Automatica 84 (2017) 149–158

Contents lists available at ScienceDirect

Automatica

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

Dual decomposition for multi-agent distributed optimization with coupling constraints[☆]

Alessandro Falsone[a], Kostas Margellos[b], Simone Garatti[a], Maria Prandini[a]

[a] Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via Ponzio 34/5, 20133 Milano, Italy
[b] Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, United Kingdom

ARTICLE INFO

Article history:
Received 11 May 2016
Received in revised form 6 March 2017
Accepted 29 May 2017
Available online 17 August 2017

Keywords:
Distributed optimization
Consensus
Dual decomposition
Proximal minimization

ABSTRACT

We study distributed optimization in a cooperative multi-agent setting, where agents have to agree on the usage of shared resources and can communicate via a time-varying network to this purpose. Each agent has its own decision variables that should be set so as to minimize its individual objective function subject to local constraints. Resource sharing is modeled via coupling constraints that involve the non-positivity of the sum of agents' individual functions, each one depending on the decision variables of one single agent. We propose a novel distributed algorithm to minimize the sum of the agents' objective functions subject to both local and coupling constraints, where dual decomposition and proximal minimization are combined in an iterative scheme. Notably, privacy of information is guaranteed since only the dual optimization variables associated with the coupling constraints are exchanged by the agents. Under convexity assumptions, jointly with suitable connectivity properties of the communication network, we are able to prove that agents reach consensus to some optimal solution of the centralized dual problem counterpart, while primal variables converge to the set of optimizers of the centralized primal problem. The efficacy of the proposed approach is demonstrated on a plug-in electric vehicles charging problem.

1. Introduction

This paper addresses optimization in multi-agent networks where each agent aims at optimizing a local performance criterion possibly subject to local constraints, but yet needs to agree with the other agents in the network on the value of some decision variables that refer to the usage of some shared resources.

Cooperative multi-agent decision making problems have been studied recently by many researchers, mainly within the control and operational research communities, and are found in various application domains such as power systems (Bolognani, Carli, Cavraro, & Zampieri, 2015; Zhang & Giannakis, 2016), wireless and social networks (Baingana, Mateos, & Giannakis, 2014; Mateos & Giannakis, 2012), robotics (Martinez, Bullo, Cortez, & Frazzoli, 2007), to name a few.

A possible approach to cooperative multi-agent optimization consists in formulating and solving a mathematical program involving the decision variables, objective functions, and constraints of the entire network. Though this centralized perspective appears sensible, it may end up being impractical for large scale systems for which the computational effort involved in the program solution can be prohibitive. Also, privacy of information is not preserved since agents are required either to share among them or to provide to a central entity their performance criteria and constraints.

Distributed optimization represents a valid alternative to centralized optimization and, in particular, it overcomes the above limitations by allowing agents to keep their information private, while distributing the computational effort. Typically, an iterative procedure is conceived, where at each iteration agents perform some local computation based on their own information and on the outcome of the local computations of their neighboring agents at the previous iteration, till convergence to some solution, possibly an optimal one for the centralized optimization problem counterpart.

Effective distributed optimization algorithms have been proposed in the literature for a general class of convex problems over time-varying, multi-agent networks. In particular, consensus-based optimization algorithms are formulated in Lee and Nedic (2013), Nedic and Ozdaglar (2009b), Nedic, Ozdaglar, and Parrilo

☆ Research was supported by the European Commission under the project UnCoVerCPS, grant number 643921. The material in this paper was partially presented at the 55th IEEE Conference on Decision and Control, December 12–14, 2016, Las Vegas, NV, USA. This paper was recommended for publication in revised form by Associate Editor Michael M. Zavlanos under the direction of Editor Christos G. Cassandras.

E-mail addresses: alessandro.falsone@polimi.it (A. Falsone), kostas.margellos@eng.ox.ac.uk (K. Margellos), simone.garatti@polimi.it (S. Garatti), maria.prandini@polimi.it (M. Prandini).

http://dx.doi.org/10.1016/j.automatica.2017.07.003
0005-1098/

(2010), and Ram, Nedic, and Veeravalli (2012) and in our recent paper Margellos, Falsone, Garatti, and Prandini (2016) for problems where agents have their own objective functions and constraints but decision variables are common.

In this paper, we address a specific class of convex optimization problems over time-varying, multi-agent networks, which we refer to as *inequality-coupled problems* for short-hand notation. In this class of problems, each agent has its own decision vector, objective function, and constraint set, and is coupled to the others via a constraint expressed as the non-positivity of the sum of convex functions, each function corresponding to one agent. We propose a novel distributed iterative scheme based on a combination of dual decomposition and proximal minimization to deal with inequality-coupled problems. Under convexity assumptions and suitable connectivity properties of the communication network, agents reach consensus with respect to the dual variables, without disclosing information about their optimal decision, local objective and constraint functions, nor about the function encoding their contribution to the coupling constraint. The proposed algorithm converges to some optimal dual solution of the centralized problem counterpart, while for the primal variables, we show convergence to the set of optimal primal solutions.

The contributions of our paper versus the existing literature are summarized in the following.

Our scheme can be seen as an extension of dual decomposition based algorithms to a distributed setting, accounting for time-varying network connectivity. As a matter of fact, if the communication networks were time-invariant and connected, then, dual decomposition techniques (see Yang and Johansson (2010), and references therein) as well as approaches based on the alternating direction method of multipliers (Boyd, Parikh, Chu, Peleato, & Eckstein, 2010; Shi, Ling, Yuan, Wu, & Yin, 2014) could be applied to the set-up of this paper, since, after dualizing the coupling constraint, the problem assumes a separable structure. However, in Boyd et al. (2010) and Yang and Johansson (2010) a central update step involving communication among all agents that are coupled via the constraints is required for the dual variables, and this prevents their usage in the distributed with time-varying connectivity set-up of this paper. In Shi et al. (2014) no central update step is needed but the constraints appearing in the dual problem cannot be handled. An interesting distributed dual decomposition based algorithm which overcomes the need for a central node and which is more in line with our scheme has been proposed in Simonetto and Jamali-Rad (2016). The main differences between Simonetto and Jamali-Rad (2016) and our algorithm are as follows:

a. the algorithm of Simonetto and Jamali-Rad (2016) requires that the communication network is time invariant, while our algorithm admits time-variability;
b. in Simonetto and Jamali-Rad (2016) a constant step-size is employed, while our algorithm uses a vanishing step-size. The constant step-size has the advantage of enhancing a faster convergence rate, but, at the same time, convergence to a neighborhood of the optimal is guaranteed only. Our algorithm instead is guaranteed to converge to the optimal solution of the original problem;
c. the algorithm of Simonetto and Jamali-Rad (2016) requires that a Slater point exists and is known to all agents, while existence only is required in our algorithm. This relaxation of the conditions for the applicability of the approach can be crucial in those cases where a Slater point is not a-priori available since the reconstruction of a Slater point in a distributed set-up seems to be as challenging as the original problem and requires extra synchronization among agents.

From another perspective, which is better explained later on in the paper, our approach can be also interpreted as a subgradient based algorithm for the resolution of the dual problem, equipped with an auxiliary sequences that allows one to recover the solution of the primal problem we are interested in. In this respect related contributions are Bertsekas (2011), Bertsekas, Nedic, and Ozdaglar (2003), and Nedic and Bertsekas (2001) where some incremental gradient/subgradient algorithms that can be adopted as an alternative to dual decomposition are proposed. These algorithms, however, require that agents perform updates sequentially, in a cyclic or randomized order, and do not really fit the distributed set-up of this paper. The recent contributions Chang et al. (2014) and Zhu and Martinez (2012) instead present primal–dual subgradient based consensus algorithms that fit our set-up and are comparable to our approach. The main differences are:

d. in Zhu and Martinez (2012) a global knowledge by all agents of the coupling constraint in the primal is required and in both Chang et al. (2014) and Zhu and Martinez (2012) information related to the primal problem is exchanged among agents while the algorithm is running. In the separable set-up of this paper, the agents local information on the primal problem (namely, the value of the local optimization variables, the local objective function, the local constraints, and the contribution of the agent to the coupling constraint) can be regarded as sensitive data and their exchange as in Chang et al. (2014) and Zhu and Martinez (2012) may raise privacy issues. In our algorithm, only the local estimates of the dual variables are exchanged, and this secures maximum privacy among agents;
e. the algorithms of Chang et al. (2014) and Zhu and Martinez (2012) require that a Slater point exists and is known to all agents, while existence only is required in our algorithm. As commented before, requiring the knowledge of a Slater point by the agents can hamper the usability of the algorithm. Moreover, the convergence to the optimal solution in Chang et al. (2014) is guaranteed only when each agents objective function is differentiable;
f. to apply the algorithm of Zhu and Martinez (2012) to our set-up, each agent has to generate local copies of the optimization variables of all the other agents, which then are optimized and exchanged. This often results in an unnecessary increase of the computational and communication efforts, which indeed scale as the number of agents in the network. In our approach instead agents need to optimize the local variables only and exchange the estimate of the dual variables, which are as many as the number of coupling constraints. The required local computational effort is thus much smaller. As for the communication effort, our approach is particularly appealing when the number of coupling constraints is low compared to the overall dimensionality of primal decision variables.

Finally, note that the approaches to distributed optimization in Lee and Nedic (2013), Margellos et al. (2016), Nedic and Ozdaglar (2009b), Nedic et al. (2010) and Ram et al. (2012) which do not resort to any dual problem, can be applied to inequality-coupled problems by introducing a common decision vector collecting all local decision variables. This, however, immediately leads to the drawback of an increased computational and communication effort as discussed in point f above. Moreover, these approaches requires an exchange of information related to the primal, which leads to the privacy issues outlined in point d above.

Table 1 summarizes the comparison between the proposed methodology and the most significant approaches that apply to the same set-up. In the table, algorithms are assessed each against the