D-ADMM: A Communication-Efficient Distributed Algorithm for Separable Optimization

João F. C. Mota, João M. F. Xavier, Pedro M. Q. Aguiar, and Markus Püschel

Abstract—We propose a distributed algorithm, named Distributed Alternating Direction Method of Multipliers (D-ADMM), for solving separable optimization problems in networks of interconnected nodes or agents. In a separable optimization problem there is a private cost function and a private constraint set at each node. The goal is to minimize the sum of all the cost functions, constraining the solution to be in the intersection of all the constraint sets. D-ADMM is proven to converge when the network is bipartite or when all the functions are strongly convex, although in practice, convergence is observed even when these conditions are not met. We use D-ADMM to solve the following problems from signal processing and control: average consensus, compressed sensing, and support vector machines. Our simulations show that D-ADMM requires less communications than state-of-the-art algorithms to achieve a given accuracy level. Algorithms with low communication requirements are important, for example, in sensor networks, where sensors are typically battery-operated and communicating is the most energy consuming operation.

Index Terms—Alternating direction method of multipliers, distributed algorithms, sensor networks.

I. INTRODUCTION

In this paper, we propose a distributed algorithm for solving *separable optimization problems*:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & f_1(x) + f_2(x) + \cdots + f_P(x) \\ \text{subject to} \quad & x \in X_1 \cap X_2 \cap \cdots \cap X_P, \end{aligned} \quad (1)$$

where $x \in \mathbb{R}^n$ is the global optimization variable, and $x^\star$ will denote any solution of (1). As illustrated in Fig. 1, we associate a network of P nodes with problem (1), where only node p has access to its private cost function f_p and to its private set X_p. Each node can only communicate with its neighbors, but all of them have to solve (1) in a cooperative way. We call any method that solves (1) without using a central node and without aggregating data at any specific location a *distributed algorithm*.

Contributions: The goal of this paper is twofold: to show that the recent distributed algorithm proposed in [1] for a specific problem called Basis Pursuit can be generalized to solve the class (1); and to show that, for many problems of interest, the resulting algorithm requires usually significantly less communications than prior distributed algorithms to

J. F. C. Mota is with the Instituto de Sistemas e Robótica (ISR), Instituto Superior Técnico (IST), Technical University of Lisbon, Lisbon 1600–011, Portugal, and also with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: joaomota@cmu.edu).

J. M. F. Xavier, and P. M. Q. Aguiar are with Instituto de Sistemas e Robótica (ISR), Instituto Superior Técnico (IST), Technical University of Lisbon, Portugal.

M. Püschel is with the Department of Computer Science at ETH Zurich, Zurich 8092, Switzerland.
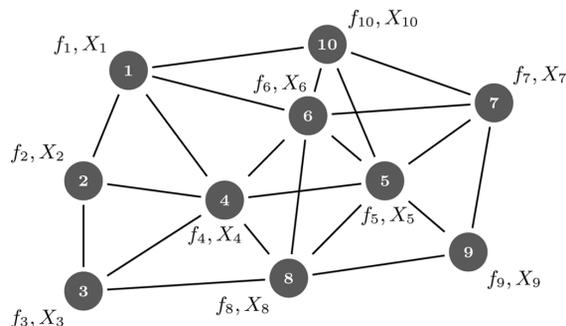
Fig. 1. Network with $P = 10$ nodes. Node p only knows f_p and X_p, but cooperates with its neighbors in order to solve (1).

achieve a given solution accuracy. This also includes algorithms that were specifically designed for a particular problem and are not applicable to the entire problem class (1). Algorithms with low communication cost are relevant, for example, in sensor networks where communication is often the most energy-consuming task and the nodes rely on batteries [2], [3].

Formal Problem Statement: Given a network with P nodes, we associate each f_p and X_p in (1) with the pth node of the network. We make the following assumptions:

Assumptions:
1) Each $f_p : \mathbb{R}^n \to \mathbb{R}$ is a convex function over $\mathbb{R}^n$, and each set X_p is closed and convex.
2) Problem (1) is solvable.
3) The network is connected and it does not vary with time.
4) A coloring scheme of the network is available.

Assumption 2) implies that (1) has at least one solution $x^\star$. In Assumption 3), a network is connected if there is a path between every pair of nodes. Finally, in Assumption 4), a coloring scheme is an assignment of numbers to the nodes of the network such that no adjacent nodes have the same number. These numbers are usually called colors, and they will be used to set up our distributed algorithm. Note that, in wireless scenarios, coloring schemes are often used in *Media Access Control* (MAC) protocols to determine the nodes' order of communication.

Under the previous assumptions, we solve the following problem: *given a network, design a distributed algorithm that solves* (1). By "distributed" we mean there is no notion of a central or special node and each node communicates only with its neighbors; also, only node p has access to f_p or X_p at any time during or before the algorithm.

Our solution for this problem relies on the *Alternating Direction Method of Multipliers* (ADMM), which has become very popular in recent years; see [4] for a survey. Specifically, we use an extended version of ADMM, whose proof of convergence was recently established in [5]. This result will also guarantee the convergence of our algorithm for some problems of interest.

Related Work: Gradient and subgradient methods, including incremental versions, are long known to yield distributed algorithms (in the sense defined before); see, e.g., [6]–[8]. Advantages of these methods are computational simplicity at each node and theoretical robustness guarantees. However, they generally require too many iterations (and hence communications) to converge.

Augmented Lagrangian methods have also been used for distributed optimization, e.g., [9]–[11]. They consist of two loops: an outer loop updating the dual variables, and an inner loop updating the primal variables. The most common outer loop algorithm is the gradient method, yielding the method of multipliers. For the inner loop, common choices are Gauss-Seidel and Jacobi methods.

The Alternating Direction Method of Multipliers (ADMM) [4] is an augmented Lagrangian-based algorithm that consists of only one loop. ADMM is not directly applicable to (1): one has to reformulate that problem first. Possible reformulations were addressed in [12] and [13], yielding algorithms that require two and one communication steps per ADMM iteration, respectively. Other work that explores these algorithms for particular instances of (1) include [14]–[18]. The algorithm we propose is also based on ADMM (on an extended version), but applied to a different reformulation of (1). Our simulations show that the proposed algorithm requires less communications than any of the previous approaches.

All the above algorithms solve (1) in a distributed way. There are, however, other algorithms that solve (1), but are not distributed in our sense. For example, the algorithm in ([4], Section 7.2) solves (1), but it requires an all-to-all communication in each iteration; this can only be accomplished in networks that are fully connected or that have a central node. In contrast, our algorithm and the ones described above are distributed and can run on any connected network topology.

II. ALGORITHM DERIVATION

To derive the algorithm, we reformulate (1) to make ADMM applicable. As mentioned before, several reformulations are possible: ours takes advantage of node coloring. First we introduce some notation.

Network Notation: Networks are represented as undirected graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, 2, \ldots, P\}$ is the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. The cardinality of these sets is represented respectively by P and E. An edge is represented by (i, j), with $i < j$, and $(i, j) \in \mathcal{E}$ means that nodes i and j can exchange data with each other. We define the neighborhood $\mathcal{N}_p$ of a node p as the set of nodes connected to node p, but excluding it; the cardinality of this set, $D_p := |\mathcal{N}_p|$, is the degree of node p.

Coloring: We assume the network is given together with a coloring scheme of C colors. The set of nodes that have color c will be denoted with $\mathcal{C}_c$, for $c = 1, \ldots, C$, and its cardinality with $C_c = |\mathcal{C}_c|$. Note that $\{\mathcal{C}_c\}_{c=1}^C$ partitions $\mathcal{V}$.

Problem Manipulations: Without loss of generality, assume the nodes are ordered such that the first C_1 nodes have color 1, the next C_2 nodes have color 2, and so on, i.e., $\mathcal{C}_1 = \{1, 2, \ldots, C_1\}, \mathcal{C}_2 = \{C_1 + 1, C_1 + 2, \ldots, C_1 + C_2\}, \ldots$. We decouple problem (1) by assigning copies of the global variable x to each node and then constrain all copies to be equal. Let $x_p \in \mathbb{R}^n$ denote the copy held by node p. As in [13], we constrain all copies to be equal in an edge-based way, and rewrite (1) as

$$
\begin{aligned}
&\underset{\bar{x} = (x_1, \ldots, x_P)}{\text{minimize}} && f_1(x_1) + f_2(x_2) + \cdots + f_P(x_P) \\
&\text{subject to} && x_p \in X_p, \ p = 1, \ldots, P, \\
& && x_i = x_j, \ (i, j) \in \mathcal{E}, \qquad\qquad (2)
\end{aligned}
$$

where $\bar{x} = (x_1, \ldots, x_P) \in (\mathbb{R}^n)^P$ is the optimization variable. Problem (2) is no longer coupled by a global variable, as (1), but instead by the new equations $x_i = x_j$, for all the pairs $(i, j) \in \mathcal{E}$. These equations enforce all copies to be equal since the network is connected (cf. Assumption 3)). Note that these constraints can be written more compactly as $(B^\top \otimes I_n)\bar{x} = 0$, where $B \in \mathbb{R}^{P \times E}$ is the node arc-incidence matrix of the graph, I_n is the identity matrix in $\mathbb{R}^n$, and $\otimes$ is the Kronecker product. Each column of B is associated with an edge $(i, j) \in \mathcal{E}$ and has 1 and −1 in the ith and jth entry, respectively; the remaining entries are zeros. Our numbering assumption induces a natural partition of B as $\begin{bmatrix} B_1^\top & B_2^\top & \cdots & B_C^\top \end{bmatrix}^\top$, where the columns of $B_c^\top$ are associated to the nodes with color c. We

partition $\bar{x}$ similarly: $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_C)$, where $\bar{x}_c \in (\mathbb{R}^n)^{C_c}$ collects the copies of all nodes with color c. This enables rewriting (2) as

$$
\begin{aligned}
&\underset{\bar{x}_1, \ldots, \bar{x}_C}{\text{minimize}} && \sum_{c=1}^C \sum_{p \in \mathcal{C}_c} f_p(x_p) \\
&\text{subject to} && \bar{x}_c \in \bar{X}_c, \ c = 1, \ldots, C \\
& && \sum_{c=1}^C (B_c^\top \otimes I_n)\bar{x}_c = 0, \qquad (3)
\end{aligned}
$$

where $\bar{X}_c := \prod_{p \in \mathcal{C}_c} X_p$. Problem (3) can be solved with the Extended ADMM, explained next.

Extended ADMM: The Extended ADMM is a natural generalization of the *Alternating Direction Method of Multipliers* (ADMM) [5]. Given C functions g_c, C sets X_c, and C matrices A_c, all with the same number of rows, the extended ADMM solves

$$
\begin{aligned}
&\underset{x_1, \ldots, x_C}{\text{minimize}} && \sum_{c=1}^C g_c(x_c) \\
&\text{subject to} && x_c \in X_c, \ c = 1, \ldots, C \\
& && \sum_{c=1}^C A_c x_c = 0, \qquad\qquad (4)
\end{aligned}
$$

where $x := (x_1, \ldots, x_C)$ is the optimization variable. The extended ADMM consists of iterating on k:

$$
x_1^{k+1} = \underset{x_1 \in X_1}{\arg\min} \ L_\rho\left(x_1, x_2^k, \ldots, x_P^k; \lambda^k\right) \qquad (5)
$$

$$
x_2^{k+1} = \underset{x_2 \in X_2}{\arg\min} \ L_\rho\left(x_1^{k+1}, x_2, x_3^k, \ldots, x_C^k; \lambda^k\right) \qquad (6)
$$

$$
\vdots
$$

$$
x_C^{k+1} = \underset{x_C \in X_C}{\arg\min} \ L_\rho\left(x_1^{k+1}, x_2^{k+1}, \ldots, x_{C-1}^{k+1}, x_C; \lambda^k\right) \qquad (7)
$$

$$
\lambda^{k+1} = \lambda^k + \rho \sum_{c=1}^C A_c x_c^{k+1}, \qquad\qquad (8)
$$

where $L_\rho(x; \lambda) = \sum_{c=1}^C (g_c(x_c) + \lambda^\top A_c x_c) + \frac{\rho}{2} \|\sum_{c=1}^C A_c x_c\|^2$ is the augmented Lagrangian of (4), λ is the dual variable, and ρ is a positive parameter. When $C = 2$, (5)–(8) becomes the ordinary ADMM and it converges under very mild assumptions. When $C > 2$, there is only a known proof of convergence when all the functions g_c are strongly convex [5]. In particular, the following theorem holds.

Theorem 1 ([19], [5]): Let $g_c : \mathbb{R}^{n_c} \to \mathbb{R}$ be a convex function over $\mathbb{R}^{n_c}$, $X_c \subseteq \mathbb{R}^{n_c}$ a closed convex set, and A_c an $m \times n_c$ matrix, for $c = 1, \ldots, C$. Assume (4) is solvable and that either

  1) $C = 2$ and each A_c has full column-rank,
  2) or $C \geq 2$ and each f_c is strongly convex.

Then, the sequence $\{(x_1^k, \ldots, x_C^k, \lambda^k)\}$ generated by (5)–(8) converges to $(x_1^\star, \ldots, x_C^\star, \lambda^\star)$, where $(x_1^\star, \ldots, x_C^\star)$ solves (4) and $\lambda^\star$ solves the dual problem of (4): $\max_\lambda G_1(\lambda) + \cdots + G_C(\lambda)$, where $G_c(\lambda) = \inf_{x_c \in X_c} (g_c(x_c) + \lambda^\top A_c x_c)$, for $c = 1, \ldots, C$.

A proof for case 1) can be found in [19], which generalizes the proofs of [9], [4]. A proof for case 2) can be found in [5]. It is believed that (5)–(7) still converges when $C > 2$ and each A_c has full column-rank, i.e., that the generalization of Theorem 1 under case 1) still holds [1], [5], [20]. Recently, [21] proved that if we replace ρ in (8) by a small constant, the resulting algorithm converges linearly.

Applying the Extended ADMM: We now apply the extended ADMM to problem (3), which has the format of (4). We start by showing that the cth optimization problem in (5)–(7) yields C_c

---

Algorithm 1: D-ADMM

Initialization: for all $p \in \mathcal{V}$, set $\gamma_p^1 = x_p^1 = 0$ and $k = 1$

1: **repeat**
2:      **for** $c = 1, \ldots, C$ **do**
3:         **for all** $p \in \mathcal{C}_c$ [in parallel] **do**

$$v_p^k = \gamma_p^k - \rho \sum_{\substack{j \in \mathcal{N}_p \\ j < p}} x_j^{k+1} - \rho \sum_{\substack{j \in \mathcal{N}_p \\ j > p}} x_j^k$$

4:           and find

$$x_p^{k+1} = \arg\min \quad f_p(x_p) + v_p^{k\top} x_p + \frac{D_p \rho}{2} \|x_p\|^2$$
$$\text{s.t.} \quad x_p \in X_p$$

5:           Send x_p^{k+1} to $\mathcal{N}_p$
6:         **end for**
7:      **end for**
8:      **for all** $p \in \mathcal{V}$ [in parallel] **do**

$$\gamma_p^{k+1} = \gamma_p^k + \rho \sum_{j \in \mathcal{N}_p} (x_p^{k+1} - x_j^{k+1})$$

9:      **end for**
10:     $k \leftarrow k + 1$
11: **until** some stopping criterion is met

---

optimization problems that can be solved in parallel. For example, $\bar{x}_1 = (x_1, \ldots, x_{C_1})$ is updated as

$$\bar{x}_1^{k+1} = \arg\min_{\bar{x}_1 \in \bar{X}_1} \sum_{p \in \mathcal{C}_1} f_p(x_p) + \lambda^{k\top} A_1 \bar{x}_1$$
$$+ \frac{\rho}{2} \left\| A_1 \bar{x}_1 + \sum_{c=2}^{C} A_c \bar{x}_c^k \right\|^2, \quad (9)$$

where $A_1 = B_1^\top \otimes I_n$. The last term in (9) can be written as

$$\frac{\rho}{2} \bar{x}_1^\top A_1^\top A_1 \bar{x}_1 + \rho \bar{x}_1^\top \sum_{c=2}^{C} A_1^\top A_c \bar{x}_c^k + \frac{\rho}{2} \left\| \sum_{c=2}^{C} A_c \bar{x}_c^k \right\|^2. \quad (10)$$

In the first term, $A_1^\top A_1 = B_1 B_1^\top \otimes I_n$, where $B_1 B_1^\top$ is a diagonal block of the graph Laplacian. Since the nodes with color 1 are not neighbors between themselves, $B_1 B_1^\top$ will be a diagonal matrix, with the degrees of the respective nodes in the diagonal. This means $\bar{x}_1^\top A_1^\top A_1 \bar{x}_1 = \sum_{p \in \mathcal{C}_1} D_p \|x_p\|^2$. Similarly, in the second term, $A_1^\top A_c = B_1 B_c^\top \otimes I_n$, where $B_1 B_c^\top$ corresponds to an off-diagonal block of the Laplacian matrix. For $i \neq j$, the ijth entry of the Laplacian matrix contains -1 if nodes i and j are neighbors, and 0 otherwise. This implies $\bar{x}_1^\top \sum_{c=2}^{C} A_1^\top A_c \bar{x}_c^k = -\sum_{p \in \mathcal{C}_1} \sum_{j \in \mathcal{N}_p} x_p^\top x_j^k$. Finally, the last term of (10) does not depend on $\bar{x}_1$ and can be ignored from the optimization problem. Thus, (9) simplifies to

$$\bar{x}_1^{k+1} = \arg\min_{\bar{x}_1 \in \bar{X}_1} \sum_{p \in \mathcal{C}_1} f_p(x_p) + \left(\gamma_p^k - \rho \sum_{j \in \mathcal{N}_p} x_j^k \right)^\top x_p$$
$$+ \frac{\rho D_p}{2} \|x_p\|^2, \quad (11)$$

where $\gamma_p^k := \sum_{j \in \mathcal{N}_p} \lambda_{pj}^k$ comes from the second term in (9): $((B_1 \otimes I_n)\lambda^k)^\top \bar{x}_1 = \sum_{p \in \mathcal{C}_1} \sum_{j \in \mathcal{N}_p} \lambda_{pj}^{k\top} x_p$. We decomposed λ edge-wise: $\lambda = (\ldots, \lambda_{ij}, \ldots)$, where λ_{ij} is defined for $i < j$ and associated to the constraint $x_i = x_j$ in (2). It is now clear that (11) decomposes into C_1 problems that can be solved in parallel. For the other colors, we can apply a similar reasoning, but we must be careful defining γ_p^k,

due to the nodes' relative numbering. Its general definition is $\gamma_p^k := \sum_{j \in \mathcal{N}_p} \text{sign}(j - p) \lambda_{pj}^k$, where $\text{sign}(a) = 1$, if $a \geq 0$, and $\text{sign}(a) = -1$, otherwise. Note that we extended the definition of λ_{ij} for $i > j$ such that $\lambda_{ij} := \lambda_{ji}$. Algorithm 1 shows the resulting algorithm, named *Distributed-ADMM*, or D-ADMM.

In Algorithm 1, the edge-wise dual variables λ_{ij} were totally replaced by the node-wise dual variables γ_p. This is because the problem in step 4 depends only on γ_p^k and not on the individual λ_{pj}^k's. The update for γ_p in step 8 stems from replacing $\lambda_{ij}^{k+1} = \lambda_{ij}^k + \text{sign}(j - i)(x_i^{k+1} - x_j^{k+1})$ in the definition of γ_p^{k+1}.

Algorithm 1 is asynchronous in the sense that nodes operate in a color-based order, with nodes with the same color operating in parallel. Since nodes with the same color are not neighbors, we would apparently need some kind of coordination to execute the algorithm. Actually, such coordination is not needed provided each node knows its own color and the colors of its neighbors. In fact, as soon as node p has received x_j^{k+1} from all its neighbors with lower colors, node p can "work," since step 4 (and subsequently step 5) can be performed. In conclusion, knowing its own and its neighbors' colors provides an automatic coordination mechanism. Regarding the convergence of D-ADMM, we have:

Corollary 1: Let Assumptions 1)–4) hold. Then, Algorithm 1 produces a sequence $(x_1^k, \ldots, x_P^k)$ convergent to $(x^\star, \ldots, x^\star)$, where $x^\star$ solves (1), whenever 1) the network is bipartite, or 2) each f_p is strongly convex.

Proof: The proof is based on showing that the conditions of Theorem 1 are satisfied. First, note that Assumptions 1) and 2) and the equivalence between (1) and (3) imply that problem (3) is solvable, that each function $\sum_{p \in \mathcal{C}_c} f_p(x_p)$ is convex over $\mathbb{R}^n$, and that each set $\bar{X}_c$ is closed and convex. Now, we will see that Assumption 3) implies that each $B_c^\top \otimes I_n$ has full column-rank. Note that it is sufficient to prove that $B_c^\top$ has full column-rank. If, on the other hand, we prove that $B_c B_c^\top$ has full rank, then the result follows because $\text{rank}(B_c B_c^\top) = \text{rank}(B_c^\top)$. Note that $B_c B_c^\top$ is a diagonal matrix, where the diagonal contains the degrees of the nodes belonging to the subnetwork composed by the nodes in $\mathcal{C}_c$. Since no node has degree 0 (cf. Assumption 3)), $B_c B_c^\top$ has full rank.

Finally, note that a bipartite network can be colored with just two colors. In that case, condition 1) of Theorem 1 is satisfied together with the remaining conditions, which ensures the convergence of Algorithm 1. When the network is non-bipartite and each f_p is strongly convex, we are in case 2) of Theorem 1, which again ensures the convergence of Algorithm 1. ∎

III. APPLICATIONS

We will now see how some important optimization problems can be recast as (1). These reformulations, except the one for LASSO, are not new: see [13]–[16]. Therefore, we refer to these references for the details of solving the optimization problem in step 4 of Algorithm 1. We note that in all the problems, except in consensus, none of the functions f_p is strongly convex. Therefore, D-ADMM is only guaranteed to converge under condition 1) of Corollary 1. Nevertheless, in Section IV, we will see that in practice D-ADMM, not only converges for all these problems, but also outperforms previous work in terms of the number of communications, including the ADMM-based algorithms [13], [12], [16].

Consensus: Consensus is a fundamental problem in networks [18], [22]. Given a network with P nodes, node p generates a number, say θ_p, and the goal is to compute the average $\theta^\star = (1/P) \sum_{p=1}^{P} \theta_p$ at every node. Consensus can be cast as [7], [18]:

$$\underset{x}{\text{minimize}} \quad \frac{1}{2} \sum_{p=1}^{P} (x - \theta_p)^2,$$
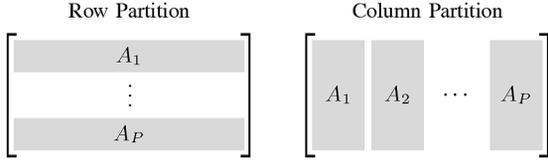
Fig. 2. Row partition and column partition of A into P blocks. A block in the row (resp. column) partition is a set of rows (resp. columns).

which is clearly an unconstrained version of (1), with $f_p(x) = (1/2)(x - \theta_p)^2$; thus, it can be solved with D-ADMM. In this case, the problem of step 4 of Algorithm 1 has a closed-form solution: $x_p^{k+1} = (\theta_p - v_p^k)/(1 + D_p\rho)$.

Sparse Solutions of Linear Systems: Finding sparse solutions of linear systems is important in many areas, including statistics, compressed sensing, and cognitive radio [23], [14]. A common approach to tackle this problem is by solving LASSO [23] or BPDN [24], respectively,

$$\text{LASSO}: \quad \underset{x}{\text{minimize}} \quad \|x\|_1$$
$$\text{subject to} \quad \|Ax - b\| \leq \sigma, \tag{12}$$

$$\text{BPDN}: \quad \underset{x}{\text{minimize}} \quad \|Ax - b\|^2 + \beta\|x\|_1, \tag{13}$$

where the matrix $A \in \mathbb{R}^{m \times n}$, the vector $b \in \mathbb{R}^m$, and the parameters $\sigma, \beta > 0$ are given. LASSO first appeared in [23] to denote a related problem, although the problem in (12) is known by the same name. We solve LASSO and BPDN in two different scenarios, visualized in Fig. 2: *row partition* (resp. *column partition*), where each node stores a block of rows (resp. columns) of A. While in the row partition vector b is partitioned similarly to A, in the column partition we assume all nodes know the full vector b.

We propose solving LASSO with a column partition and BPDN with a row partition. The reverse cases, i.e., LASSO with a row partition and BPDN with a column partition, cannot be trivially recast as (1). However, in our previous work [1], we solved Basis Pursuit (i.e., LASSO with $\sigma = 0$) for both the row and the column partition.

LASSO: Column Partition: Assume A is partitioned by columns, and the pth block is only known at node p. Also, assume vector b, parameter σ, and the number of nodes P are available at all nodes. LASSO in this scenario cannot be directly recast as (1): we will have to do it through duality. However, only solving the ordinary dual of LASSO will not allow us to recover a primal solution afterwards, since its objective is not strictly convex. We thus start by regularizing LASSO, making it strictly convex:

$$\underset{x}{\text{minimize}} \quad \|x\|_1 + \frac{\delta}{2}\|x\|^2$$
$$\text{subject to} \quad \|Ax - b\| \leq \sigma, \tag{14}$$

where $\delta > 0$ is small enough. This regularization is inspired by [25], which establishes exact regularization conditions. By exact, we mean there exists $\bar{\delta} > 0$ such that the solution of (14) is always a LASSO solution, for $\delta \leq \bar{\delta}$. One of these conditions is that the objective is linear and the constraint set is the intersection of a linear system with a closed polyhedral cone. Although LASSO can be recast as

$$\underset{x,t,u,v}{\text{minimize}} \quad 1_n^\top t$$
$$\text{subject to} \quad \|u\| \leq v,$$
$$u = Ax - b, \quad v = \sigma,$$
$$x \leq t, -x \leq t, \tag{15}$$

where $1_n \in \mathbb{R}^n$ is the vector of ones, the closed convex cone $\{(u, v) : \|u\| \leq v\}$ is not polyhedral; thus, there is not a proof of exact regularization for (15). However, experimental results in [25] suggest that exact regularization might occur for non-polyhedral cones. In the simulations discussed in the next section, we solved (14) always with $\delta = 10^{-2}$ and the corresponding solutions never differed more than 0.5% from the "true" solution of LASSO.

We now introduce a variable $y \in \mathbb{R}^m$ in (14) and rewrite it as:

$$\underset{x,y}{\text{minimize}} \quad \sum_{p=1}^{P}(\|x_p\|_1 + \frac{\delta}{2}\|x_p\|^2)$$
$$\text{subject to} \quad \|y\| \leq \sigma$$
$$y = \sum_{p=1}^{P} A_p x_p - b. \tag{16}$$

If we only dualize the last constraint of (16), we get the dual problem of minimizing $\sum_{p=1}^{P} g_p(\lambda)$, where $g_p(\lambda) := \frac{1}{P}(b^\top \lambda + \sigma\|\lambda\|) - \inf_{x_p}(\|x_p\|_1 + (A_p^\top \lambda)^\top x_p + \frac{\delta}{2}\|x_p\|^2)$ is the function associated to node p. This problem is clearly an unconstrained version of (1).

BPDN: Row Partition: In BPDN, A and b are partitioned by rows, with the blocks A_p and b_p stored at node p. In this scenario, BPDN can be readily rewritten as

$$\underset{x}{\text{minimize}} \sum_{p=1}^{P} \left(\|A_p x - b_p\|^2 + \frac{\beta}{P}\|x\|_1 \right), \tag{17}$$

which is an unconstrained version of (1): just make $f_p(x) := \|A_p x - b_p\|^2 + \frac{\beta}{P}\|x\|_1$.

Distributed Support Vector Machines: A *Support Vector Machine* (SVM) is an optimization problem that arises in machine learning in the context of classification and regression ([26] Ch. 7). While there are several possible formulations for an SVM, here we solve ([26], Section 7.1.1)

$$\underset{s,r,\xi}{\text{minimize}} \quad \frac{1}{2}\|s\|^2 + \beta 1_K^\top \xi$$
$$\text{subject to} \quad y_k(s^\top x_k - r) \geq 1 - \xi_k, \quad k = 1, \dots, K,$$
$$\xi \geq 0, \tag{18}$$

where the parameter $\beta > 0$ and the pairs $(x_k, y_k), k = 1, \dots, K$, are given. Each point x_k belongs to one of two classes: $y_k = 1$ or $y_k = -1$. The goal in solving (18) is to find an hyperplane $\{x \in \mathbb{R}^n : s^\top x = r\}$ that best separates the two classes. The optimization variables in (18) are $s \in \mathbb{R}^n$, the vector orthogonal to the hyperplane, $r \in \mathbb{R}$, the hyperplane offset, and $\xi \in \mathbb{R}^K$, the vector of slack variables. We assume that K is divisible by the number of nodes P, and that each node knows $m := K/P$ pairs of points (x_k, y_k). The resulting problem can be formulated as an unconstrained version of (1) by setting

$$f_p(s, r) = \inf_{\bar{\xi}_p} \frac{1}{2P}\|s\|^2 + \beta 1_m^\top \bar{\xi}_p$$
$$\text{s.t. } Y_p(X_p s - r 1_m) \geq 1_m - \bar{\xi}_p$$
$$\bar{\xi}_p \geq 0,$$

where Y_p is a diagonal matrix with the y_k's corresponding to node p in the diagonal, and X_p is an $m \times n$ matrix with each row containing $x_k^\top$. Note that the size of the variable to be transmitted among the nodes is $n + 1$, corresponding to the size of the global variable (s, r): the variables $\bar{\xi}_p$ are internal to each node.

IV. SIMULATION RESULTS

This section shows simulation results of D-ADMM and related algorithms solving the problems presented in the previous section. We focus on the ADMM-based algorithms [12] and [13], since these are the best
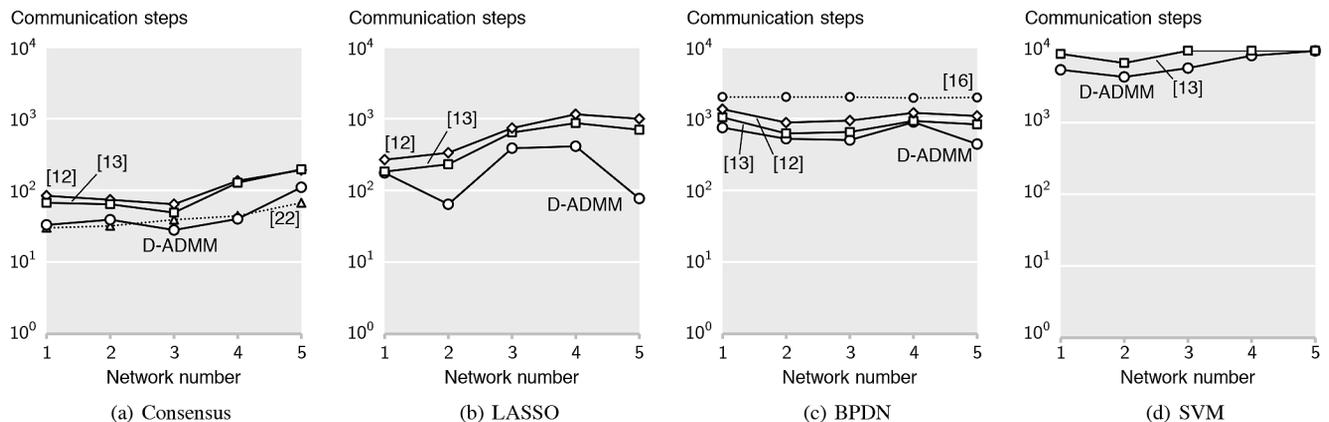
Fig. 3.   Results of the simulations for (a) consensus, (b) LASSO, (c) BPDN, and (d) SVM.

among the distributed algorithms for (1). We start by discussing the performance measure.

Performance Measure: Communication Steps: We say that a *Communication Step* (CS) has occurred when all the nodes have transmitted a vector of size n to its neighbors. All the algorithms we consider here, including D-ADMM, consist of one iterative loop. One iteration of D-ADMM, as well as of [13], corresponds to one CS; one iteration of [12] corresponds to two CSs, since each node transmits two vectors of size n per iteration. The number of CSs is proportional to the number of total communications. Thus, in a wireless scenario, the smaller the number of CSs, the lower the energy consumption.

Note, however, that the CS measure does not take into account the computational complexity at each node. (Actually, D-ADMM, and the algorithms in [12], [13] have similar computational complexities.) Also, this measure is not necessarily related with execution time. In fact, while D-ADMM requires less CSs than competing algorithms (as we will see), it may be slower than some of them. The reason is because D-ADMM is asynchronous, while all the other algorithms are synchronous. Scenarios allowing synchronous transmissions are, however, limited to very controlled environments, such as computer clusters or super-computers, where approaches like ([4], Section 7.2) would probably be more appropriate than distributed algorithms. On the other hand, in wireless networks, the single fact that one node cannot transmit and receive messages at the same time forces synchronous algorithms to operate asynchronously.

Experimental Setup: We generated 5 networks with $P = 50$ nodes according to the models of Table I; see [1] for a description of these models. Table I also gives the number of colors for each network. The results of our simulations are in Fig. 3, where each plot depicts the number of CSs as a function of the network. Having computed the solution $x^\star$ of (1) beforehand and in a centralized way, each algorithm stopped whenever $\|x^k - x^\star\|/\|x^\star\| \le \epsilon$, or when a maximum number of M CSs were reached. In the case of consensus, BPDN, and SVM, x^k denotes the estimate of $x^\star$ at an arbitrary node; in the case of LASSO, it represents the global estimate of the network, since each node only estimates some components of $x^\star$. We used the following values for the pair (ϵ, M): $(10^{-4}, 10^3)$ for consensus, $(5 \times 10^{-3}, 10^3)$ for LASSO, $(10^{-4}, 2 \times 10^3)$ for BPDN, and $(10^{-3}, 10^3)$ for SVM. Since the problem in step 4 of Algorithm 1 does not have a closed-form solution for all the applications we consider, except for consensus, it has to be solved iteratively in each of the algorithms we compare. To make a fair comparison in terms of CSs, we thus use the same solver in all the algorithms, i.e., the problem in step 4 of Algorithm 1 is solved with the same precision in all the algorithms we compare.

TABLE I
NETWORK MODELS

Network	Model (parameters)	# Colors
1	Erdős-Rényi (0.12)	5
2	Watts-Strogatz $(4, 0.4)$	4
3	Barabasi (2)	3
4	Geometric (0.23)	10
5	Lattice (5×10)	2

It is known that the parameter ρ affects strongly the performance of ADMM-based algorithms. Hence, to make a fair comparison, we ran each (ADMM-based) algorithm for several values of ρ and chose always the best result, i.e., the smaller number of CSs. The values for ρ were taken from the set $\{10^{-4}, 10^{-3} 10^{-2} 10^{-1}, 1, 10, 10^2\}$.

Consensus: For the consensus problem, we generated each θ_p randomly from a Gaussian distribution: $\theta_p \overset{\text{i.i.d.}}{\sim} \mathcal{N}(10, 10^4)$. Fig. 3(a) shows the results for D-ADMM, the ADMM-based algorithms [12], [13], and the algorithm [22], which is considered to be the fastest consensus algorithm [18]. Note that [22] was designed for consensus only and cannot be easily generalized to solve (1). Fig. 3(a) shows that D-ADMM has a performance very similar to that of [22].

LASSO and BPDN: The matrix A for the problems LASSO and BPDN was taken from problem 902 of the Sparco toolbox [27]. The vector b was generated as $b = As + n$, where s is a sparse vector and n is Gaussian noise. We chose $\sigma = 0.1$ and $\beta = 0.3$ for the noise parameters, and $\delta = 10^{-2}$ for the approximation parameter in LASSO. The results of these experiments for LASSO and BPDN are shown, respectively, in Fig. 3(b) and 3(c). Additionally, we show the performance of Algorithm 3 of [16], which is an ADMM-based algorithm specifically designed to solve BPDN. This algorithm has the advantage of requiring much simpler computations at each node, but in our simulations it achieved the maximum number of CSs in all but the last two networks. In both LASSO and BPDN, D-ADMM was always the algorithm requiring fewer CSs to converge.

SVM: For the SVM problem (18), we used data from [28], namely two overlapping sets of points from the Iris dataset. The parameter β was always set to 1. In this case, the algorithm from [12] achieved always the maximum number of CSs and thus is not represented in Fig. 3(d), which shows the simulation results. Again, we see that D-ADMM was the algorithm requiring the smallest number of CSs to converge.

V. CONCLUSION

We proposed an algorithm for solving separable problems in networks, in a distributed way. Each node has a private cost and a private constraint set, but all nodes cooperate to solve the optimization problem that minimizes the sum of all costs and that has the intersection of all sets as a constraint. Our algorithm hinges on a coloring scheme of the network, according to which the nodes operate asynchronously. This results in an algorithm with fewer communication requirements than previous algorithms, as shown experimentally for several problems. Although we proved the convergence of the algorithm, it still remains an open question to explain theoretically why the algorithm is more efficient than previous algorithms.

REFERENCES

[1] J. Mota, J. Xavier, P. Aguiar, and M. Püschel, "Distributed basis pursuit," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 1942–1956, 2012.

[2] I. Akyildiz, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: A survey," *Comput. Netw.*, vol. 38, pp. 393–422, 2002.

[3] P. Fischione, P. Park, and K. Johansson*, Wireless Network Based Control, Chapter Design Principles of Wireless Sensor Network Protocols for Control Applications*. New York, NY, USA: Springer, 2011.

[4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.

[5] D. Han and X. Yuan, "A note on the alternating direction method of multipliers," *J. Optim. Theory Appl.*, pp. 227–238, 2012.

[6] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. AC-31, no. 9, pp. 803–812, 1986.

[7] M. Rabbat and R. Nowak, "Distributed optimization in sensor networks," in *Proc. IPSN*, 2004, pp. 20–27.

[8] P. Palomar and Y. Eldar*, Convex Optimization in Signal Processing and Communications*. Cambridge, U.K.: Cambridge Univ. Press, 2010.

[9] D. Bertsekas and J. Tsitsiklis*, Parallel and Distributed Computation: Numerical Methods*. Belmont, MA, USA: Athena Scientific, 1997.

[10] D. Jakovetić, J. Xavier, and J. Moura, "Cooperative convex optimization in networked systems: Augmented lagrangian algorithms with directed gossip communication," *IEEE Trans. Signal Process.*, vol. 59, no. 8, pp. 3889–3902, 2011.

[11] M. Rabbat, R. Nowak, and J. Bucklew, "Generalized consensus algorithms in networked systems with erasure links," in *Proc. IEEE Workshop Signal Process. Adv. Wireless Commun.*, 2005, pp. 1088–1092.

[12] I. Schizas, A. Ribeiro, and G. Giannakis, "Consensus in *ad hoc* wsns with noisy links—Part I: Distributed estimation of deterministic signals," *IEEE Trans. Signal Process.*, vol. 56, no. 1, pp. 350–364, 2008.

[13] H. Zhu, G. Giannakis, and A. Cano, "Distributed in-network channel decoding," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 3970–3983, 2009.

[14] J. Bazerque and G. Giannakis, "Distributed spectrum sensing for cognitive radio networks by exploiting sparsity," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pt. 2, pp. 1847–1862, 2010.

[15] P. Forero, A. Cano, and G. Giannakis, "Consensus-based distributed support vector machines," *J. Mach. Learn. Res.*, vol. 11, pp. 1663–1707, 2010.

[16] G. Mateos, J. Bazerque, and G. Giannakis, "Distributed sparse linear regression," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5262–5276, 2010.

[17] J. Bazerque, G. Mateos, and G. Giannakis, "Group-Lasso on splines for spectrum cartography," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4648–4663, 2011.

[18] T. Erseghe, D. Zennaro, E. Dall'Anese, and L. Vangelista, "Fast consensus by the alternating direction multipliers method," *IEEE Trans. Signal Process.*, vol. 59, no. 11, pp. 5523–5537, 2011.

[19] J. Mota, J. Xavier, P. Aguiar, and M. Püschel, "A proof of convergence for the alternating direction method of multipliers applied to polyhedral constrained functions," 2011 [Online]. Available: http://arxiv.org/abs/1112.2295

[20] B. He, M. Tao, and X. Yuan, "Alternating direction method with Gaussian back substitution for separable convex programming," *SIAM J. Optim.*, vol. 22, no. 2, pp. 313–340, 2012.

[21] Z. Luo, "On the linear convergence of the alternating direction method of multipliers," 2012 [Online]. Available: arxiv.org/abs/1208.3922

[22] B. Oreshkin, M. Coates, and M. Rabbat, "Optimization and analysis of distributed averaging with short node memory," *IEEE Trans. Signal Process.*, vol. 58, no. 5, pp. 2850–2865, 2010.

[23] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. R. Statist. Soc. B*, vol. 58, no. 1, pp. 267–288, 1996.

[24] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.

[25] M. Friedlander and P. Tseng, "Exact regularization of convex programs," *SIAM J. Optim.*, vol. 18, no. 4, pp. 1326–1350, 2007.

[26] C. Bishop*, Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.

[27] E. Berg, M. Friedlander, G. Hennenfent, F. Herrmann, R. Saab, and Ö. Yilmaz, Sparco: A Testing Framework for Sparse Reconstruction Dept. Comput. Sci., Univ. of British Columbia, Vancouver, BC, Canada, Tech. Rep., 2007.

[28] A. Frank and A. Asuncion, UCI Machine Learning Repository, Univ. of California, Irvine, CA, USA, 2010.