Signal Processing 142 (2018) 330–339

Contents lists available at ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

Decentralized adaptive search using the noisy 20 questions framework in time-varying networks

Theodoros Tsiligkaridis

MIT Lincoln Laboratory, 244 Wood St., Lexington MA, 02420, USA

ARTICLE INFO

Article history:
Received 16 November 2016
Revised 21 June 2017
Accepted 30 June 2017
Available online 1 July 2017

Keywords:
Adaptive query
Adaptive search
Consensus
Probabilistic bisection
Asymptotic consistency
Social learning

ABSTRACT

This paper considers the problem of adaptively searching for an unknown target using multiple agents connected through a time-varying network topology. Agents are equipped with sensors capable of fast information processing, and we propose a decentralized collaborative algorithm for controlling their search given noisy observations. Specifically, we propose decentralized extensions of the adaptive query-based search strategy that combines elements from the 20 questions approach and social learning. Under standard assumptions on the time-varying network dynamics, we prove convergence to correct consensus on the value of the parameter as the number of iterations go to infinity. The convergence analysis takes a novel approach using martingale-based techniques combined with spectral graph theory. Our results establish that stability and consistency can be maintained even with one-way updating and randomized pairwise averaging, thus providing a scalable low complexity method with performance guarantees. We illustrate the effectiveness of our algorithm for random network topologies.

1. Introduction

Consider a set of agents that try to estimate a parameter, e.g., estimate a target state or location, collectively. The agents are connected by a time-varying information sharing network and can periodically query one of their local neighbors about the target location. In this paper we adopt a generic observation model based on query-response models where the queries are functions of agents' local information and successive queries are determined by a feedback control policy. Specifically, in the 20 questions-type model considered in this paper, the observation of each agent is coupled with the query region chosen by that agent, which is a function of its current local belief.

A centralized collaborative 20 questions framework was proposed and studied in [1], where a global centralized controller jointly or sequentially formulates optimal queries about target location for all agents. This work was later extended in [2] to the decentralized setting, in which each agent formulates his own query based on his local information and exchanges beliefs with its neighbors in a synchronous fashion (i.e, updating the beliefs of all agents simultaneously at each update step). The proposed decentralized algorithm therein consisted of two stages: (1) local belief update; and (2) local information sharing. In stage 1 each agent implements the bisection query policy of Tsiligkaridis et al. [1] to update their local belief function. In stage 2 the local belief functions are averaged over nearest neighborhoods in the information sharing network. This two-stage algorithm was proven to converge to a consensus estimate of the true state, assuming synchronous updating of all agents' beliefs in the network and irreducibility of the social interaction graph.

The decentralized collaborative 20 questions problem is applicable to large scale collaborative stochastic search applications where there is no centralized authority. Examples include: object tracking in camera networks [3]; road tracking from satellite remote sensing networks [4]; and wide area surveillance networks [5]. The 20 questions approach is motivated by adaptively asking appropriate questions and is of primary importance when computational effort and time are limited resources.

In this paper, we consider a variant of the two-stage decentralized collaborative algorithm of Tsiligkaridis et al. [2] by relaxing the assumption of a fixed network topology. We consider time-varying network topologies in which two randomly chosen agents interact at each update step, giving rise to asynchronous updates (i.e., beliefs of all agents are not updated at each update step). We analyze the convergence properties of the two-stage asynchronous decentralized collaborative 20 questions algorithm under appropriate conditions. This asynchronous model is applicable to practical sensor networks where agents may be located in large geometric distances, and as a result their wireless communications are unreliable or intermittent due to path occlusions or other environmental effects, and lead to time-varying network topologies. Our analysis

E-mail addresses: ttsili@ll.mit.edu, ttsili@umich.edu

http://dx.doi.org/10.1016/j.sigpro.2017.06.030
0165-1684/

is based on smoothing techniques and martingale convergence theory in similar spirit to Tsiligkaridis et al. [2]. However, the randomness due to the time-varying network topology and lack of strong connectivity at each time instant introduce additional complications in the analysis that were absent when analyzing the static network case of adaptive agents in [2]. We remark that specific applications and case studies of the proposed methods are out of the scope of this work. This paper is primarily concerned with the convergence analysis of decentralized algorithms for probabilistic bisection-based search in time-varying networks.

In addition to theoretical analysis of the convergence of the proposed algorithm, numerical studies of performance are provided showing interesting information behavior that information sharing yields. The benefit of our asynchronous approach is that the asynchronous decentralized algorithm attains similar performance as its synchronous counterpart introduced and analyzed in [2]. The lack of synchronization our approach offers is of great practical interest because synchronization of a large number of agents can be difficult [6]. Furthermore, the popular time-division-multiple-access (TDMA) communication protocol for distributed networks is only applicable to synchronized networks. Even though the synchronous update scheme in [2] is proven to be convergent to the correct limit, it is still an open question whether the same algorithm converges with an asynchronous implementation. Counterexamples showing that asynchronous updates do not converge although synchronous updates converge are presented in [7] for the standard consensus problem.

1.1. Prior work

The noisy 20 questions problem, also known as Ulam's game, was introduced by Renyi [8] and was later rediscovered by Ulam [9]. The first work making the connection between communication with feedback and noisy search appeared in [10]. The probabilistic bisection algorithm dates back to the work of Horstein [11], where it was originally proposed and analyzed heuristically in the contest of communication with noiseless feedback over the binary symmetric channel. This algorithm was shown to achieve capacity for arbitrary memoryless channels in [12,13]. The probabilistic bisection algorithm was generalized to multiple players in [1] in the centralized setting, and decentralized algorithms for probabilistic bisection search were proposed in [2].

Our work also differs from the works on 20 questions/active stochastic search of Jedynak et al. [14], Castro and Nowak [5], Waeber et al. [15], and Tsiligkaridis et al. [1] because we consider intermediate local belief sharing between agents after each local bisection and update. In addition, in contrast to previous work, in the proposed framework each agent incorporates the beliefs of its neighbors in a way that is agnostic of its neighbors' error probabilities. The analysis of Jadbabaie and co-workers [16,17] does not apply to our model since we consider controlled observations, although we use a form of the social learning model of Jadbabaie and co-workers [16,17]. While a randomized distributed averaging/consensus problem was analyzed in [18], the convergence analysis is not applicable because we consider new information injected in the dynamical system at each iteration (controlled information gathering) in addition to randomized information sharing. The material in this paper was presented in part as a conference [19], and this extended version contains a more thorough presentation, full proofs for convergence analysis, and extensive simulations.

Although consensus to the true target location holds for the degenerate case of no agent collaboration by using results in the existing literature [15], collaboration improves the rate of convergence of the estimation error. As shown in the numerical results in this paper, the error decays faster as a function of iterations. This is the primary motivation for studying such collaborative signal processing algorithms. However, proving convergence to the correct consensus is the first step in analyzing such algorithms, and it is by no means a trivial one. In fact, even for the simple single-agent case, only very recently, Waeber et al. [15] were able to prove the first rigorous convergence rate result for the continuous probabilistic bisection algorithm. The focus of this paper is to establish convergence of decentralized algorithms for probabilistic bisection in time-varying networks.

2. Notation

We define X^* the true parameter, the target state in the sequel, and its domain as the unit interval $\mathcal{X} = [0, 1]$. Let $\mathcal{B}(\mathcal{X})$ be the set of all Borel-measurable subsets $B \subseteq \mathcal{X}$. Let $\mathcal{N} = \{1, \dots, M\}$ index the M agents in an interaction network, denoted by the vertex set $\mathcal{N}$ and the directed edges joining agents at time $t \in \mathbb{N}$ are captured by $E(t)$. Let $\mathbf{A}_t = \{a_{i,j}(t)\}$ denote the interaction matrix at time t, which is a stochastic matrix (i.e., nonnegative entries with rows summing to unity). At each time t, the time-varying network structure is modeled by the directed graph $(\mathcal{N}, E(t))$, where

$$E(t) = \{(j, i) : [\mathbf{A}_t]_{i,j} > 0\}$$

Let $\mathbf{A}_{i \to j}$ denote the interaction matrix when agent i performs a Bayesian update based on its query and averages beliefs with agent j. In our model, a random agent i is chosen with probability q_i and a collaborating agent j is chosen with probability $P_{i,j}$ at each update step. Thus, the interaction matrix $\mathbf{A}_t = \mathbf{A}_{i \to j}$ is chosen with probability $q_i P_{i,j}$, and nodes i and j collaborate. The matrix $\mathbf{P} = \{P_{i,j}\}$ contains the probabilistic weights for collaboration between agents and are zero when there is no edge in the information sharing network at any time; if $P_{i,j} = 0$, then agent i cannot collaborate with agent j at any time.

Define the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ consisting of the sample space Ω generating the unknown state X^* and the observations $\{Y_{i,t+1}\}$ at times $t = 0, 1, \dots,$ an event space $\mathcal{F}$ and a probability measure $\mathbb{P}$. The expectation operator $\mathbb{E}$ is defined with respect to $\mathbb{P}$.

Define the belief of the ith agent at time t on $\mathcal{X}$ as the posterior density $p_{i,t}(x)$ of target state $x \in \mathcal{X}$ based on all of the information available to this agent at this time. Define the $M \times 1$ vector $\mathbf{p}_t(x) = [p_{1,t}(x), \dots, p_{M,t}(x)]^T$ for each $x \in \mathcal{X}$. For any $B \in \mathcal{B}(\mathcal{X})$, define $\mathbf{P}_t(B)$ as the vector of probabilities with i-th element equal to $\int_B p_{i,t}(x)dx$. We define the query point/target estimate of the i-th agent as $\hat{X}_{i,t}$. The query point is the right boundary of the region $A_{i,t} = [0, \hat{X}_{i,t}]$. We let $F_{i,t}(a) = P_{i,t}([0, a]) = \int_0^a p_{i,t}(x)dx$ denote the cumulative distribution function associated with the density $p_{i,t}(\cdot)$.

We assume that a randomly chosen agent i constructs a query at time t of the form "does X^* lie in the region $A_{i,t} \subset \mathcal{X}$?". We indicate this query with the binary variable $Z_{i,t} = I(X^* \in A_{i,t})$ to which agent i responds with a binary response $Y_{i,t+1}$, which is correct with probability $1 - \epsilon_i$, and without loss of generality $\epsilon_i \le 1/2$. This error model is equivalent to a binary symmetric channel (BSC) with crossover probability ϵ_i. The query region $A_{i,t}$ depends on the accumulated information up to time t at agent i. Define the nested sequence of event spaces $\mathcal{F}_t$, $\mathcal{F}_{t-1} \subset \mathcal{F}_t$, for all $t \ge 0$, generated by the sequence of queries and responses. The queries $\{A_{i,t}\}_{t \ge 0}$ are measurable with respect to this filtration. Define the canonical basis vectors $\mathbf{e}_i \in \mathbb{R}^M$ as $[\mathbf{e}_i]_j = I(j = i)$. The notation *i.p.* denotes convergence in probability and *a.s.* denotes almost-sure convergence.

3. Asynchronous decentralized 20 questions

Motivated by the work of Tsiligkaridis et al. [1,2] and Jadbabaie et al. [16], we proceed as follows. As in the fixed-topology decen-