in wireless networks. There exists many approaches encouraging coordination among base stations in cellular mobile systems, *e.g.*, coordinated multi-point transmission (CoMP) [2, 10, 13]. In this context, coordination between base stations allows power gain, channel rank advantage, and/or diversity gains. In the area of sensor networks, distributed coordination schemes have been proposed recently, where sensors adaptively select to be coordinators or not, *i.e.*, stay awake and forward information or not, while turning off redundant sensors. This often (i) preserves capacity [6], (ii) improves the network's lifetime and communication latency, by using a geo-location information of sensors [27], or (iii) builds a self-configuring localization system [5].

Many questions in online/offline social networks are based on the social interactions among individuals. Coordinating actions of two individuals in social networks is of importance, since the power of interactions often determines how the network evolves. The authors in [26] have highlighted an importance of a coordination mechanism such as a social structure among individuals, for efficient knowledge sharing. The authors in [1] have conducted large field experiments to identify the effect of coordination, *e.g.*, strength of social ties, in word-of-mouth advertising application. Many works including [12, 18, 20] studied how to control node activations or seeding set of nodes for speeding up a diffusion and mathematically analyzed diffusion time, where individuals' interaction is modeled by a coordination game [18, 20] or Ising model in statistical physics [12].

Our paper is also based on the importance of the coordination powers among nodes or individuals. However, our main interest is to find a sequence of node activations (generated by a *distributed algorithm*) leading to the solution of the problem in (1) that maximizes the network-wide coordination gain with each node's activation incurring some cost.

Organization. The rest of the paper is organized as follows. In Section 2, we introduce a network model and formulate a coordination maximization problem. Section 3 provides

the description of our algorithm and its convergence and optimality analysis, including backgrounds on Ising model and stochastic approximation procedure. We then present simulation results in Section 4 to support the theoretical findings, and conclude in Section 5.

2. MODEL AND PRELIMINARIES

2.1 Network Model

Network model. In this paper, we consider a network $G = (V, E)$ consisting of a set V of nodes and a set $E \subset V \times V$ of edges. With this graphical representation, each node corresponds to an agent in social networks or a sensor node in wireless sensor networks, and each edge corresponds to a physical connectivity or a social relationship between nodes, *i.e.*, $(i, j) \in E$ means that node i and node j are connected and have an interaction. Note that we study undirected networks where interaction requires mutual consent, *i.e.*, (i, j) is equivalent to (j, i). Let $\mathcal{N}(i) = \{j \in V : (i, j) \in E\}$ denote the neighbors of node i.

Configuration and coordination scheduling. We consider a continuous time framework. Let $\sigma_i(t) \in \{0, 1\}$ indicate whether node i is *active* at time t or not, *i.e.*, $\sigma_i(t) = 1$ means that the node i is active at time t, and 0 otherwise. We say that nodes i and j are (or edge (i, j) is) *coordinated* when $\sigma_i(t)\sigma_j(t) = 1$. We also denote by $\boldsymbol{\sigma}(t) = [\sigma_i(t)]_{i\in V}$ a *node configuration* at time t, and it is clear that a set of possible configurations of the graph G is defined as $\mathcal{I}(G) := \{0, 1\}^{|V|}$. To formally discuss a coordination gain, which we will introduce later, we define a *coordination configuration* vector as follows:

$$\boldsymbol{\phi}(\boldsymbol{\sigma}) := ([\sigma_i]_{i\in V}, [\sigma_i\sigma_j]_{(i,j)\in E}), \tag{2}$$

which is an augmented configuration vector capturing the activation status of nodes and coordination status of edges. Then, every coordination configuration belongs to $\Phi(G) := \{0, 1\}^{|V|+|E|}$. Now, a *coordination scheduling* (or simply scheduling) algorithm is a mechanism that selects $\boldsymbol{\sigma}(t) \in \mathcal{I}(G)$ (thus a coordination configuration $\boldsymbol{\phi}(\boldsymbol{\sigma}(t)) \in \Phi(G)$ is also determined) over time $t \in \mathbb{R}_+$.

Coordination capacity region. We now define the maximum achievable *coordination region* (also called coordination capacity region) $\Lambda \subset [0, 1]^{|V|+|E|}$ of the network, which is the convex hull of the feasible coordination configuration set $\Phi(G)$, *i.e.*,

$$\Lambda := \Big\{ \sum_{\boldsymbol{\sigma}\in\mathcal{I}(G)} \alpha(\boldsymbol{\sigma})\boldsymbol{\phi}(\boldsymbol{\sigma}) : \sum_{\boldsymbol{\sigma}\in\mathcal{I}(G)} \alpha(\boldsymbol{\sigma}) = 1, \alpha(\cdot) \geq 0 \Big\}.$$

The intuition of the notion of coordination capacity region comes from the fact that any coordination scheduling algorithm has to choose a node configuration from $\mathcal{I}(G)$ over time (thus a coordination $\boldsymbol{\phi}(\boldsymbol{\sigma})$ is determined), and $\alpha(\boldsymbol{\sigma})$ denotes the fraction of time selecting a node configuration $\boldsymbol{\sigma}$ (and thus a coordination $\boldsymbol{\phi}(\boldsymbol{\sigma})$). Hence, the long-term average time portion of node activation and edge coordination induced by any scheduling algorithm must belong to Λ.

2.2 Problem Formulation

Objective. We require nodes and edges to control its long-term time portion of node activation and edge coordination close to some boundary of Λ. Specifically, we aim at designing a coordination scheduling algorithm that makes decisions $\boldsymbol{\sigma}(t) \in \mathcal{I}(G)$ over time t so that the long-term time portion of node activation and edge coordination converges to a solution of the following optimization problem:

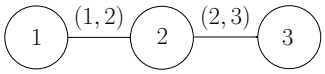

Figure 1: An example line network with 3 nodes and 2 edges, where there are 8 feasible node configurations $\boldsymbol{\sigma} \in \{0, 1\}^3$.

$$\textbf{(OPT)} \quad \max_{\boldsymbol{\lambda}\in\Lambda} \quad \sum_{(i,j)\in E} U_{ij}(\lambda_{ij}) - \sum_{i\in V} C_i(\lambda_i). \tag{3}$$

This mathematical problem captures dependencies among nodes and a trade-off between coordination gain and node activation cost, where $U_{ij} : [0, 1] \to \mathbb{R}$ is a strictly concave and twice-differentiable coordination gain function of edge $(i, j) \in E$, and $C_i : [0, 1] \to \mathbb{R}$ is a strictly convex, twice-differentiable cost function of node $i \in V$. It is easy to see that **OPT** has the unique solution $\boldsymbol{\lambda}^* := ([\lambda_i^*]_{i\in V}, [\lambda_{ij}^*]_{(i,j)\in E})$. The network utility is defined as a total coordination utility subtracted by a total incurring cost. More coordination gain is generated as nodes i and j are coordinated more, but it also incurs more cost of nodes i and j to be activated.

One of the choices of U_{ij} and C_i is a *linear scalarization* of multi-objective optimization, *i.e.*, a weighted coordination gain and weighted incurring activation cost with weights $\boldsymbol{w} = [w_{ij}]_{(i,j)\in E}$ and $\boldsymbol{\gamma} = [\gamma_i]_{i\in V}$, respectively. For example, one can choose $\forall i \in V, \forall (i, j) \in E$,

$$U_{ij}(x) = w_{ij} \cdot \log(x), \quad C_i(x) = \gamma_i \cdot \frac{1}{1 - x},$$

so that **OPT** becomes a weighted proportional fairness of coordination as well as a weighted cost incurred by node activation, where the trade-off between those two is externally controlled by $\boldsymbol{w}$ and $\boldsymbol{\gamma}$.

In this work, our goal is to design a distributed coordination scheduling algorithm using only local information, *e.g.*, one-hop message passing, that converges to the optimal solution of **OPT**, *i.e.*,

$$\begin{aligned} \lim_{T\to\infty} \frac{1}{T} \int_0^T \sigma_i(t)dt &= \lambda_i^* \quad \text{for all} \quad i \in V, \\ \lim_{T\to\infty} \frac{1}{T} \int_0^T \sigma_i(t)\sigma_j(t)dt &= \lambda_{ij}^* \quad \text{for all} \quad (i, j) \in E. \end{aligned}$$

Example. To motivate further and illustrate, we provide an example of a coordination maximization problem and its solution structure, where we use a line topology with 3 nodes and 2 edges, as depicted in Figure 1. The coordination maximization problem **OPT** in this case is expressed by:

$$\max_{\boldsymbol{\lambda}\in\Lambda} \Big[U_{12}(\lambda_{12}) + U_{23}(\lambda_{23}) - \Big(C_1(\lambda_1) + C_2(\lambda_2) + C_3(\lambda_3) \Big) \Big].$$

Let the long-term time portion of the activation and coordination of nodes and edges be characterized by the distribution $\pi(\boldsymbol{\sigma})$ over each node configuration $\boldsymbol{\sigma} \in \{0, 1\}^3$, *i.e.*,

$$\lambda_1 = \pi(1, 0, 0) + \pi(1, 0, 1) + \pi(1, 1, 0) + \pi(1, 1, 1),$$

$$\begin{aligned}\lambda_2 &= \pi(0,1,0)+\pi(0,1,1)+\pi(1,1,0)+\pi(1,1,1),\\ \lambda_3 &= \pi(0,0,1)+\pi(0,1,1)+\pi(1,0,1)+\pi(1,1,1),\\ \lambda_{12} &= \pi(1,1,0)+\pi(1,1,1),\\ \lambda_{23} &= \pi(0,1,1)+\pi(1,1,1).\end{aligned}$$

Note that the total coordination gain is generated according to the long-term coordination time portion of two edges, *i.e.*, $\lambda_{12}, \lambda_{23}$, and the total incurring cost is due to the long-time activation of three nodes, *i.e.*, $\lambda_1, \lambda_2, \lambda_3$.

A smart scheduling is required since each node's activation should be coordinated with its neighboring nodes in order to produce enough gain at the cost of activation. For expositional convenience, let us choose the following utility and cost functions: $U_{12}(x) = U_{23}(x) = \log(x)$ and $C_1(x) = C_2(x) = C_3(x) = x^2$, then a simple algebra gives the following distributions and the resulting optimal solution:

$$\pi^*(0,0,0) = 0.4226, \quad \pi^*(1,1,1) = 0.5774,$$

$$(\lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_{12}^*, \lambda_{23}^*) = (0.5774, 0.5774, 0.5774, 0.5774, 0.5774),$$

where the most and the only "efficient" node configuration turns out to be $(1,1,1)$ with some cost balancing by avoiding the activation of any node, *i.e.*, scheduling $(0,0,0)$. We now choose a different cost function only for node 3 by $C_3(x) = 3x^2$ (*i.e.*, more cost is incurred for node 3), then the optimal distribution and the resulting optimal solution become:

$$\pi^*(0,0,0) = 0.5,\ \pi^*(1,1,0) = 0.0915,\ \pi^*(1,1,1) = 0.4085,$$

$$(\lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_{12}^*, \lambda_{23}^*) = (0.5, 0.5, 0.4085, 0.5, 0.4085),$$

where the optimal solution is attained by assigning some probability to the configuration $(1,1,0)$ rather than giving a high priority only to $(1,1,1)$ in the first case. A lot of challenges arise for (i) more complex, general topologies, (ii) general form of utility and cost functions $U_{ij}(\cdot)$ and $C_i(\cdot)$, and more importantly (iii) the solution should be found in a distributed manner.

3. ALGORITHM AND ANALYSIS

In this section, we develop an adaptive scheduling algorithm that operates in a distributed manner, yet asymptotically converges to the solution of **OPT** problem (3), which we call **CoordMax** algorithm. To describe the algorithm, we introduce a parameter $\boldsymbol{\theta} \in \mathbb{R}^{|V|+|E|}$ as:

$$\boldsymbol{\theta} = ([\theta_i]_{i\in V}, [\theta_{ij}]_{(i,j)\in E}),$$

where θ_i is a parameter for node i and θ_{ij} is that for edge (i,j). A configuration decision is made according to this parameter $\boldsymbol{\theta}$, which we call *algorithm parameter*, and thus **CoordMax** is an algorithm that makes decisions of node configuration via controlling $\boldsymbol{\theta}$, over time t. Moreover, by "distributed", we mean that each node and edge updates its parameter with one-hop message passing, relying only on pure local observations.

3.1 Algorithm Description

3.1.1 Configuration Change Dynamics

We first present a time-by-time dynamics of changing configurations of the system for a fixed algorithm parameter, simply called **CCD** (Configuration Change Dynamics). Under the dynamics with a fixed $\boldsymbol{\theta}$, every node has a Poisson clock with unit rate, and nodes try to decide a new configuration $\boldsymbol{\sigma}'$ based on a current configuration $\boldsymbol{\sigma}$ by the following procedure, consisting of two steps **S1** and **S2**.

Configuration Change Dynamics: **CCD**($\boldsymbol{\theta}$)

Input: Algorithm parameter $\boldsymbol{\theta}$, current configuration $\boldsymbol{\sigma}$.
Output: New configuration $\boldsymbol{\sigma}'$.

S1. Select a node i when node i's clock ticks.

S2. Node i changes its configuration from σ_i to σ_i'

$$\sigma_i' = \begin{cases} 1, & \text{with probability } \frac{\exp(\theta_i + \sum_{j\in\mathcal{N}(i)} \sigma_j\theta_{ij})}{1+\exp(\theta_i+\sum_{j\in\mathcal{N}(i)}\theta_{ij}\sigma_j)} \\ 0, & \text{with probability } \frac{1}{1+\exp(\theta_i+\sum_{j\in\mathcal{N}(i)}\theta_{ij}\sigma_j)} \end{cases} \quad (4)$$

and $\sigma_j' = \sigma_j$ for all $j \neq i$.

Note that Poisson clock of each node leads to the uniform node selection, and given a graph G, **CCD** changes the configuration over time in a distributed manner that requires only one-hop message passing. In particular, when node i's clock ticks in **S1**, it requires to know (i) configuration status of neighboring nodes, *i.e.*, $\{\sigma_j : j \in \mathcal{N}(i)\}$, and (ii) parameter of neighboring edges, *i.e.*, $\{\theta_{ij} : j \in \mathcal{N}(i)\}$, in **S2**.

We can easily check that **CCD** for a given parameter $\boldsymbol{\theta}$ leads to a continuous-time Markov chain $\{\boldsymbol{\sigma}(t)\}_{t=0}^{\infty}$ achieving the following stationary distribution $p_{\boldsymbol{\theta}} = [p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})]_{\boldsymbol{\sigma}\in\mathcal{I}(G)}$ on the finite state space $\mathcal{I}(G)$:

$$\begin{aligned} p_{\boldsymbol{\theta}}(\boldsymbol{\sigma}) &\propto \exp\{\langle\boldsymbol{\theta}, \boldsymbol{\phi}(\boldsymbol{\sigma})\rangle\} \\ &= \exp\{\sum_{i\in V}\theta_i\sigma_i + \sum_{(i,j)\in E}\theta_{ij}\sigma_i\sigma_j\}, \end{aligned} \quad (5)$$

where $\langle \boldsymbol{a}, \boldsymbol{b}\rangle$ is the inner product of two vectors $\boldsymbol{a}$ and $\boldsymbol{b}$. Moreover, $\{\boldsymbol{\sigma}(t)\}_{t=0}^{\infty}$ is an irreducible, aperiodic, and reversible Markov process [14].

Given the parameter $\boldsymbol{\theta}$, the ergodicity and reversibility of the Markov process imply that the marginal probability of nodes and edges under the stationary distribution $p_{\boldsymbol{\theta}}$, denoted by $\boldsymbol{s}(\boldsymbol{\theta}) = ([s_i(\boldsymbol{\theta})]_{i\in V}, [s_{ij}(\boldsymbol{\theta})]_{(i,j)\in E})$, become the long-term time portion of node activation and edge coordination, and can be characterized as: $\forall i \in V, (i,j) \in E$,

$$\begin{aligned} s_i(\boldsymbol{\theta}) &= \mathbb{E}_{p_{\boldsymbol{\theta}}}[\sigma_i] = \sum_{\boldsymbol{\sigma}\in\mathcal{I}(G):\sigma_i=1} p_{\boldsymbol{\theta}}(\boldsymbol{\sigma}), \\ s_{ij}(\boldsymbol{\theta}) &= \mathbb{E}_{p_{\boldsymbol{\theta}}}[\sigma_i\sigma_j] = \sum_{\boldsymbol{\sigma}\in\mathcal{I}(G):\sigma_i\sigma_j=1} p_{\boldsymbol{\theta}}(\boldsymbol{\sigma}). \end{aligned} \quad (6)$$

3.1.2 CoordMax Algorithm

We now describe a **CoordMax** algorithm in Algorithm 1. Time is divided into frames $k = 0, 1, \cdots$, of fixed durations T, and each node i updates parameter θ_i and $\{\theta_{ij} : j \in \mathcal{N}(i)\}$ at the end of each frame in **S3**. To this end, at the beginning of each frame k, each node i sends a message including its own information, *i.e.*, configuration status and parameter of node i, to its neighboring nodes in **S1**, and then measures (i) its node activation rate $\hat{s}_i[k]$, and (ii) edge coordination rate of its connecting edges $\{\hat{s}_{ij}[k] : j \in \mathcal{N}(i)\}$ in **S2**, *i.e.*,

$$\hat{s}_i[k] = \frac{1}{T}\int_{kT}^{(k+1)T} \sigma_i(x)dx,$$

$$\hat{s}_{ij}[k] \;=\; \frac{1}{T}\int_{kT}^{(k+1)T} \sigma_i(x)\sigma_j(x)dx.$$

This is interpreted as the number of activations and coordinations at nodes and edges at the frame k, and it can be measured from the local observations by exchanging messages about its own configuration status with neighbors.

Algorithm 1 CoordMax: At each frame $k = 0, 1, \ldots,$

Input: Efficiency factor β, boundary values $\theta_{\min}, \theta_{\max}$
Output: Algorithm parameter $\boldsymbol{\theta}[k+1]$
Initialize: Set $\boldsymbol{\theta}[0]$ arbitrarily, and $a[0] = 0$.

S1. Each node i sends a message $(\sigma_i[k], \theta_i[k])$ to its neighboring nodes $\mathcal{N}(i)$.

S2. The system runs the configuration change dynamics $\mathbf{CCD}(\boldsymbol{\theta})$, and each node i records the number of its activations and coordinations of its connecting edges.

S3. Each node i updates local parameters, θ_i and $\{\theta_{ij} : j \in \mathcal{N}(i)\}$, with the step-size constant $a[k]$ as follows:

$$\theta_i[k+1] = \left[\theta_i[k] + a[k]\left(C_i'^{-1}\left(\frac{-\theta_i[k]}{\beta}\right) - \hat{s}_i[k]\right)\right]_{\theta_{\min}}^{\theta_{\max}}, \tag{7}$$

$$\theta_{ij}[k+1] = \left[\theta_{ij}[k] + a[k]\left(U_{ij}'^{-1}\left(\frac{\theta_{ij}[k]}{\beta}\right) - \hat{s}_{ij}[k]\right)\right]_{\theta_{\min}}^{\theta_{\max}}. \tag{8}$$

In **CoordMax** algorithm, β is an efficiency parameter, as shown later on; $a : \mathbb{N} \to \mathbb{R}_+$ is a step-size function; $\theta_{\min}, \theta_{\max}$ are the given constants; and $[\cdot]_x^y := \max(y, \min(x, \cdot))$. The choice of a step-size a is of critical importance for convergence of this algorithm, and β controls the efficiency of the algorithm, which is elaborated in Section 3.2. As mentioned before, **CoordMax** operates in a distributed manner in the way that each node updates its local parameters based only on local information, *i.e.*, one-hop message passing and local observations.

3.2 Convergence and Optimality Analysis

We now analyze the convergence and asymptotic optimality of **CoordMax** algorithm. We consider an approximated problem **A-OPT** (parametrized by $\beta > 0$) of **OPT**:

(A-OPT)

$$\begin{aligned} \max \quad & H(\boldsymbol{\mu}) + \beta\Big(\sum_{(i,j)\in E} U_{ij}(\lambda_{ij}) - \sum_{i\in V} C_i(\lambda_i)\Big) \\ \text{over} \quad & \boldsymbol{\mu} \in \mathcal{M}, \quad \boldsymbol{\lambda} \in [0,1]^{|V|+|E|} \\ \text{subject to} \quad & \lambda_i = \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i], \quad \forall i \in V, \\ & \lambda_{ij} = \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i\sigma_j], \quad \forall (i,j) \in E, \end{aligned} \tag{9}$$

where we let $\mathcal{M}$ be a space of all probability distributions over $\mathcal{I}(G)$, and denote by $H(\boldsymbol{\mu})$ the *entropy* of $\boldsymbol{\mu} \in \mathcal{M}$, *i.e.*,

$$H(\boldsymbol{\mu}) = -\sum_{\boldsymbol{\sigma}\in\mathcal{I}(G)} \mu(\boldsymbol{\sigma}) \log \mu(\boldsymbol{\sigma}).$$

Since the entropy term $H(\boldsymbol{\mu})$ is bounded above and below, a solution of **A-OPT** provides an approximate solution of **OPT**, where the gap due to this approximation decays for increasing β (we will discuss the trade-off between approximation efficiency and convergence speed in Section 4 via numerical results). We also denote by $\boldsymbol{\lambda}[k] = ([\lambda_i[k]]_{i\in V}, [\lambda_{ij}[k]]_{(i,j)\in E})$ the vector representing the cumulative average node activation rate and edge coordination rate up to frame k, *i.e.*,

$$\begin{aligned} \lambda_i[k] &= \frac{1}{k}\sum_{m=0}^{k-1} \hat{s}_i[m], \quad \forall i \in V, \\ \lambda_{ij}[k] &= \frac{1}{k}\sum_{m=0}^{k-1} \hat{s}_{ij}[m], \quad \forall (i,j) \in E. \end{aligned}$$

To prove the convergence and optimality of **CoordMax**, we need following assumption (A1), meaning that we choose $\theta_{\min}$ and $\theta_{\max}$, such that the interval $[\theta_{\min}, \theta_{\max}]$ is large enough to include the optimal solution of **A-OPT**.

(A1) If $\boldsymbol{\theta}^0 \in \mathbb{R}^{|V|+|E|}$ solves for all $i \in V$ and $(i,j) \in E$,

$$\begin{aligned} \theta_i^0 &= -\beta C_i'\Big(\sum_{\boldsymbol{\sigma}\in\mathcal{I}(G)} \sigma_i p_{\boldsymbol{\theta}^0}(\boldsymbol{\sigma})\Big), \\ \theta_{ij}^0 &= \beta U_{ij}'\Big(\sum_{\boldsymbol{\sigma}\in\mathcal{I}(G)} \sigma_i\sigma_j p_{\boldsymbol{\theta}^0}(\boldsymbol{\sigma})\Big), \end{aligned}$$

then $\theta_{\min} \le \theta_i^0 \le \theta_{\max}$ for all $i \in V$, and $\theta_{\min} \le \theta_{ij}^0 \le \theta_{\max}$ for all $(i,j) \in E$. Note that, for example, if the coordination gain function $U_{ij}(\cdot)$ is such that $U_{ij}'(0) < \infty$, then (A1) for θ_{ij}^0 is satisfied when $\theta_{\min} \le \beta U_{ij}'(1)$ and $\theta_{\max} \ge \beta U_{ij}'(0)$, and we can similarly check it for other edges, and other nodes with cost function $C_i(\cdot)$. One can easily verify that this assumption provides a guarantee that the convergent point of our algorithm actually belongs to bounded region.

The next theorem is our main result, which states the convergence of **CoordMax** to a point arbitrarily close to the coordination-optimal solution, under the decreasing step-size function.

Theorem 1 (Convergence and Optimality) Choose a positive step-size function $a(\cdot)$ satisfying

$$\sum_{k=0}^{\infty} a[k] = \infty, \quad \sum_{k=0}^{\infty} a[k]^2 \le \infty. \tag{10}$$

Under (A1), for any initial condition $\boldsymbol{\theta}[0]$,

(i) *Convergence.* **CoordMax** converges to $(\boldsymbol{\lambda}^o, \boldsymbol{\theta}^o)$, *i.e.*,

$$\lim_{k\to\infty} \boldsymbol{\theta}[k] = \boldsymbol{\theta}^o \;\text{ and }\; \lim_{k\to\infty} \boldsymbol{\lambda}[k] = \boldsymbol{\lambda}^o, \;\text{ almost surely,}$$

where $\boldsymbol{\theta}^o$ and $\boldsymbol{\lambda}^o$ are such that $(p_{\boldsymbol{\theta}^o}, \boldsymbol{\lambda}^o)$ is the solution of **A-OPT** in (9) (over $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$).

(ii) *Optimality.* Furthermore, **CoordMax** approximately solves **OPT** in the following sense:

$$\begin{aligned} &\Big(\sum_{(i,j)\in E} U_{ij}(\lambda_{ij}^o) - \sum_{i\in V} C_i(\lambda_i^o)\Big) \ge \\ &\Big(\sum_{(i,j)\in E} U_{ij}(\lambda_{ij}^*) - \sum_{i\in V} C_i(\lambda_i^*)\Big) - \frac{\log|\mathcal{I}(G)|}{\beta} \end{aligned} \tag{11}$$

The proof of Theorem 1 is presented in Section 3.3. **CoordMax** is interpreted as a stochastic approximation procedure with controlled Markov noise, and a main technical challenge lies in handling a non-trivial coupling between Markov process of **CCD** and parameter $\boldsymbol{\theta}$ updates. Simply, the conditions of the step-size function in Theorem 1 play a crucial role to provide a provable convergence, in the way that we intuitively expect that with decreasing step-sizes, the speed of variations of the algorithm parameter tends to zero after sufficiently long time, see Section 3.3 for the mathematical detail. Moreover, Theorem 1 states that **CoordMax** converges towards a point that is arbitrarily close to the coordination-optimizer, as we choose sufficiently large β. A typical example choice of the step-size function is $a[k] = \frac{1}{k}$. Moreover, θ_i may be interpreted as an aggressiveness (or willingness) of node i to be activated, and θ_{ij} may be interpreted as a power of interaction between nodes i and j. We may expect that a parameter θ_{ij} of a dominant edge that has a large coordination gain, will reach to a large value in result.

3.3 Proof of Theorem 1

The convergence analysis of **CoordMax** is on the strength of stochastic approximation theory. As we will verify later, our algorithm is interpreted as a stochastic approximation procedure with controlled continuous-time Markov process, where the stationary distribution of the underlying Markov process from **CCD** indeed corresponds to an Ising model. In Section 3.3.1, we first describe Ising model, and provide the convergence analysis of a general stochastic approximation procedure with a controlled Markov process, based on an ordinary differential equation approach. Then, we show the convergence and asymptotic optimality of **CoordMax** in Section 3.3.2.

3.3.1 Preliminaries

Ising model. Under our coordination scheduling algorithm represented in Section 3.1, a change of a configuration from $\boldsymbol{\sigma}$ to $\boldsymbol{\sigma}'$ is affected by a current configuration $\boldsymbol{\sigma}$ and current algorithm parameter $\boldsymbol{\theta}$, and the stationary distribution of the resulting Markov chain, *i.e.*, $p_{\boldsymbol{\theta}}$, is a parametrized by $\boldsymbol{\theta}$. To mathematically represent and understand such conditional dependencies among nodes, we can explore a graphical model, which has been emerged as a powerful tool in variety of fields for succinct representations of joint probability distributions of some random variables by a graph structure. *Ising model* is a class of the graphical models which capture a singleton and pair-wise interactions among binary random variables via a parameter vector $\boldsymbol{\eta}$, so-called Ising parameter throughout this paper, *i.e.*, nodes accommodate their configurations in such a way to maximize the number of weighted singleton and pair-wise agreements, formally defined in the following way: for any feasible configuration $\boldsymbol{x}$,

$$p_{\boldsymbol{\eta}}(\boldsymbol{x}) \;:=\; \exp\{\sum_{i\in V}\eta_i x_i + \sum_{(i,j)\in E}\eta_{ij}x_i x_j - A(\boldsymbol{\eta})\}, \quad (12)$$

where $A(\boldsymbol{\eta})$ is a normalizing constant. We can now easily check that the stationary distribution of Markov chain resulting from **CCD** dynamics for a given $\boldsymbol{\theta}$, *i.e.*, $p_{\boldsymbol{\theta}} = [p_{\boldsymbol{\theta}}(\boldsymbol{\sigma})]_{\boldsymbol{\sigma}\in\mathcal{I}(G)}$, indeed corresponds to Ising model with Ising parameter $\boldsymbol{\theta}$ over configuration space $\mathcal{I}(G)$. Moreover, **CCD** may be regarded as a Glauber dynamics over an Ising model under continuous-time setting, and **CoordMax** can be interpreted as a way of controlling the Ising parameter of the **CCD** dynamics over time.

Stochastic approximation with controlled Markov noise. Consider a general discrete-time stochastic process with the following form:

$$x[k+1] = x[k] + a[k]\cdot v(x[k], Y[k]), \quad \forall k\in\mathbb{N}, \quad (13)$$

where $x[k]\in\mathbb{R}^L$ is the system state at the iteration k; $a[k]$ corresponds to the step-size; and $Y[k]$ is a random variable representing the random observation during iteration k to update the system state. This is often called a stochastic approximation with *controlled continuous-time Markov process*, in [3, 4]. Here, (i) $\{z(s)\}_{s\geq 0}$ is a stochastic process taking values in a finite set, (ii) for $s\in[k, k+1)$, $z(s)$ evolves as a continuous-time Markov process $z^{x[k]}(s)$ with a transition kernel controlled by $x[k]$, and (iii) the observation $Y[k]$ is a function of $\{z(s)\}_{k\leq s<k+1}$, *i.e.*, $Y[k] = \int_k^{k+1} f(z(s))ds$, where f is a bounded function, and $v(x, Y)$ is a bounded, continuous, Lipschitz in x and uniformly over Y. We assume that for any $x\in\mathbb{R}^L$, the controlled Markov kernel G^x is irreducible and ergodic with stationary distribution π^x, and furthermore, the mapping $x\mapsto G^x$ is continuous and $x\mapsto\pi^x$ is Lipschitz continuous. In the following, for all $x\in\mathbb{R}^L$, $\xi^x(dy)$ denotes the stationary distribution of $\int_0^1 f(z^x(s))ds$, where $z^x(\cdot)$ is a Markov process with a transition kernel G^x. We also assume that $x[k]$ remains bounded, which can be imposed by projecting the algorithm to a bounded subset.

We use a positive monotonically decreasing function $a[k]$ such that $\sum_k a[k] = \infty$ and $\sum_k a[k]^2 < \infty$. From the step-size, we define a virtual time scale $\alpha(k) = \sum_{i=0}^{k-1} a[i]$. Then, a continuous-time interpolation of the system state for the time scale α is defined as follows: $\forall k\in\mathbb{N}$, for all $\tau\in[\alpha(k), \alpha(k+1))$,

$$x_\alpha(\tau) = x[k] + (x[k+1] - x[k])\times\frac{\tau-\alpha(k)}{\alpha(k+1)-\alpha(k)}. \quad (14)$$

Under the afore-mentioned assumptions, Theorem 2 states the convergence guarantee of the iterative procedure (13).

Theorem 2 (Theorem 1 of [22]) Let $T>0$, and denote by $\tilde{x}^s(\cdot)$ the solution on $[s, s+T]$ of the following ordinary differential equation (ODE):

$$\dot{x}(\tau) = \int_y v(x(\tau), y)\cdot\xi^{x(\tau)}(dy) \quad (15)$$

with $\tilde{x}^s(s) = x_\alpha(s)$. Then, we have almost surely,

$$\lim_{s\to\infty}\sup_{\tau\in[s,s+T]}\|x_\alpha(\tau) - \tilde{x}^s(\tau)\| = 0.$$

Intuitively, we expect that due to the decreasing step-sizes, the speed of variations of the system state decreases and tends to 0 when time sufficiently grows. As a consequence, the dynamics of a stochastic process (13) are close to those of an irreducible and ergodic Markov process with fixed generator (as if the system state was frozen), and has time to converge to its ergodic behavior. Hence, when time grows large, we have that the system behaves as if the observation was averaged, *i.e.*, for continuous and bounded f, we have that component-wise, almost surely,

$$\int_y v(x,y)\xi^x(dy) \;=\; \sum_{z\in\mathcal{Z}} v(x, f(z))\pi^x(z).$$

Remark. Theorem 2 states that as time evolves, the dynamics of the underlying Markov process is averaged due to the decreasing step-size, thus "almost reaching the stationary status." Thus it suffices to see how the ODE (15) behaves. We note that the convergence analysis in Theorem 2 follows other known results on stochastic approximation procedures [4]: under the given step-size conditions, the dynamics of (13) converges to the internally chain transitive invariant set of the ODE (15) via continuous-time interpolation (14), Corollary 8 of [4][pp.74]. In particular, when the ODE (15) has a unique fixed stable equilibrium point x^*, we have almost surely: $x[k] \to x^*$ as $k \to \infty$.

3.3.2 Proof of Theorem 1

We first prove that **CoordMax** converges to the optimal solution of the approximated problem **A-OPT** in (9), and then show that the solution is asymptotically optimal of the problem **OPT** in (3). Our main proof strategy follows the stochastic approximation procedure whose limiting behavior is understood by an ordinary differential equation (ODE) [4].

PROOF. We show that in ***Step 1***, the dynamics (7), (8) asymptotically approach deterministic trajectory. In ***Step 2***, we prove that the resulting deterministic trajectory converges to the solution of **A-OPT**, and finally in ***Step 3***, we show that the solution of **A-OPT** is asymptotically optimal of the **OPT**. To do this, we take the afore-mentioned results Theorem 2 in classical stochastic approximation theory and Corollary 8 of [4][pp.74].

Step 1: Converging and averaging. We first define a virtual time scale as:

$$\alpha(k) = \sum_{m=0}^{k-1} a[m].$$

From the discrete-time sequence $\{\boldsymbol{\theta}[k], k \in \mathbb{N}\}$ of (7) and (8), we take a continuous-time interpolation of the algorithm parameter in the following way. Define $\{\boldsymbol{\theta}(\tau) : \tau \in \mathbb{R}_+\}$ as: $\forall k \in \mathbb{N}$, for all $\tau \in [\alpha(k), \alpha(k+1))$,

$$\boldsymbol{\theta}(\tau) = \boldsymbol{\theta}[k] + (\boldsymbol{\theta}[k+1] - \boldsymbol{\theta}[k]) \times \frac{\tau - \alpha(k)}{\alpha(k+1) - \alpha(k)}. \quad (16)$$

We also define $\hat{\boldsymbol{s}}(\tau) := \hat{\boldsymbol{s}}[k] \cdot \mathbf{1}_{\alpha(k) \le \tau \le \alpha(k+1)}$, where $\mathbf{1}_A$ is the indicator function for the event A. It should be clear that **CoordMax** is a stochastic approximation algorithm with controlled Markov noise as defined in (13). The equivalence is obtained by: $x[k] \equiv \boldsymbol{\theta}[k]$; $Y[k] \equiv \hat{\boldsymbol{s}}[k]$; $\{z(s)\}_{k \le s < k+1} \equiv \{\boldsymbol{\sigma}(s)\}_{k \le s < k+1}$ is the process recording the configuration from **CCD** with $\boldsymbol{\theta}[k]$ during frame k; $f(z(s)) \equiv \boldsymbol{\phi}(\boldsymbol{\sigma}(s))$ is a coordination configuration; $\pi^x \equiv p_{\boldsymbol{\theta}}$ is the stationary distribution (5) of the **CCD**($\boldsymbol{\sigma}$); finally

$$\begin{aligned} v(x,y) &\equiv U_{ij}'^{-1}\left(\frac{x}{\beta}\right) - y, \quad \text{for edge component,} \\ v(x,y) &\equiv C_i'^{-1}\left(\frac{-x}{\beta}\right) - y, \quad \text{for node component.} \end{aligned}$$

One can now easily verify that the assumptions of (13) in Section 3.3.1 are satisfied under our setup. Note that under our setup of $U_{ij}(\cdot)$ and $C_i(\cdot)$, v is a bounded Lipschitz continuous function, Markov process generated by **CCD** is continuous function of $\boldsymbol{\theta}$, and moreover $\boldsymbol{\theta} \mapsto p_{\boldsymbol{\theta}}$ is Lipschitz continuous. Then, from Theorem 2, we have following lemma:

Lemma 1 Let $T > 0$, and denote by $\tilde{\boldsymbol{\theta}}^w(\cdot)$ the solution on $[w, w+T]$ of the following ODE: for all $i \in V$ and $(i,j) \in E$,

$$\begin{aligned} \dot{\theta}_i(\tau) &= C_i'^{-1}\left(\frac{-\theta_i(\tau)}{\beta}\right) - \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \sigma_i p_{\boldsymbol{\theta}(\tau)}(\boldsymbol{\sigma}), \\ \dot{\theta_{ij}}(\tau) &= U_{ij}'^{-1}\left(\frac{\theta_{ij}(\tau)}{\beta}\right) - \sum_{\boldsymbol{\sigma} \in \mathcal{I}(G)} \sigma_i \sigma_j p_{\boldsymbol{\theta}(\tau)}(\boldsymbol{\sigma}), \quad (17) \end{aligned}$$

with $\tilde{\boldsymbol{\theta}}^w(w) = \boldsymbol{\theta}(w)$. Then, we have almost surely,

$$\lim_{w \to \infty} \sup_{\tau \in [w, w+T]} \|\boldsymbol{\theta}(\tau) - \tilde{\boldsymbol{\theta}}^w(\tau)\| = 0.$$

Lemma 1 is a direct result from Theorem 2, and it states that the interpolated continuous trajectory $\boldsymbol{\theta}(\cdot)$ from the discrete sequence $\boldsymbol{\theta}[\cdot]$ of **CoordMax** asymptotically approaches the trajectory of $\tilde{\boldsymbol{\theta}}$. Note that if the ODE system (17) has a unique fixed point $\boldsymbol{\theta}^o$, then we would have $\lim_{\tau \to \infty} \tilde{\boldsymbol{\theta}}(\tau) = \boldsymbol{\theta}^o$, and thus we have almost surely, $\lim_{k \to \infty} \boldsymbol{\theta}[k] = \boldsymbol{\theta}^o$.

Step 2: Converging to optimality. To complete the proof of optimality, we show that (17) may be interpreted as a sub-gradient algorithm solving the dual of the convex problem (9), similarly in [9]. We now consider the Lagrangian $\mathcal{L}$ of **A-OPT** with dual variables θ_i and θ_{ij} to constraints $\lambda_i = \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i]$ and $\lambda_{ij} = \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i \sigma_j]$, respectively, as follows:

$$\begin{aligned} & \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\lambda}; \boldsymbol{\theta}) \\ = \; & H(\boldsymbol{\mu}) + \beta\Big(\sum_{(i,j) \in E} U_{ij}(\lambda_{ij}) - \sum_{i \in V} C_i(\lambda_i)\Big) \\ & + \sum_{i \in V} \theta_i\Big(\mathbb{E}_{\boldsymbol{\mu}}[\sigma_i] - \lambda_i\Big) + \sum_{(i,j) \in E} \theta_{ij}\Big(\mathbb{E}_{\boldsymbol{\mu}}[\sigma_i \sigma_j] - \lambda_{ij}\Big) \\ = \; & H(\boldsymbol{\mu}) + \sum_{i \in V} \theta_i \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i] + \sum_{(i,j) \in E} \theta_{ij} \mathbb{E}_{\boldsymbol{\mu}}[\sigma_i \sigma_j] \\ & + \sum_{i \in V}\Big(-\beta C_i(\lambda_i) - \theta_i \lambda_i\Big) + \sum_{(i,j) \in E}\Big(\beta U_{ij}(\lambda_{ij}) - \theta_{ij}\lambda_{ij}\Big). \end{aligned}$$

The solution of **A-OPT** is the minimum point of the dual function, which is given by:

$$\begin{aligned} \mathcal{D}(\boldsymbol{\theta}) &= \sup \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\lambda}; \boldsymbol{\theta}) \\ &\qquad \text{over } \boldsymbol{\mu} \in \mathcal{M},\ \boldsymbol{\lambda} \in [0,1]^{|V|+|E|}. \quad (18) \end{aligned}$$

Finally, the dual optimization is given by

$$\min \quad \mathcal{D}(\boldsymbol{\theta}) \qquad \text{over} \qquad \boldsymbol{\theta} \in \mathbb{R}^{|V|+|E|}. \quad (19)$$

Note that the primal optimization problem (9) is a concave maximization and the dual optimization problem (19) is a convex minimization due to concavity of entropy and coordination functions, and convexity of cost functions under our setup. There is no duality gap and hence both have the same, unique optimal solution, and its sub-gradient algorithm will converge to the solution of (9).

Given the feasible dual $\boldsymbol{\theta} \in \mathbb{R}^{|V|+|E|}$, let $\boldsymbol{\mu}^o(\boldsymbol{\theta}), \boldsymbol{\lambda}^o(\boldsymbol{\theta})$ be the corresponding primal feasible solutions that maximize the Lagrangian $\mathcal{L}$. Given structure of $\mathcal{L}$, from the Karush-Kuhn-Tucker (KKT) conditions of (9), it follows that $\boldsymbol{\lambda}^o(\boldsymbol{\theta})$ must be such that: $\forall i \in V, \forall (i,j) \in E$,

$$\begin{aligned} \lambda_i^o(\boldsymbol{\theta}) &= \arg\max_{y \in [0,1]}\Big[-\beta C_i(y) - \theta_i y\Big], \quad (20) \\ \lambda_{ij}^o(\boldsymbol{\theta}) &= \arg\max_{y \in [0,1]}\Big[\beta U_{ij}(y) - \theta_{ij} y\Big]. \quad (21) \end{aligned}$$

For $\boldsymbol{\mu}^o(\boldsymbol{\theta})$, observe that

$$\frac{\partial \mathcal{L}}{\partial \mu(\boldsymbol{\sigma})} = -1 - \log \mu(\boldsymbol{\sigma}) + \sum_{i \in V} \theta_i \sigma_i + \sum_{(i,j) \in E} \theta_{ij} \sigma_i \sigma_j.$$

Since $\boldsymbol{\mu}^o(\boldsymbol{\theta})$ is maximizing $\mathcal{L}$, from the above, it follows that $\boldsymbol{\mu}^o(\boldsymbol{\theta}) \in (0,1)^{|V|+|E|}$. Therefore, for any $\boldsymbol{\sigma}, \boldsymbol{\rho} \in \mathcal{I}(G)$ and $\boldsymbol{\sigma} \neq \boldsymbol{\rho}$, it must be that

$$\frac{\partial \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}), \boldsymbol{\lambda}(\boldsymbol{\theta}); \boldsymbol{\theta})}{\partial \mu(\boldsymbol{\sigma})} = \frac{\partial \mathcal{L}(\boldsymbol{\mu}(\boldsymbol{\theta}), \boldsymbol{\lambda}(\boldsymbol{\theta}); \boldsymbol{\theta})}{\partial \mu(\boldsymbol{\rho})},$$

and thus

$$\mu(\boldsymbol{\sigma}) \propto \exp(\sum_{i \in V} \sigma_i \theta_i + \sum_{(i,j) \in E} \sigma_i \sigma_j \theta_{ij}), \quad \forall \boldsymbol{\sigma} \in \mathcal{I}(G).$$

From the above derivatives of Lagrangian function $\mathcal{L}$ with respect to $\boldsymbol{\mu}$, we can conclude that $\boldsymbol{\mu}^o(\boldsymbol{\theta}) = p_{\boldsymbol{\theta}}$.

Now, it follows that the dual function is characterized as:

$$\mathcal{D}(\boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{\mu}^o(\boldsymbol{\theta}), \boldsymbol{\lambda}^o(\boldsymbol{\theta}); \boldsymbol{\theta}),$$

and the dual variables $\boldsymbol{\theta}$ capture the slack in the corresponding constraints of (9). Specifically, the slack in each constraint is given by

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\mu}^o(\boldsymbol{\theta})}[\sigma_i] &- \lambda_i^o(\boldsymbol{\theta}), \quad \forall i \in V, \\ \mathbb{E}_{\boldsymbol{\mu}^o(\boldsymbol{\theta})}[\sigma_i \sigma_j] &- \lambda_{ij}^o(\boldsymbol{\theta}), \quad \forall (i,j) \in E. \end{aligned} \tag{22}$$

Accounting for (20) and (22), the sub-gradient algorithm solving the dual problem is given by following ODEs:

$$\begin{aligned} \dot{\theta_i} &= \left(C_i'\left(\frac{-\theta_i}{\beta}\right) - \mathbb{E}_{\boldsymbol{\mu}^o(\boldsymbol{\theta})}[\sigma_i] \right), \quad \forall i \in V, \\ \dot{\theta_{ij}} &= \left(U_{ij}'\left(\frac{\theta_{ij}}{\beta}\right) - \mathbb{E}_{\boldsymbol{\mu}^o(\boldsymbol{\theta})}[\sigma_i \sigma_j] \right), \quad \forall (i,j) \in E, \end{aligned} \tag{23}$$

which is equivalent to (17), provided that $\boldsymbol{\theta}(\tau)$ remains between $[\theta_{\min}, \theta_{\max}]$ component-wisely. Note that the dual solution $\boldsymbol{\theta}^o$ actually belongs to the interval $[\theta_{\min}, \theta_{\max}]$ component-wisely, as a fixed point of (23), under (A1). Since both primal and dual problem have unique solution $(\boldsymbol{\mu}^o, \boldsymbol{\lambda}^o)$, and $\boldsymbol{\theta}^o$, respectively, the sub-gradient algorithm converges to $\boldsymbol{\theta}^o$, and hence the unique solution of the ODE (17) converges to $\boldsymbol{\theta}^o$. Combining with ***Step 1*** and ***Step 2***, we can conclude that $\lim_{k \to \infty} \boldsymbol{\theta}[k] = \boldsymbol{\theta}^o$.

Step 3: Asymptotic optimality. Finally, to establish goodness of the result of coordination optimization algorithm, $\boldsymbol{\lambda}^o$, note that $(\boldsymbol{\lambda}^o, \boldsymbol{\nu}^o)$ is the optimal solution of **A-OPT**. Now, the optimal solution of **OPT**, $\boldsymbol{\lambda}^*$, along with an appropriate distribution $\boldsymbol{\nu}^* \in \mathcal{M}$ is one feasible solution of **A-OPT**. Therefore, it follows that

$$\begin{aligned} & \beta\Big(\sum_{(i,j) \in E} U_{ij}(\lambda_{ij}^*) - \sum_{i \in V} C_i(\lambda_i^*) \Big) \\ \overset{(a)}{\leq} & H(\boldsymbol{\nu}^*) + \beta\Big(\sum_{(i,j) \in E} U_{ij}(\lambda_{ij}^*) - \sum_{i \in V} C_i(\lambda_i^*) \Big) \\ \overset{(b)}{\leq} & H(\boldsymbol{\mu}^o) + \beta\Big(\sum_{(i,j) \in E} U_{ij}(\lambda_{ij}^o) - \sum_{i \in V} C_i(\lambda_i^o) \Big) \\ \overset{(c)}{\leq} & \log |\mathcal{I}(G)| + \beta\Big(\sum_{(i,j) \in E} U_{ij}(\lambda_{ij}^o) - \sum_{i \in V} C_i(\lambda_i^o) \Big) \end{aligned}$$

In the above, the first inequality (a) comes from the fact that the entropy is non-negative, the inequality (b) holds since $(\boldsymbol{\lambda}^o, \boldsymbol{\nu}^o)$ is the optimal solution of **A-OPT**, and finally we have used the fact that the maximum value of a discrete valued random variable's entropy is at most the logarithm of the cardinality of the support set $|\mathcal{I}(G)|$, for the last inequality (c). □

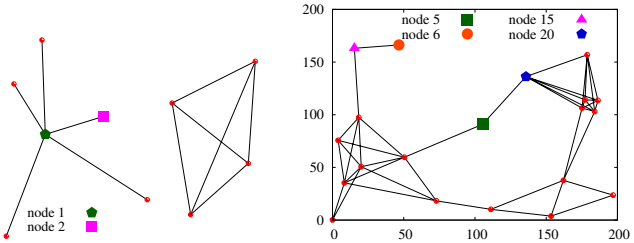

(a) `STAR` (left) and `COMP` (right) topology (b) `RAND` topology

Figure 2: Network topologies

4. NUMERICAL RESULTS

In this section, we carry out numerical experiments to assess the theoretical performance of **CoordMax** by considering networks with various topologies and cost functions.

Setup. In this paper, we consider "basic" topologies to show that our **CoordMax** converges to the accurate solution, and a random topology that is regarded as a collection of such basic topologies for more general results. The network topologies under which our results are presented here are: star, complete and random graphs. We consider a proportional fairness across edges with symmetric weights w_{ij} for coordination gain: $U_{ij}(x) = \log(x)$ for all edges $(i,j) \in E$, and consider two cost functions for nodes with symmetric weights γ_i: $C_i(x) = 2x^2$, or $C_i(x) = \frac{1}{1-x}$ for all nodes $i \in V$, as classified into the following 6 topologies:

- `STAR-S1`: Star graph with 6 nodes and 5 edges, and nodes are symmetric: $C_i(x) = 2x^2, \forall i \in V$.
- `STAR-A`: Star graph with 6 nodes and 5 edges, and nodes are asymmetric: only $C_2(x) = \frac{1}{1-x}$.
- `COMP-S1`: Complete graph with 4 nodes and 6 edges, and nodes are symmetric: $C_i(x) = 2x^2, \forall i \in V$
- `COMP-A`: Complete graph with 4 nodes and 6 edges, and nodes are asymmetric: only $C_2(x) = \frac{1}{1-x}$.
- `RAND-S1`: Random graph with 20 nodes and 42 edges, and nodes are symmetric with cost functions: $C_i(x) = 2x^2, \forall i \in V$.
- `RAND-S2`: Random graph with 20 nodes and 42 edges, and nodes are symmetric with cost functions: $C_i(x) = \frac{1}{1-x}, \forall i \in V$.

The above six topologies are depicted in Figure 2, *i.e.*, the left of Figure 2(a) for `STAR-S1`, `STAR-A`, the right of Figure 2(a) for `COMP-S1`, `COMP-A`, and Figure 2(b) for `RAND-S1`, `RAND-S2`. Moreover, we choose a step-size function of the **CoordMax** as $a[k] = 1/k$, which satisfies the conditions of Theorem 1, and take various values of β from 0.5 to 5.0, which is a factor of efficiency.

(i) Convergence to the optimal solution: To demonstrate our analytical findings of convergence and optimality, we first consider simple cases that show **CoordMax** finds the "accurate" solution, where the exact solution can be numerically solved. Then, we show **CoordMax**'s performance with two cost functions, under more general topology.

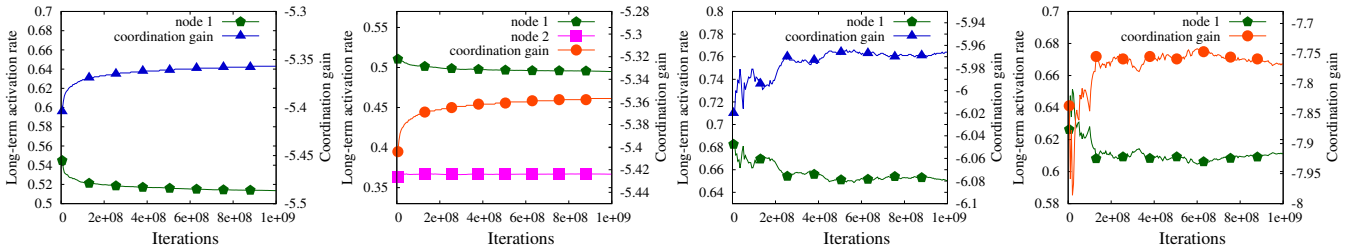

(a) Convergence of `STAR-S1` (b) Convergence of `STAR-A` (c) Convergence of `COMP-S1` (d) Convergence of `COMP-A`

Figure 3: Convergence of long-term rate (left axis) and coordination gain (right axis) to optimality on simple cases.

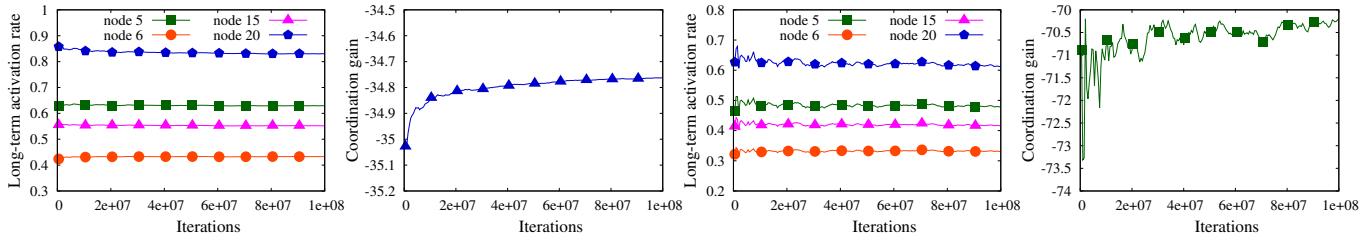

(a) Long-term rate of `RAND-S1` (b) Coordination gain of `RAND-S1` (c) Long-term rate of `RAND-S2` (d) Coordination gain of `RAND-S2`

Figure 4: Convergence of long-term rate and coordination gain on `RAND` topology, for two types of cost functions.

Simple cases: 4 types of environments are investigated here: `STAR-S1`, `STAR-A`, `COMP-S1`, and `COMP-A`, and let G^* denote the maximum coordination gain. We first solve the exact coordination-optimal solution of each case: For `STAR-S1`, $\lambda_1^* = 0.447, G^* = -5.218$, and for `STAR-A`, $\lambda_1^* = 0.433, \lambda_2^* = 0.382, G^* = -6.5905$, whose results with $\beta = 5.0$ are shown in Figure 3(a) and 3(b), respectively. We see that our **CoordMax** does converge to the solution after long iterations within a range of $O(1/\beta)$ gap. Under `STAR-A`, since node 2 has higher cost, the optimal activation rate of node 2 is attained at a less value than that under `STAR-S1`, and less coordination gain is generated on edge $(1, 2)$, as we see in Figure 3(b). Under `COMP-S1` and `COMP-A`, the exact optimal solution is attained at $\lambda_1^* = 0.6125, G^* = -5.942$ and $\lambda_1^* = 0.538, G^* = -7.619$, whose convergence results with $\beta = 5.0$ are illustrated in Figure 3(c) and 3(d), respectively. Note that the algorithm takes shorter time for convergence in star topology because each node has only 1 edge except the hub node, *i.e.*, node 1, thus pairwise interactions are less complex, while all nodes are interacted with each other, and thus **CoordMax** converges after longer time, as we see in Figure 3. The impact of asymmetries of cost functions is similarly shown under the complete topology.

Cost functions: We provide numerical results of two types of cost functions which we mentioned above, under a general random topology of Figure 2(b). For both cost functions, our **CoordMax** does converge, as we see in Figure 4. Note that at the convergent status, long-term activation rate of a node is dependent on (i) how many neighbors it has, and (ii) how powerful its neighbor is in terms of degree of coordination. As we see in Figure 4(a) and 4(c), node 6, *i.e.*, who has very little contribution to the coordination gain since it has only one neighbor, achieves the lowest long-term activation rate, while node 20 has very high long-term rate. Comparing nodes 5 and 15, even though both have two neighbors, node 5 achieves a higher long-term rate since one of its neighbors (node 20) is a hub so that node 5 may implicitly contribute to coordination gain of the network via node 20. Next, from Figure 4(b) and 4(d), we see that the network becomes less aggressive to be coordinated if nodes have cost functions $C_i(x) = \frac{1}{1-x}$, since this cost function plays a role of preventing exclusive activation of nodes, *i.e.*, infinite amount of cost is generated when activation rate becomes close to 1.

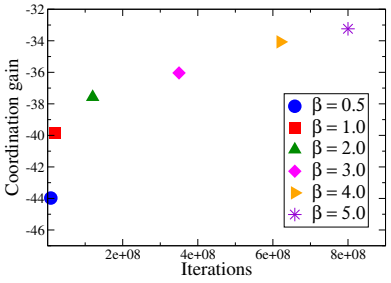

Figure 5: Trade-off between efficiency(β) and convergence speed.

(ii) Trade-off between efficiency and convergence speed: As stated in Theorem 1, incurring coordination gain gap of our algorithm is asymptotically $\frac{1}{\beta}$. To support it through numerical examples, we vary β and plot the coordination gain at the convergence point, and measure the convergence speed. Figure 5 shows that, as β grows, **CoordMax** requires longer time to converge, but the corresponding convergent point becomes closer to the optimal solution. From the numerical results under `RAND-S1`, coordination gain with $\beta = 2.0$ is -37.570 and converges after 9×10^7 iterations, while that with $\beta = 0.5$ is -43.974 and converges after 8.5×10^6 iterations.

5. CONCLUSION

In many networked environments, a variety of gains among nodes and users are generated from coordinating their actions. In this paper, we formulate an optimization problem that captures the amount of coordination gain at the cost

of node activation over networks, and we develop a fully-distributed algorithm relying only on one-hop message passing and local observations, which we call **CoordMax**. Our algorithm is inspired by a control of Ising model in statistical physics, and theoretical findings of convergence to optimality of **CoordMax** takes a stochastic approximation method that runs a Markov chain incompletely over time with a smartly designed step-size function.

6. REFERENCES

[1] E. Bakshy, D. Eckles, Y. Dean, and I. Rosenn. Social influence in social advertising: evidence from field experiments. In *Proceedings of ACM Conference on Electronic Commerce*, 2012.

[2] A. Barbieri, P. Gaal, S. Geirhofer, T. Ji, D. Malladi, Y. Wei, and F. Xue. Coordinated downlink multi-point communications in heterogeneous cellular networks. In *Proceedings of Information Theory and Applications Workshop*, 2012.

[3] V. S. Borkar. Stochastic approximation with 'controlled' markov noise. *Systems & control letters*, 55(2):139–145, 2006.

[4] V. S. Borkar. Stochastic approximation. *Cambridge Books*, 2008.

[5] N. Bulusu, D. Estrin, L. Girod, and J. Heidemann. Scalable coordination for wireless sensor networks: self-configuring localization systems. In *Proceedings of International Symposium on Communication Theory and Applications*, 2001.

[6] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of ACM Mobicom*, 2001.

[7] J. Dziarmaga. Dynamics of a quantum phase transition: Exact solution of the quantum ising model. *Physical review letters*, 95(24):245701, 2005.

[8] E. Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925.

[9] L. Jiang and J. Walrand. A distributed csma algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transaction on Networking*, 18(3):960–972, 2010.

[10] M. K. Karakayali, G. J. Foschini, and R. Valenzuela. Network coordination for spectrally efficient communications in cellular systems. *IEEE Wireless Communications*, 13(4), 2006.

[11] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society*, pages 237–252, 1998.

[12] C. E. Laciana and S. L. Rovere. Ising-like agent-based technology diffusion model: Adoption patterns vs. seeding strategies. *Physica A: Statistical Mechanics and its Applications*, 390(6):1139–1149, 2011.

[13] D. Lee, H. Seo, B. Clerckx, E. Hardouin, D. Mazzarese, S. Nagata, and K. Sayana. Coordinated multipoint transmission and reception in lte-advanced: deployment scenarios and operational challenges. *IEEE Communications Magazine*, 50(2), 2012.

[14] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.

[15] N. Li, L. Chen, and S. H. Low. Optimal demand response based on utility maximization in power networks. In *Proceedings of IEEE Power and Energy Society General Meeting*, 2011.

[16] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor. Towards utility-optimal random access without message passing. *Wireless Communications and Mobile Computing*, 10(1):115–128, 2010.

[17] B. M. McCoy and T. T. Wu. *The two-dimensional Ising model*. Courier Corporation, 2014.

[18] A. Montanari and A. Saberi. The spread of innovations in social networks. *Proceedings of National Academy of Sciences*, 107(47):20196–20201, 2010.

[19] J. Ni, B. Tan, and R. Srikant. Q-csma: queue-length-based csma/ca algorithms for achieving maximum throughput and low delay in wireless networks. *IEEE/ACM Transactions on Networking*, 20(3):825–836, 2012.

[20] J. Ok, Y. Jin, J. Shin, and Y. Yi. On maximizing the diffusion speed in social networks: Impact of random seeding and clustering. In *Proceedings of ACM Sigmetrics*, 2001.

[21] D. P. Palomar and M. Chiang. A tutorial on decomposition methods for network utility maximization. *IEEE Journal on Selected Areas in Communications*, 24(8):1439–1451, 2006.

[22] A. Proutiere, Y. Yi, T. Lan, and M. Chiang. Resource allocation over network dynamics without timescale separation. In *Proceedings of IEEE INFOCOM*, 2010.

[23] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: An efficient randomized protocol for contention resolution. In *Proceedings of ACM Sigmetrics*, 2009.

[24] H. E. Stanley, D. Stauffer, J. Kertesz, and H. J. Herrmann. Dynamics of spreading phenomena in two-dimensional ising models. *Physical review letters*, 59(20):2326, 1987.

[25] M. Suzuki and R. Kubo. Dynamics of the ising model near the critical point. i. *Journal of the Physical Society of Japan*, 24(1):51–60, 1968.

[26] W. Tsai. Social structure of "cooperation" within a multiunit organization: Coordination, competition, and intraorganizational knowledge sharing. *Organization science*, 12(2):179–190, 2002.

[27] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of annual international conference on Mobile computing and networking*, 2001.