

Optimal Scaling of a Gradient Method for Distributed Resource Allocation¹

L. XIAO² AND S. BOYD³

Communicated by P. Tseng

Published Online: 29 November 2006

Abstract. We consider a class of weighted gradient methods for distributed resource allocation over a network. Each node of the network is associated with a local variable and a convex cost function; the sum of the variables (resources) across the network is fixed. Starting with a feasible allocation, each node updates its local variable in proportion to the differences between the marginal costs of itself and its neighbors. We focus on how to choose the proportional weights on the edges (scaling factors for the gradient method) to make this distributed algorithm converge and on how to make the convergence as fast as possible.

We give sufficient conditions on the edge weights for the algorithm to converge monotonically to the optimal solution; these conditions have the form of a linear matrix inequality. We give some simple, explicit methods to choose the weights that satisfy these conditions. We derive a guaranteed convergence rate for the algorithm and find the weights that minimize this rate by solving a semidefinite program. Finally, we extend the main results to problems with general equality constraints and problems with block separable objective function.

Key Words. Distributed optimization, resource allocations, weighted gradient methods, convergence rates, semidefinite programming.

1. Introduction

We consider an optimal resource allocation problem over a network of autonomous agents. The network is modeled as a directed graph $(\mathcal{V}, \mathcal{E})$ with node

¹The authors are grateful to Professor Paul Tseng and the anonymous referee for their valuable comments that helped us to improve the presentation of this paper.

²Postdoctoral Scholar, Center for the Mathematics of Information, California Institute of Technology, Pasadena, California.

³Professor, Department of Electrical Engineering, Stanford University, Stanford, California.

set $\mathcal{V} = \{1, \dots, n\}$ and edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. Each edge (i, j) is an ordered pair of distinct nodes. We define \mathcal{N}_i , the set of (oriented) neighbors of node i , as

$$\mathcal{N}_i = \{j \mid (i, j) \in \mathcal{E}\};$$

in other words, $j \in \mathcal{N}_i$ if there is an edge from node i to node j .

With node i , we associate a variable $x_i \in \mathfrak{R}$ and a corresponding convex cost function $f_i : \mathfrak{R} \rightarrow \mathfrak{R}$. We consider the following optimization problem:

$$\min \sum_{i=1}^n f_i(x_i), \tag{1a}$$

$$\text{s.t.} \quad \sum_{i=1}^n x_i = c, \tag{1b}$$

where $c \in \mathfrak{R}$ is a given constant. We can think of x_i as the amount of some resource located at node i and interpret $-f_i$ as the local (concave) utility function. The problem (1) is to fund an allocation of the resource that maximizes the total utility $\sum_{i=1}^n -f_i(x_i)$. In this paper, we are interested in distributed algorithms for solving this problem, where each node is only allowed to communicate with its neighbors and conduct local computation. Thus, the local information structure imposed by the graph should be considered as part of the problem formulation. This simple model and its variations have many applications in economic systems (e.g., Refs. 1–2) and distributed computer systems (Ref. 3).

We assume that the functions f_i are convex and twice continuously differentiable with second derivatives that are bounded below and above,

$$l_i \leq f_i''(x_i) \leq u_i, \quad x_i \in \mathfrak{R}, \quad i = 1, \dots, n, \tag{2}$$

where $l_i > 0$ and u_i are known (the functions are strictly convex). Let $x = (x_1, \dots, x_n) \in \mathfrak{R}^n$ denote the vector of the variables and let

$$f(x) = \sum_{i=1}^n f_i(x_i)$$

denote the objective function. We use f^* to denote the optimal value of this problem; i.e.,

$$f^* = \inf\{f(x) \mid \mathbf{1}^T x = c\},$$

where $\mathbf{1}$ denotes the vector with all components one. Under the above assumption, the convex optimization problem (1) has a unique optimal solution x^* . Let

$$\nabla f(x) = (f_1'(x_1), \dots, f_n'(x_n))$$

denote the gradient of f at x . The optimality conditions for this problem are

$$\mathbf{1}^T x^* = c, \quad \nabla f(x^*) = p^* \mathbf{1}, \tag{3}$$

where p^* is the (unique) optimal Lagrange multiplier.

In a centralized setup, many methods can be used to solve the problem (1) or equivalently the optimality conditions (3). If the functions f_i are all quadratic, the optimality conditions (3) are a set of linear equations in x^*, p^* and can be solved directly. More generally, the problem can be solved by iterative methods, e.g., the projected gradient method, Newton’s method, quasi-Newton methods (e.g. BFGS method); see e.g. Refs. 4–5.

The design of decentralized mechanisms for resource allocation has a long history in economics (Ref. 6) and there are two main classes of mechanisms: price-directed (Ref. 1) and resource-directed (Ref. 2). However, most of the methods are not fully distributed because they either need a central price coordinator or need a central resource dispatcher. In this paper, we will focus on a class of center-free algorithms first proposed in Ref. 7.

1.1. Center-Free Algorithm for Resource Allocation. Assume that we have an initial allocation of the resource $x(0)$ that satisfies $\underline{1}^T x(0) = c$. The center-free algorithm for solving problem (1) has the following iterative form:

$$x_i(t + 1) = x_i(t) - W_{ii} f'_i(x_i(t)) - \sum_{j \in \mathcal{N}_i} W_{ij} f'_j(x_j(t)), \quad i = 1, \dots, n, \quad (4)$$

for $t = 0, 1, \dots$. In other words, at each iteration, each node computes the derivative of its local function, queries the derivative values from its neighbors, and then updates its local variable by a weighted sum of the values of the derivatives. Here, W_{ii} is the self-weight at node i , and $W_{ij}, j \in \mathcal{N}_i$, is the weight associated with the edge $(i, j) \in \mathcal{E}$. Setting $W_{ij} = 0$ for $j \notin \mathcal{N}_i$, this algorithm can be written in vector form as

$$x(t + 1) = x(t) - W \nabla f(x(t)), \quad (5)$$

where $W \in \mathfrak{R}^{n \times n}$ is the weight matrix. Thus, this algorithm can be thought of as a weighted gradient descent method, in which the weight matrix W has a sparsity constraint,

$$W \in \mathcal{S} = \{Z \in \mathfrak{R}^{n \times n} | Z_{ij} = 0, \text{ if } i \neq j \text{ and } (i, j) \notin \mathcal{E}\}. \quad (6)$$

Throughout this paper, we focus on the following question: How should we choose W ?

We consider first two basic requirements on W . First, we require that all iterates $x(t)$ of the algorithm are feasible, i.e., satisfy $\underline{1}^T x(t) = c$ for all t . With the assumption that $x(0)$ is feasible, this requirement will be met provided the weight matrix satisfies

$$\underline{1}^T W = 0, \quad (7)$$

since we then have

$$\underline{1}^T x(t + 1) = \underline{1}^T x(t) - \underline{1}^T W \nabla f(x(t)) = \underline{1}^T x(t).$$

We will also require, naturally, that the optimal point x^* is a fixed point of (5), i.e.,

$$x^* = x^* - W \nabla f(x^*) = x^* - p^* W \underline{1}.$$

This will hold in the general case (with $p^* \neq 0$) provided

$$W \underline{1} = 0. \tag{8}$$

The requirements (7) and (8) show that the vector $\underline{1}$ must be both a left and right eigenvector of W associated with the eigenvalue zero. One special case of interest is when the weight matrix W is symmetric. In this case, of course, the requirements (7) and (8) are the same and simply state that $\underline{1}$ is in the null space of W .

Assuming that the weights satisfy (8), we have

$$W_{ii} = - \sum_{j \in \mathcal{N}_i} W_{ij},$$

which can be substituted into equation (4) to get

$$x_i(t + 1) = x_i(t) - \sum_{j \in \mathcal{N}_i} W_{ij} (f'_j(x_j(t)) - f'_i(x_i(t))), \quad i = 1, \dots, n. \tag{9}$$

Thus, the change in the local variable at each step is given by a weighted sum of the differences between its own derivative value and those of its neighbors. The equation (9) has a simple interpretation: at each iteration, each connected pair of nodes shifts resources from the node with higher marginal cost to the one with lower marginal cost, in proportion to the difference in marginal costs. The weight $-W_{ij}$ gives the proportionality constant on the edge $(i, j) \in \mathcal{E}$.

Remark 1.1. This interpretation suggests that the edge weights should be negative, but we will see examples where a few positive edge weights actually enhance the convergence rate.

1.2. Previous Work. Distributed resource allocation algorithms of the form (9) were first proposed and studied by Ho, Servi, and Suri in Ref. 7. They considered an undirected graph with symmetric weights on the edges and called algorithms of this form center-free algorithms (the term center-free refers to the absence of a central coordinating entity). In the notation of this paper, they assumed

$$W = W^T, \quad W \underline{1} = 0,$$

and derived the following additional conditions on W that are sufficient for the algorithm (9) to converge to the optimal solution x^* :

$$(a) \text{ } W \text{ is irreducible,} \quad (10a)$$

$$(b) \text{ } W_{ij} \leq 0, \quad (i, j) \in \mathcal{E}, \quad (10b)$$

$$(c) \sum_{j \in \mathcal{N}_i} |W_{ij}| < 1/u_{\max}, \quad i = 1, \dots, n, \quad (10c)$$

where u_{\max} is an upper bound on the second derivatives of the functions f_i , i.e., $u_{\max} \geq \max_i u_i$. The first condition, that W is irreducible, is equivalent to the statement that the subgraph consisting of all the nodes and edges with nonzero weights is connected. We will show that these conditions are implied by those established in this paper.

It should be noted that the problem considered in Ref. 7 has nonnegativity constraints on the variables: $x_i \geq 0$, $i = 1, \dots, n$, with $c > 0$. They gave a separate initialization procedure, which identifies and eliminates some nodes that will have zero value at optimality (not necessarily all such nodes). As a result of this initialization procedure and some additional conditions on the initial point $x(0)$, all the following iterates of the center-free algorithm (9) automatically satisfy the nonnegativity constraints. In Ref. 3, the second derivatives of the functions f_i are used to modify the algorithm (9) (with a constant weight on all edges) to obtain faster convergence. An interesting analogy between various iterative algorithms for solving problem (1) and the dynamics of several electrical networks can be found in Ref. 8.

Many interesting similarities exist between the resource allocation problem and network flow problems with convex separable cost (see e.g. Refs. 9–11 and references therein). In particular, by ignoring the local information structure, problem (1) can be formulated as a simple network flow problem with two nodes and n links connecting them. Thus, many distributed algorithms for network flow problems such as those in Refs. 12–13 can be used; see also Ref. 14 for a convergence rate analysis of such an algorithm. However, with the imposed local information structure on a graph, the above mentioned algorithms cannot be applied directly. The center-free algorithm considered in this paper belongs to a more general class of gradient-like algorithms studied in Ref. 15.

In this paper, we give sufficient conditions weaker than (10) for the center-free algorithm to converge and optimize the edge weights to get fast convergence. Our method is closely related to the approach in Ref. 16, where the problem of finding the fastest mixing Markov chain on a graph is considered. In Ref. 17, the same approach was used to find fast linear iterations for a distributed average consensus problem.

1.3. Outline. In Section 2, we give sufficient conditions on the weight matrix W under which the algorithm (5) converges to the optimal solution monotonically.

These conditions involve a linear matrix inequality (LMI) in the weight matrix. Moreover, we quantify the convergence by deriving a guaranteed convergence rate for the algorithm. In Section 3, we give some simple, explicit choices for the weight matrix W that satisfy the convergence conditions. In Section 4, we propose to minimize the guaranteed convergence rate obtained in Section 2 in order to get fast convergence of the algorithm (5). We observe that the optimal weights (in the sense of minimizing the guaranteed convergence rate) can be found by solving a semidefinite program (SDP). In Section 5, we show some numerical examples that demonstrate the effectiveness of the proposed weight selection methods. Finally, in Section 6, we extend the main results to problems with general equality constraints and problems with block separable objective functions. We give our conclusions and some final remarks in Section 7.

2. Convergence Conditions

In this section, we state and prove the main theorem. We use the following notation: L and U denote diagonal matrices in $\mathfrak{R}^{n \times n}$ whose diagonal entries are the lower bounds l_i and upper bounds u_i given in (2). Note that L and U are positive definite ($L \succ 0$ and $U \succ 0$). For a symmetric matrix Z , we list its eigenvalues (all real) in nonincreasing order,

$$\lambda_1(Z) \geq \lambda_2(Z) \geq \dots \geq \lambda_n(Z),$$

where $\lambda_i(Z)$ denotes the i th largest eigenvalue of Z .

Theorem 2.1. If the weight matrix W satisfies $\underline{1}^T W = 0$, $W \underline{1} = 0$, and

$$\lambda_{n-1}(L^{1/2}(W + W^T - W^T U W)L^{1/2}) > 0, \tag{11}$$

then the algorithm (5) converges to the optimal solution x^* of problem (1) and the objective function decreases monotonically. In fact, we have

$$f(x(t)) - f^* \leq \eta(W)^t [f(x(0)) - f^*] \tag{12}$$

with a guaranteed convergence rate

$$\eta(W) = 1 - \lambda_{n-1}(L^{1/2}(W + W^T - W^T U W)L^{1/2}). \tag{13}$$

Moreover, the condition (11) is equivalent to the strict linear matrix inequality (LMI)

$$\begin{bmatrix} W + W^T + (1/n)\underline{1}\underline{1}^T & W \\ W^T & U^{-1} \end{bmatrix} \succ 0. \tag{14}$$

Proof. Let

$$\Delta x(t) = x(t + 1) - x(t).$$

From (5), we have

$$\Delta x(t) = -W \nabla f(x(t)).$$

The Taylor expansion of f_i at $x_i(t)$ yields

$$f_i(x_i(t + 1)) = f_i(x_i(t)) + f'_i(x_i(t))\Delta x_i(t) + (1/2)f''_i(z_i(t))\Delta x_i(t)^2,$$

where $z_i(t)$ is between $x_i(t)$ and $x_i(t + 1)$. Let

$$\nabla^2 f(x) = \text{diag}(f''_1(x_1), \dots, f''_n(x_n))$$

be the Hessian matrix of f at x . The objective function at iterate $t + 1$ can be written as

$$\begin{aligned} f(x(t + 1)) &= f(x(t)) + \nabla f(x(t))^T \Delta x(t) + (1/2)\Delta x(t)^T \nabla^2 f(z(t))\Delta x(t) \\ &= f(x(t)) - \nabla f(x(t))^T W \nabla f(x(t)) \\ &\quad + (1/2)\nabla f(x(t))^T W^T \nabla^2 f(z(t))W \nabla f(x(t)) \\ &= f(x(t)) - (1/2)\nabla f(x(t))^T [W + W^T - W^T \nabla^2 f(z(t))W] \\ &\quad \times \nabla f(x(t)). \end{aligned}$$

Using the assumption (2), we have

$$\nabla^2 f(z(t)) \preceq U,$$

so

$$\begin{aligned} f(x(t + 1)) &\leq f(x(t)) - (1/2)\nabla f(x(t))^T (W + W^T - W^T U W) \nabla f(x(t)) \\ &= f(x(t)) - (1/2)\nabla f(x(t))^T L^{-1/2} V L^{-1/2} \nabla f(x(t)), \end{aligned} \tag{15}$$

where

$$V = L^{1/2}(W + W^T - W^T U W)L^{1/2}. \tag{16}$$

From conditions (7) and (8), i.e., $\underline{1}^T W = 0$ and $W \underline{1} = 0$, we conclude that $L^{-1/2} \underline{1}$ is an eigenvector of the symmetric matrix V associated with the eigenvalue zero. Let $e(t)$ be the projection of $L^{-1/2} \nabla f(x(t))$ onto the subspace that is orthogonal to $L^{-1/2} \underline{1}$,

$$\begin{aligned} e(t) &= [I - (1/\underline{1}^T L^{-1} \underline{1})L^{-1/2} \underline{1} \underline{1}^T L^{-1/2}]L^{-1/2} \nabla f(x(t)) \\ &= L^{-1/2} \left[\nabla f(x(t)) - \frac{\underline{1}^T L^{-1} \nabla f(x(t))}{\underline{1}^T L^{-1} \underline{1}} \underline{1} \right]. \end{aligned}$$

Replacing $L^{-1/2} \nabla f(x(t))$ by $e(t)$ does not change the equation (15), so we have

$$f(x(t + 1)) \leq f(x(t)) - (1/2)e(t)^T V e(t). \tag{17}$$

Now the condition (11), i.e., $\lambda_{n-1}(V) > 0$ means that the matrix V is positive semidefinite and $\lambda_n(V) = 0$ is a simple eigenvalue, with associated eigenvector $L^{-1/2}\underline{1}$. This implies that

$$f(x(t + 1)) < f(x(t)),$$

i.e., that the objective function decreases provided $e(t) \neq 0$. When $e(t) = 0$, the gradient $\nabla f(x(t))$ must be a multiple of $\underline{1}$, i.e., the optimality conditions (3) are satisfied; in this case, we conclude that $x(t)$ is optimal.

Next, we derive the guaranteed convergence rate (13). Note that, with the condition (11) and the inequality (17), we have

$$f(x(t + 1)) - f(x(t)) \leq -(1/2)\lambda_{n-1}(V)\|e(t)\|^2. \tag{18}$$

We shall derive another inequality relating $f(x(t))$, f^* , and $\|e(t)\|^2$. To do so, we use again the Taylor expansion of f at $x(t)$. Using the assumption (2), we have

$$f(y) \geq f(x(t)) + \nabla f(x(t))^T(y - x(t)) + (1/2)(y - x(t))^T L(y - x(t)),$$

for any $y \in \mathfrak{R}^n$. If y is feasible (true for any iterate of the proposed algorithm), then

$$\underline{1}^T(y - x(t)) = 0.$$

So, we have

$$\begin{aligned} f(y) &\geq f(x(t)) + \left[\nabla f(x(t)) - \frac{\underline{1}^T L^{-1} \nabla f(x(t))}{\underline{1}^T L^{-1} \underline{1}} \underline{1} \right]^T (y - x(t)) \\ &\quad + (1/2)(y - x(t))^T L(y - x(t)) \\ &= f(x(t)) + e(t)^T L^{1/2}(y - x(t)) + (1/2)(y - x(t))^T L(y - x(t)) \\ &= f(x(t)) + e(t)^T z + (1/2)z^T z, \end{aligned}$$

where

$$z = L^{1/2}(y - x(t)).$$

Minimizing the right-hand side over z yields

$$f(y) \geq f(x(t)) - (1/2)\|e(t)\|^2.$$

Since this is true for all feasible y , it is certainly true for x^* . In other words, we have

$$f^* \geq f(x(t)) - (1/2)\|e(t)\|^2. \tag{19}$$

Now combining the inequalities (18) and (19) yields

$$f(x(t + 1)) - f^* \leq [1 - \lambda_{n-1}(V)][f(x(t)) - f^*], \tag{20}$$

which gives the desired results (12) and (13). Since all $x(t)$ are feasible, we always have

$$f(x(t)) - f(x(t + 1)) \leq f(x(t)) - f^*.$$

Applying the two inequalities (18) and (19) on two sides of the above inequality respectively yields $\lambda_{n-1}(V) \leq 1$, hence $0 < \eta(W) < 1$. It follows that

$$\lim_{t \rightarrow \infty} f(x(t)) = f^*,$$

i.e., the algorithm converges to the optimal solution under condition (11).

It remains to show that the eigenvalue condition (11) is equivalent to the strict LMI (14). Since L is diagonal and positive definite, $\lambda_{n-1}(V) > 0$ if and only if

$$\lambda_{n-1}(W + W^T - W^T U W) > 0,$$

which can be expressed as the quadratic matrix inequality

$$W + W^T - W^T U W + (1/n)\underline{1}\underline{1}^T \succ 0.$$

Here, the rank-one matrix $(1/n)\underline{1}\underline{1}^T$ has its only nonzero eigenvalue one associated with the eigenvector $\underline{1}$. Finally, using Schur complements, the above quadratic matrix inequality is equivalent to the LMI (14). Note that we do not need to know L to test the convergence condition (11), although it is needed in calculating the guaranteed convergence rate. □

Remark 2.1. The derivation of the equation (20) holds without assuming the convergence condition (11), so long as we interpret $\lambda_{n-1}(V)$ as the smallest eigenvalue of V excluding the zero eigenvalue associated with the eigenvector $\underline{1}$. It is evident from (20) that $\lambda_{n-1}(V) > 0$ is necessary for $\eta(W) < 1$ and sufficient for the algorithm to converge to the optimal solution.

2.1. Conditions for Symmetric Weights. When the weight matrix W is symmetric, the convergence conditions reduce to

$$W = W^T, \quad W\underline{1} = 0, \tag{21}$$

$$2W + (1/n)\underline{1}\underline{1}^T \succ 0, \tag{22}$$

$$2U^{-1} - W \succ 0. \tag{23}$$

To see this, we first rewrite the LMI (14) for symmetric W ,

$$\begin{bmatrix} 2W + (1/n)\underline{1}\underline{1}^T & W \\ W & U^{-1} \end{bmatrix} \succ 0.$$

Applying Schur complements, this LMI is equivalent to

$$2W + (1/n)\underline{1}\underline{1}^T > 0,$$

$$U^{-1} - W[2W + (1/n)\underline{1}\underline{1}^T]^{-1}W > 0.$$

The first inequality is precisely (22). The second inequality is in fact (23) because, under the conditions (21) and (22), we have

$$W[2W + (1/n)\underline{1}\underline{1}^T]^{-1}W = (1/2)W.$$

The conditions (21)–(23) are generalizations of (10). In particular, we consider the conditions that the matrix $2U^{-1} - W$ is strictly diagonally dominant, which is sufficient for (23) to hold. The strictly diagonal dominance property can be expressed as

$$|2/u_i - W_{ii}| > \sum_{j \in \mathcal{N}_i} |W_{ij}|, \quad i = 1, \dots, n.$$

If all the off-diagonal elements of W are nonpositive and

$$W_{ii} = - \sum_{j \in \mathcal{N}_i} W_{ij},$$

the above condition becomes

$$\sum_{j \in \mathcal{N}_i} |W_{ij}| < 1/u_i, \quad i = 1, \dots, n. \tag{24}$$

When all the numbers u_i are equal or simply take their maximum, as assumed in Ref. 7, the inequality (24) is exactly the third condition in (10).

3. Simple Weight Selections

In this section, we give two simple methods to select the weight matrix so that it satisfies the convergence conditions established in the previous section. Here, we consider only symmetric weight matrices and the methods are based on the sufficient condition (24). For symmetric weight matrices, each edge of the graph is bidirectional and has the same weight in both directions, so each can be considered as an undirected edge with a single weight.

3.1. Constant Weight on Edges. The simplest and most commonly used method is to have constant weight on all the edges of the graph and obtain the self-weights W_{ii} from the equality constraint $W\underline{1} = 0$,

$$W_{ij} = \begin{cases} \alpha, & (i, j) \in \mathcal{E}, \\ -d_i\alpha, & i = j, \\ 0, & \text{otherwise,} \end{cases}$$

where $d_i = |\mathcal{N}_i|$ is the degree of node i . From the condition (24), we can deduce a range of α that guarantees the convergence of the algorithm,

$$-1/(\max_{i \in \mathcal{N}} d_i u_i) < \alpha < 0. \tag{25}$$

Actually, it is also safe to set

$$\alpha = -1/(\max_{i \in \mathcal{N}} d_i u_i) \tag{26}$$

unless all the $d_i u_i$'s are equal. This can be verified by considering the irreducibly diagonally dominant property of the matrix $2U^{-1} - W$; see e.g., Ref. 18, Section 6.2.

We call the constant weight in (26) the max-degree weight. It comes from the fact that, when the functions f_i are appropriately scaled such that $u_i = 1$ for all i , then α is determined solely by the maximum degree of the graph. In fact, the range (25) is usually very conservative. In Section 4.2, we will find the best constant α^* that minimizes the guaranteed convergence rate η in (13). We will see that α^* is often outside the range (25).

3.2. Weights Determined by Local Information. A slightly more sophisticated method is to determine a range for the weight on each edge of the graph. From the condition (24), it is straightforward to obtain the following ranges for the edge weights that guarantee the convergence of the algorithm:

$$-\min\{1/d_i u_i, 1/d_j u_j\} < W_{ij} < 0, \quad (i, j) \in \mathcal{E}.$$

As before, after choosing the edge weights, we can determine the self-weights using

$$W_{ii} = - \sum_{j \in \mathcal{N}_i} W_{ij}.$$

Again, unless all the nodes have the same value of $d_i u_i$, we can set

$$W_{ij} = -\min\{1/d_i u_i, 1/d_j u_j\}, \quad (i, j) \in \mathcal{E}. \tag{27}$$

We call these weights the Metropolis weights, because the main idea of this method relates to the Metropolis algorithms for choosing transition probabilities on a graph to make the associated Markov chain mix rapidly (see Ref. 19; see also e.g. Refs. 16, 20). In Section 5, we will see that the weights selected by

this method often perform better than the max-degree constant weight on all edges.

Note that this method of choosing edge weights relies only on local information in the graph: the degrees and the upper bounds on the second derivatives of the functions at its two incident nodes. This property allows the weights to be selected in a distributed manner, which is ideal for a distributed algorithm running on the graph such as (9).

4. Optimal Scaling of the Center-Free Algorithm

In this section, we pose the question of how to choose the weight matrix to make the algorithm (5) converge as fast as possible. Clearly, the convergence properties depend on the particular objective functions and initial condition. When the lower and upper bounds L and U are the only information available, it is reasonable to choose the weight matrix to minimize the guaranteed convergence rate η established in Theorem 2.1. This is equivalent to maximizing the second smallest eigenvalue of the matrix V in (16), i.e.,

$$\max \lambda_{n-1}[L^{1/2}(W + W^T - W^T U W)L^{1/2}], \tag{28a}$$

$$\text{s.t. } W \in S, \quad \underline{1}^T W = 0, \quad W \underline{1} = 0, \tag{28b}$$

where the optimization variable is W . In this section, we show that this problem is convex, can be converted into a semidefinite program (SDP), and so can be efficiently and globally solved using numerical algorithms such as interior-point methods (see e.g. Ref. 21). We discuss also the special case of finding the best constant weight.

4.1. Weight Design via SDP. We show first that the eigenvalue optimization problem (28) can be converted to an SDP. Let s be a lower bound on the eigenvalues of $L^{1/2}(W + W^T - W^T U W)L^{1/2}$ on the subspace that is orthogonal to $L^{-1/2}\underline{1}$ (which corresponds to the eigenvalue zero). Then, we have

$$L^{1/2}(W + W^T - W^T U W)L^{1/2} \succeq s[I - (1/\underline{1}^T L^{-1}\underline{1})L^{-1/2}\underline{1}\underline{1}^T L^{-1/2}],$$

which is equivalent to (by multiplying $L^{-1/2}$ on the left and right),

$$W + W^T - W^T U W \succeq s[L^{-1} - (1/\underline{1}^T L^{-1}\underline{1})L^{-1}\underline{1}\underline{1}^T L^{-1}]. \tag{29}$$

Using the Schur complements, the above quadratic matrix inequality is equivalent to the LMI

$$\begin{bmatrix} W + W^T - s[L^{-1} - (1/\underline{1}^T L^{-1} \underline{1})L^{-1} \underline{1}\underline{1}^T L^{-1}] & W^T \\ W & U^{-1} \end{bmatrix} \succeq 0. \tag{30}$$

Therefore the eigenvalue optimization problem (28) is equivalent to the SDP

$$\max s, \tag{31a}$$

$$\text{s. t. } W \in S, \quad \underline{1}^T W = 0, \quad W \underline{1} = 0, \tag{31b}$$

$$\begin{bmatrix} W + W^T - s[L^{-1} - (1/\underline{1}^T L^{-1} \underline{1})L^{-1} \underline{1}\underline{1}^T L^{-1}] & W^T \\ W & U^{-1} \end{bmatrix} \succeq 0, \tag{31c}$$

with optimization variables s and W .

Note that the matrices on both sides of the inequality (29) have the common eigenvalue zero associated with the eigenvector $\underline{1}$. As a result, the LMI (30) has empty interior, which can cause trouble for some classes of interior-point methods (see e.g. Refs. 5 and 22–24). But this problem is readily avoided: we simply add the rank-one matrix $(1/n)\underline{1}\underline{1}^T$ to the left hand side of (29) to get the inequality

$$W + W^T - W^T U W + (1/n)\underline{1}\underline{1}^T \succeq s[L^{-1} - (1/\underline{1}^T L^{-1} \underline{1})L^{-1} \underline{1}\underline{1}^T L^{-1}].$$

This leads to the SDP

$$\max s, \tag{32a}$$

$$\text{s. t. } W \in S, \quad \underline{1}^T W = 0, \quad W \underline{1} = 0, \tag{32b}$$

$$\begin{bmatrix} W + W^T + (1/n)\underline{1}\underline{1}^T - s[L^{-1} - (1/\underline{1}^T L^{-1} \underline{1})L^{-1} \underline{1}\underline{1}^T L^{-1}] & W^T \\ W & U^{-1} \end{bmatrix} \succeq 0, \tag{32c}$$

This SDP is equivalent to (31), but here the LMI constraint has nonempty interior, so general interior-point methods can be applied to solve this problem.

Since W is often very sparse, exploiting sparsity can further improve the efficiency of numerical algorithms and allow the solution of very large-scale problems. We discuss how to exploit sparsity in interior-point methods and a simple subgradient method for solving a similar class of SDPs in Ref. 17. The same techniques apply in this case.

4.2. Best Constant Weight. If the weight matrix is symmetric and all the edge weights are equal to a constant α (as described in Section 3.1), we can write the weight matrix as

$$W = -\alpha\mathcal{L}, \tag{33}$$

where \mathcal{L} is the Laplacian matrix of the graph, which is defined as

$$\mathcal{L}_{ij} = \begin{cases} -1, & (i, j) \in \mathcal{E}, \\ d_i, & i = j, \\ 0, & \text{otherwise.} \end{cases}$$

The Laplacian matrix is positive semidefinite. Since the graph is assumed connected, it has a simple eigenvalue $\lambda_n(\mathcal{L}) = 0$ with associated eigenvector $\underline{1}$ (see e.g. Ref. 25–26). From the expression (33), we must have $\alpha < 0$ for W to be positive semidefinite [see condition (22)].

We can substitute (33) into the SDP (32) to solve for the best constant α^* . Here, we consider a simple case where α^* can be found analytically. We assume that the functions f_i are appropriately scaled such that the upper bound on the second derivatives is $U = I$ (the identity matrix) and the lower bound is $L = \beta I$ with $0 < \beta < 1$. In this case,

$$\begin{aligned} V &= \beta(W + W^T - W^T U W) \\ &= \beta(-2\alpha\mathcal{L} - \alpha^2\mathcal{L}^2) \\ &= \beta I - \beta(I + \alpha\mathcal{L})^2. \end{aligned}$$

Now, we can express $\lambda_{n-1}(V)$ in terms of β and the two extreme eigenvalues of \mathcal{L} ,

$$\lambda_{n-1}(V) = \beta - \beta[\max\{1 + \alpha\lambda_{n-1}(\mathcal{L}), -1 - \alpha\lambda_1(\mathcal{L})\}]^2.$$

For $\lambda_{n-1}(V) > 0$, the weight α must lie in the range

$$-2/\lambda_1(\mathcal{L}) < \alpha < 0.$$

The optimal constant weight that maximizes $\lambda_{n-1}(V)$ is

$$\alpha^* = -2/[\lambda_1(\mathcal{L}) + \lambda_{n-1}(\mathcal{L})].$$

It is interesting to note that this result is very similar to the best constant weight found in the fast averaging problem; see Ref. 17, Section 4.1. The only difference is the sign, which is due to different sign conventions used.

5. Numerical Example

In this section, we use simulation to demonstrate the effectiveness of the weight matrices determined by various methods proposed in the previous two

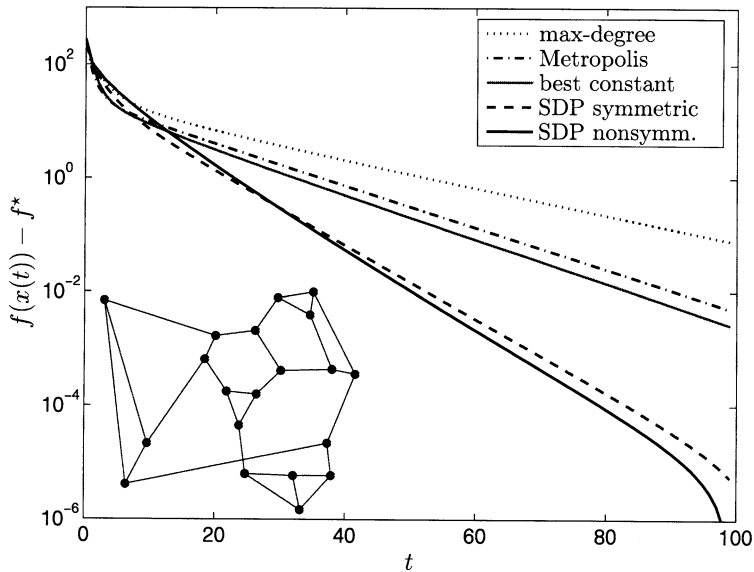


Fig. 1. Simulation on a randomly generated graph (lower left corner): objective function values $f(x(t)) - f^*$ versus iteration number t , with five different weight matrices and the common initial condition $x(0) = 0$.

sections. We consider the graph shown at the lower left corner of Figure 1. This is a randomly generated regular graph with $n = 20$ nodes, where each node has degree three. Each edge is bidirectional. In this example, we consider both symmetric and nonsymmetric weight matrices.

We use the following family of functions:

$$f_i(x_i) = (1/2)a_i(x_i - c_i)^2 + \log[1 + \exp(b_i(x_i - d_i))], \quad i = 1, \dots, n,$$

with the coefficients a_i, b_i, c_i, d_i generated randomly with uniform distributions on $[0, 2], [-2, 2], [-10, 10], [-10, 10]$ respectively. The second derivatives of these functions are

$$f_i''(x_i) = a_i + b_i^2 \exp[b_i(x_i - d_i)]/[1 + \exp(b_i(x_i - d_i))]^2, \quad i = 1, \dots, n,$$

which have the following lower and upper bounds:

$$l_i = a_i, \quad u_i = a_i + (1/4)b_i^2, \quad i = 1, \dots, n.$$

We assume that the sum of the variables is fixed to zero, i.e.,

$$\sum_{i=1}^n x_i = 0.$$

Table 1. Guaranteed convergence rates of different weight matrices.

Method	Max-degree	Metropolis	Best constant	SDP Symmetric	SDP Nonsymmetric
$\eta(W)$	0.9503	0.9237	0.9217	0.8750	0.8729

Using these bounds, we find the weights using the five methods that have been described in Sections 3 and 4. In particular, we find the best constant edge weight by solving the SDP (32), yielding $\alpha^* = -0.2030$, which is outside the range (25) derived from the diagonally dominant condition (24). The smallest value for the diagonally dominant condition is given by the max-degree weight (26), which is $\alpha^{md} = -0.1251$.

Table 1 shows the guaranteed convergence rates $\eta(W)$ obtained with different weight selection methods. These show that the max-degree has the largest value of η , with Metropolis and best constant about the same and smaller. The optimal symmetric and nonsymmetric weights (obtained by solving SDPs) give even faster convergence. The nonsymmetric weight matrix W found by solving the SDP (32) has two positive offdiagonal entries; see Remark 1.1. Of course, it must be remembered that η is only a guaranteed bound on the convergence rate, so small differences in η (such as between the optimal symmetric and optimal nonsymmetric weights) probably have no meaning in terms of the actual convergence.

Figure 1 shows the objective values of the algorithm (5) using the five different weight matrices, all starting with the initial condition $x_i(0) = 0$ for all i . The plot shows that, in this case, the bound η does predict the actual convergence rate; in particular, the convergence with the optimal weights is substantially faster than, for example, the max-degree weights.

6. Extensions

In this section, we extend the main results to problems with more complex constraints and problems with vector variables at each node.

6.1. Problems with General Equality Constraints. We consider the following optimization problem with general linear equality constraints:

$$\min \quad f(x), \tag{34a}$$

$$\text{s. t.} \quad Ax = b, \tag{34b}$$

where $x \in \mathfrak{R}^n$ is the variable and $A \in \mathfrak{R}^{m \times n}$, $b \in \mathfrak{R}^m$ are given parameters. Without loss of generality, we assume that $m < n$ and A is full rank. We assume that the

function $f(x)$ is strictly convex and

$$L \preceq \nabla^2 f(x) \preceq U,$$

for some symmetric positive-definite matrices L and U . In fact, the above partial ordering only needs to hold on the nullspace of A , i.e., $\{x \in \mathfrak{R}^n \mid Ax = 0\}$. Here, we do not assume explicitly additional structure such as those imposed by a graph; we note only that the sparsity pattern of W could depend on the sparsity pattern of $\nabla^2 f(x)$. A point x^* is optimal for (34) if and only if there is a $v^* \in \mathfrak{R}^m$ such that

$$Ax^* = b, \quad \nabla f(x^*) = A^T v^*.$$

We consider solving the problem (34) using the weighted gradient descent algorithm (5). Here, the feasibility and fixed point conditions translate into the following constraints on the weight matrix W :

$$AW = 0, \quad WA^T = 0, \tag{35}$$

which correspond to conditions (7) and (8) respectively. The rows of A form the left and right eigenspace of W associated with the eigenvalue zero, which has multiplicity m . The additional sufficient condition for monotonic convergence becomes [see condition (11)]

$$\lambda_{n-m}(V) = \lambda_{n-m}[L^{1/2}(W + W^T - W^T U W)L^{1/2}] > 0,$$

which can be expressed as the strict LMI condition [see condition (14)]

$$\begin{bmatrix} W + W^T + A^T(AA^T)^{-1}A & W \\ W^T & U^{-1} \end{bmatrix} \succ 0.$$

As before, the convergence condition does not rely on L . In this case, the guaranteed convergence rate is given by

$$\eta(W) = 1 - \lambda_{n-m}(V).$$

Similarly, for optimization-based weight design, we need only to carry out the following replacements in the SDPs (31) and (32):

$$\begin{aligned} (1/n)\underline{1}\underline{1}^T &\longleftrightarrow A^T(AA^T)^{-1}A, \\ (1/\underline{1}^T L^{-1}\underline{1})L^{-1}\underline{1}\underline{1}^T L^{-1} &\longleftrightarrow L^{-1}A^T(AL^{-1}A^T)^{-1}AL^{-1}. \end{aligned}$$

In this case, however, it is difficult to come up with simple methods to select the weight matrix, like those given in Section 3. While the diagonally dominant condition (24) is still useful, it is difficult to directly construct weight matrices satisfying the eigenspace conditions in (35).

6.2. Problems with Vector Variables at Each Node. We consider a modified version of problem (1)

$$\min \sum_{i=1}^n f_i(x_i), \tag{36a}$$

$$\text{s. t. } \sum_{i=1}^n x_i = b, \tag{36b}$$

where each x_i is an m -dimensional vector, i.e.,

$$x_i = (x_i^1, \dots, x_i^m), \quad \text{for } i = 1, \dots, n.$$

Here we do impose the local information structure given by the graph. In this case, we assume the following condition on the Hessians of the functions f_i

$$L_i \preceq \nabla^2 f(x_i) \preceq U_i \quad i = 1, \dots, n,$$

where L_i and U_i are positive-definite matrices.

The center-free algorithm now takes the form

$$x_i(t + 1) = x_i(t) + \sum_{j \in \mathcal{N}_i} W_{ij} [\nabla f_i(x_i(t)) - \nabla f_j(x_j(t))], \tag{37}$$

where $W_{ij} \in \mathfrak{R}^{m \times m}$ is the weight matrix between two adjacent nodes i and j . The n by n blocks of W (each with size m by m) are either zero (if the corresponding two nodes are not connected) or a full m by m matrix (for two adjacent nodes). A simplified version of (37) was suggested in Ref. 7, where the weight matrix W_{ij} between two adjacent node is diagonal.

The convergence condition and optimization-based weight design can be obtained by applying the results in Section 6.1. Here, $x = (x_1^T, \dots, x_n^T)$ is an augmented vector of length mn and the equality constraint can be written as $Ax = b$, where $A = [I_m, \dots, I_m]$ and I_m is the m -dimensional identity matrix. The weight matrix W should satisfy the eigenspace conditions (35). The lower and upper bounding matrices are block diagonal,

$$L = \text{diag}(L_1, \dots, L_m), \quad U = \text{diag}(U_1, \dots, U_n).$$

7. Conclusions

We have considered a class of distributed gradient algorithms for optimally allocate a fixed amount of resource over a network, to minimize the sum of the convex cost functions at each node. In these algorithms, each node updates the amount of its local resource in proportion to the differences between the marginal costs of itself and its neighbors. We focused on how to choose the proportional

weights on the edges to make the algorithm converge and on how to optimize the weights (scaling of the gradient method) to get fast convergence. We should make several comments about the algorithm and the methods for selecting the weights.

First, the algorithm is designed for distributed implementation, with a simple protocol between each pair of connected nodes. In general, it is not competitive with offline centralized optimization algorithms, such as BFGS or Newton's methods. These methods generally give faster (and sometimes much faster) convergence than the center-free algorithm; but on the other hand, they cannot be implemented in a simple, fully distributed way.

The second point concerns the cost of computing the optimal weights by solving an SDP. While these weights evidently yield faster convergence (compared to, say, the max-degree or Metropolis weights), solving the SDP often involves substantially more computational effort than solving the original resource allocation problem. So if the resource allocation problem is to be solved just once or a few times (say), this extra cost is certainly not justified. On the other hand, if we are to solve the resource allocation problem multiple times, on a network with the same topology, with (say) different initial conditions or different cost functions (with the same lower and upper bounds on their second derivatives), then the extra effort of solving an SDP to find the optimal weights can be justified.

In contrast, the Metropolis methods (which are based only on local information) are well suited for distributed weight selection and can potentially be used with time-varying network topologies. One can imagine that, at a slower time scale, nodes can join or leave the network, carrying in or out their current share of total resource. The algorithms presented in this paper, with the Metropolis weights determined in real-time, can run at a faster time scale to track the optimal solution of each network configuration.

References

1. ARROW, K. J., and HURWICZ, L., *Decentralization and Computation in Resource Allocation*, Essays in Economics and Econometrics, Edited by R. W. Pfouts, University of North Carolina Press, Chapel Hill, North Carolina, pp. 34–104, 1960.
2. HEAL, G. M., *Planning without Prices*, Review of Economic Studies, Vol. 36, pp. 347–362, 1969.
3. KUROSE, J. F., and SIMHA, R., *A Microeconomic Approach to Optimal Resource Allocation in Distributed Computer Systems*, IEEE Transactions on Computers, Vol. 38, pp. 705–717, 1989.
4. BERTSEKAS, D. P., *Nonlinear Programming*, 2nd Edition, Athena Scientific, Belmont, Massachusetts, 1999.
5. BOYD, S., and VANDENBERGHE, L., *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004.
6. HURWICZ, L., *The Design of Mechanisms for Resource Allocation*, American Economic Review, Vol. 63, pp. 1–30, 1973.

7. Ho, Y. C., SERVI, L. D., and SURI, R., *A Class of Center-Free Resource Allocation Algorithms*, Large Scale Systems, Vol. 1, pp. 51–62, 1980.
8. SERVI, L. D., *Electrical Networks and Resource Allocation Algorithms*, IEEE Transactions on Systems, Man, and Cybernetics, Vol. 10, pp. 841–848, 1980.
9. ROCKAFELLAR, R. T., *Network Flows and Monotropic Optimization*, John Wiley and Sons, New York, NY, 1984.
10. BERTSEKAS, D. P., and TSITSIKLIS, J. N., *Parallel and Distributed Computation*, Prentice-Hall, Englewood, Cliffs, New Jersey, 1989.
11. BERTSEKAS, D. P., *Network Optimization: Continuous and Discrete Models*, Athena Scientific, Belmont, Massachusetts, 1998.
12. TSITSIKLIS, J. N., and BERTSEKAS, D. P., *Distributed Asynchronous Optimal Routing in Data Networks*, IEEE Transactions on Automatic Control, Vol. 31, pp. 325–332, 1986.
13. EL BAZ, D., *Asynchronous Gradient Algorithms for a Class of Convex Separable Network Flow Problems*, Computational Optimization and Applications, Vol. 5, pp. 187–205, 1996.
14. LUO, Z. Q., and TSENG, P., *On the Rate of Convergence of a Distributed Asynchronous Routing Algorithm*, IEEE Transactions on Automatic Control, Vol. 39, pp. 1123–1129, 1994.
15. TSITSIKLIS, J. N., BERTSEKAS, D. P., and ATHANS, M., *Distributed Asynchronous Deterministic and Stochastic Gradient Optimization Algorithms*, IEEE Transactions on Automatic Control, Vol. 31, pp. 803–812, 1986.
16. BOYD, S., DIACONIS, P., and XIAO, L., *Fastest Mixing Markov Chain on a Graph*, SIAM Review, Vol. 46, pp. 667–689, 2004.
17. XIAO, L., and BOYD, S., *Fast Linear Iterations for Distributed Averaging*, Systems and Control Letters, Vol. 53, pp. 66–78, 2004.
18. HORN, R. A., and JOHNSON, C. A., *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1985.
19. METROPOLIS, N., ROSENBLUTH, A., ROSENBLUTH, M., TELLER, A., and TELLER, E., *Equations of State Calculations by Fast Computing Machines*, Journal of Chemical Physics, Vol. 21, pp. 1087–1092, 1953.
20. DIACONIS, P., and SALOFF-COSTE, L., *What Do We Know About the Metropolis Algorithms*, Journal of Computer and System Sciences, Vol. 57, pp. 20–36, 1998.
21. VANDENBERGHE, L., and BOYD, S., *Semidefinite Programming*, SIAM Review, Vol. 38, pp. 49–95, 1996.
22. NESTEROV, Y., and NEMIROVSKII, A., *Interior-Point Polynomial Algorithms in Convex Programming*, SIAM Studies in Applied Mathematics, SIAM, Philadelphia, Pennsylvania, 1994.
23. Ye, Y., *Interior-Point Algorithms: Theory and Analysis*, John Wiley and Sons, New York, NY, 1997.
24. WOLKOWICZ, H., SAIGAL, R., and VANDENBERGHE, L., Editors, *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*, Kluwer Academic Publishers, Boston, Massachusetts, 2000.
25. MERRIS, R., *Laplacian Matrices of Graphs: A Survey*, Linear Algebra and Its Applications, Vol. 197, pp. 143–176, 1994.
26. GODSIL, C., and ROYLE, G., *Algebraic Graph Theory*, Springer, New York, NY, 2001.