Cloud-Based Optimization: A Quasi-Decentralized Approach to Multi-Agent Coordination

M.T. Hale and M. Egerstedt[†]

Abstract— **New architectures and algorithms are needed to reflect the mixture of local and global information that is available as multi-agent systems connect over the cloud. We present a novel architecture for multi-agent coordination where the cloud is assumed to be able to gather information from all agents, perform centralized computations, and disseminate the results in an intermittent manner. This architecture is used to solve a multi-agent optimization problem in which each agent has a local objective function unknown to the other agents and in which the agents are collectively subject to global inequality constraints. Leveraging the cloud, a dual problem is formulated and solved by finding a saddle point of the associated Lagrangian.**

I. INTRODUCTION

Distributed optimization and algorithms have received significant attention during the last decade, e.g., [1], [20], [10], [8], [15], [22], [6], due to the emergence of a number of application domains in which individual decision makers have to collectively arrive at a decision in a distributed manner. Examples of these applications include communication networks [12], [5], sensor networks [13], [23], [2], multi-robot systems [26], [21], and smart power grids [3].

Distributed algorithms are needed mainly because the scale of large distributed systems is such that no central, global decision maker can collect all relevant information, perform all required computations, and then disseminate the results back to individual nodes in the network in a timely fashion. However, one can envision a scenario in which such globally obtained information can be used in conjunction with local computations performed across the network. This could, for example, be the case when a cloud computer is available to collect information, as was envisioned in [9]. The question then becomes that of designing the appropriate architecture and algorithms that can leverage this mix of prompt decentralized computations with intermittent centralized computations.

One approach to multi-agent optimization that will prove useful towards achieving this hybrid architecture is based on primal-dual methods to find saddle points of a problem's Lagrangian [19], [7]. In fact, the study of saddle point dynamics in optimization can be traced back to earlier results from Uzawa in [24], which will provide the starting point for the work in this paper. The primary difference between this paper and the established literature is the cloud-based architecture used to solve the problem; indeed the architecture is this paper's main contribution. The architecture we introduce uses a cloud computer in order to receive information from each agent, perform global computations, and transmit this information to other agents. We will see that this division of labor results in globally asymptotic convergence to an ϵ-ball about a Lagrangian's saddle point.

The goal of this paper is to serve as a first attempt at understanding how centralized, cloud-based information might be injected in an intermittent but useful manner into a network of agents where such information would otherwise be absent. In order to highlight how the cloud might prove useful to such a system, we choose to consider an extreme case where no inter-agent communication occurs at all, in contrast to existing distributed multi-agent optimization techniques, e.g., [14], [16], [17], [18]. Under this architecture the cloud handles all communications, and computations are divided between the cloud and the agents in the network.

The rest of the paper is organized as follows: Section II gives a detailed problem statement and describes the cloud architecture, and then Section III provides the convergence analysis for the given problem. Next, Section IV provides numerical results to demonstrate the viability of this approach, and finally Section V concludes the paper.

II. PROBLEM STATEMENT AND ARCHITECTURE

A. Architecture Motivation

We now explain the interplay between the cloud architecture and the problem under consideration here. A detailed explanation is given below, with a summary and example following at the end of this section. Consider a collection of N agents indexed by $i \in A$, $A = \{1, \dots, N\}$, where each agent is associated with a scalar state $x_i \in \mathbb{R}$ and where there is no communication at all between the agents. Let the task agent i is trying to solve be encoded in a strictly convex objective functions in C^2, $f_i : \mathbb{R} \to \mathbb{R}$. Each agent is assumed to have no knowledge of other agents' objective function and each agent's only goal is to minimize its own objective function.

To that end, agent i is assumed to have immediate access to its own state, which seemingly makes this problem very simple. However, what prevents agent i from simply computing $\frac{df_i}{dx_i}$ and setting this equal to zero – a completely decentralized operation as f_i only depends on x_i – is that the agents need to coordinate their actions through a globally defined constraint, that can, for example, represent finite resources that must be shared across the team. In this paper it is assumed that agent i cannot measure the state of any other agents and, as mentioned above, that there is

no communication between agents. Instead, this information must be obtained in some other manner, which is where the cloud will enter into the picture.

The team-wide coordination is encoded through the global constraint

$$g(x) = \begin{pmatrix} g_1(x) \\ g_2(x) \\ \vdots \\ g_m(x) \end{pmatrix} \leq 0, \tag{1}$$

where $x = (x_1, \ldots, x_N)^T$ is a state vector containing the states of all agents in the network. It is further assumed that each $g_j(x) \in C^2$ is convex. The cloud architecture discussed here applies to any problem in which the user has selected functions f_i and g_j that meet the above criteria and the forthcoming analysis fully characterizes all such problem formulations.

Let

$$F(x) = \sum_{i=1}^{N} f_i(x_i). \tag{2}$$

Then F is strictly convex and the problem under consideration becomes that of minimizing F subject to g. The Kuhn-Tucker Theorem on concave programming (e.g., [25]) states that the optimum of this constrained problem is a saddle point of the Lagrangian

$$L(x, \mu) = F(x) + \mu^T g(x), \tag{3}$$

where the Kuhn-Tucker (KT) multipliers μ_j satisfy $\mu_j \geq 0$ for all $j \in \{1, \ldots, m\}$. We assume that the minimizer of L with respect to x, denoted $\hat{x}$, is a regular point of g so that there is a unique saddle point, $(\hat{x}, \hat{\mu})$, of L [4]. Using that L is convex in x and concave in μ, the saddle point $(\hat{x}, \hat{\mu})$ can be shown to satisfy the inequalities

$$L(\hat{x}, \mu) \leq L(\hat{x}, \hat{\mu}) \leq L(x, \hat{\mu}) \tag{4}$$

for all admissible x and μ.

Using Uzawa's algorithm [24], the problem of finding $(\hat{x}, \hat{\mu})$ can be solved from the initial point $(x(0), \mu(0))$ using the difference equations

$$x(k) = x(k-1) - \rho \frac{\partial L}{\partial x}(x(k-1), \mu(k-1)) \tag{5}$$

$$\mu(k) = \max \left\{ 0, \ \mu(k-1) + \rho \frac{\partial L}{\partial \mu}(x(k-1), \mu(k-1)) \right\} \tag{6}$$

where $\rho > 0$ is a constant, and where the maximum defining μ is taken component-wise so that each component of μ is projected onto the non-negative orthant of $\mathbb{R}^m$, denoted by $\mathbb{R}^m_+$. In the context of Uzawa's algorithm, the i^{th} element of the state vector x is updated according to

$$x_i(k) = x_i(k-1) - \rho \frac{\partial L}{\partial x_i}(x(k-1), \mu(k-1)). \tag{7}$$

Under the envisioned organization of the agents and the lack of inter-agent communication, Uzawa's algorithm cannot be directly applied. To see this, observe that if agent i is to compute its own state update using Equation (7),

a fundamental problem is encountered: computing $\frac{\partial L}{\partial x_i}$ will require knowledge of states of (possibly all) other agents and agent i cannot directly access this information. Furthermore, determining μ at each timestep using (6) will also require the full state vector x, which no single agent has direct access to.

To account for the need of each agent for global information in applying Equation (7) and to compute μ using aggregated global information, the cloud computer is used. The cloud computer is taken to be capable of large batch computations and receives periodic transmissions from each agent containing each agent's own state. The cloud computer uses the agents' states to compute the next value of μ using Equation (6) and then transmits the states it received and the newly computed μ vector to each agent. Each agent then uses the information from the cloud to update its own state in the vein of (7).

B. Formal Architecture Description

We first describe the actions taken to initialize the system and then explain its operation. Let the agents each be programmed with their objective functions onboard and let them either be programmed with an initial state or else be able to sense it (e.g., if it corresponds to some physical quantity). The agents are assumed to be identifiable according to their indices in A so that the cloud knows the source of each transmission it receives. Each agent stores and manipulates a state vector onboard and we denote the state vector stored onboard agent i by x^i; agent i's copy of its own state is denoted x_i^i and when we are referring to a specific point in time, say timestep k, we denote agent i's copy of its own state at this time by $x_i^i(k)$. The vector of KT multipliers stored onboard agent i at time k is denoted $\mu^i(k)$, though we emphasize that agent i does not compute any KT vectors but instead relies on the cloud for these computations.

Before the optimization process begins, let agent i send its initial state, $x_i^i(0)$, to the cloud and let the cloud store these states in the vector $x^c(0) \in \mathbb{R}^N$, with the superscript 'c' denoting "cloud" and the timestep 0 reflecting that this is the initial state. In this notation, the cloud's copy of agent i's state at time k is denoted $x_i^c(k)$. Similarly, we denote the KT vector stored in the cloud at time k by $\mu^c(k)$. Let the cloud be programmed by the user with the constraint functions, $g(x^c)$. Upon receiving the each agent's state, the cloud symbolically computes $\frac{\partial g}{\partial x_i^c}$ and sends this function to agent i along with some initial KT multiplier vector, $\mu(0)$, a stepsize $\rho > 0$, and the vector $\boldsymbol{y}^i \in \mathbb{R}^{N-1}$ defined as

$$\boldsymbol{y}^i = \begin{pmatrix} x_1^c \\ \vdots \\ x_{i-1}^c \\ x_{i+1}^c \\ \vdots \\ x_N^c \end{pmatrix}. \tag{8}$$

This vector contains states stored by the cloud in the vector x^c and contains information originally from time 0 (though

in Equation (8) explicit timesteps are intentionally omitted). The subscripts in (8) denote that agent i does not receive its own old state value from the cloud, which is logical since agent i always knows its own most recent state. In $\boldsymbol{y}^i$, then, the cloud sends to agent i the most recent state information it has about each *other* agent. In this notation, agent j's state in $\boldsymbol{y}^i$ is denoted $\boldsymbol{y}^i_j$. In the forthcoming analysis, $\boldsymbol{y}^i$ always refers to the most recent state information that agent i has received from the cloud and it will not be written as an explicit function of any time step. Similarly, the notation $\boldsymbol{\mu}^i$ refers to the most recent KT vector sent to agent i and will be written without an explicit timestep. We use the notation $\boldsymbol{z}^i$ to denote the most recent transmission to agent i containing both $\boldsymbol{y}^i$ and $\boldsymbol{\mu}^i$.

After receiving $\boldsymbol{z}^i$ for this first time, all N agents and the cloud have the same information onboard, and each agent begins the optimization process. At timestep 0, each agent takes one gradient step to update its own state according to Equation (7). Simultaneously, and also at timestep 0, the cloud takes one gradient step to update the KT multipliers in the cloud according to Equation (6). Then at timestep 1, agent i sends its state, $x^i_i(1)$, to the cloud. These transmissions are received at timestep 2. In timestep 2, the cloud sends $\boldsymbol{y}^i$ and $\mu^c(1)$ to agent i. These vectors are received in timestep 3 at which point the cloud updates μ^c as before and each agent takes a step to update its own state as before, thus repeating this cycle of communication and computation. It is important to note that communications cycles do not overlap and that the agents do not send their states to the cloud at every timestep, but instead do so every 3^{rd} timestep. In addition, we emphasize that each agent's objective function is assumed to be private throughout this process.

Due to the communications structure of the system, it is often the case that $x^i_i(k) \neq x^j_i(k)$, namely that agents i and j will have different values for agent i's state beacuse agent j must wait to received agent i's state from the cloud. Due these differences, Equation (5) is modified to reflect that each agent stores and manipulates a local copy of the problem. The global system therefore contains N copies of the system in Equations (5) and (6) and the state vector of agent i at time k, $x^i(k)$, is assumed to be different from that of agent j at time k, $x^j(k)$, when $i \neq j$.

Using the fact that agent i will only update its state in timesteps just after it receives an update from the cloud, Equation (5) is modified so that onboard agent i it is

$$ x^i(k) = \begin{cases} \bar{\boldsymbol{y}}^i - \rho\nabla^i L(\bar{\boldsymbol{y}}^i, \boldsymbol{\mu}^i) & \boldsymbol{z}^i \text{ received at time } k{-}1 & \text{(9a)} \\ x^i(k-1) & \text{else,} & \text{(9b)} \end{cases} $$

where we define

$$ \nabla^i L(\bar{\boldsymbol{y}}^i, \boldsymbol{\mu}^i) = \begin{pmatrix} 0 \\ \vdots \\ \frac{df_i}{dx_i}(x^i_i(k-1)) + (\boldsymbol{\mu}^i)^T \frac{\partial g}{\partial x_i}(\bar{\boldsymbol{y}}^i) \\ \vdots \\ 0 \end{pmatrix}. \quad (10) $$

Here, $\bar{\boldsymbol{y}}^i$ is defined as a vector onboard agent i which contains $\boldsymbol{y}^i$ and the most recent state of agent i inserted in the appropriate place. In essence, $\bar{\boldsymbol{y}}^i$ is the most up-to-date information about all of the agents that agent i has access to and contains the correct value of each other agent's state when it is received. Note that $\nabla^i L$ is simply $\frac{\partial L}{\partial x}$ with all entries except the i^{th} set to 0. This is because agent i does not itself compute any updates for the other agents' states which it stores onboard, but instead waits for the cloud to provide such updates.

Under the architecture of this problem, only the cloud computes values of μ and there is therefore only a single update equation needed for μ. Bearing in mind that updates to μ are only made in timesteps immediately after those in which the cloud receives each agent's state, Equation (6) is modified to take the form

$$ \mu^c(k) = \begin{cases} \left[\mu^c(k{-}1) + \rho\frac{\partial L}{\partial \mu}(x^c(k{-}1), \mu^c(k{-}1)) \right]_+ & \begin{subarray}{l}\text{update}\\\text{at } k{-}1\end{subarray} & \text{(11a)} \\ \mu^c(k-1) & \text{else,} & \text{(11b)} \end{cases} $$

where $[\cdot]_+$ denotes the projection onto $\mathbb{R}^m_+$ and the update referred to in Equation (11a) is an update of each state's value sent to the cloud.

We note that Equation (11) is not indexed on a per-agent basis since only the cloud computes values of μ. However, we will continue to use the notation $\mu^i(k)$ to denote the μ vector stored on agent i at time k (which may be different from the μ vector stored in the cloud at time k). It is important to note that the argument of $\mu^i(k)$ is intended to reflect the time at which agent i has μ^i onboard and *does not* imply that μ^i was computed at time k or that agent i computed it. In this notation $\boldsymbol{\mu}^i$ represents the μ vector most recently sent from the cloud to agent i, while μ^i represents the μ vector stored on agent i.

With this model in mind, instead of considering the system defined in Equations (5) and (6), we consider N copies of the system defined by (9) and (11). Using the notation that $\mu^i(k)$ represents the vector μ as stored on agent i at time k, we can write the full update equations onboard agent i as

$$ x^i(k) = \begin{cases} \bar{\boldsymbol{y}}^i - \rho\nabla^i L(\bar{\boldsymbol{y}}^i, \boldsymbol{\mu}^i) & \boldsymbol{z}^i \text{ received at time } k{-}1 & \text{(12a)} \\ x^i(k-1) & \text{else,} & \text{(12b)} \end{cases} $$

$$ \mu^i(k) = \begin{cases} \boldsymbol{\mu}^i & \boldsymbol{z}^i \text{ was received at time } k-1 & \text{(13a)} \\ \mu^i(k-1) & \text{else,} & \text{(13b)} \end{cases} $$

where all changes in μ will result from the cloud using Equation (11).

To illustrate the communications cycle described above, Table 1 contains a sample schedule for a single cycle. Each timestep is listed on the left and the corresponding actions taken at that timestep are listed on the right.

III. CONVERGENCE ANALYSIS

A. Ultimate Boundedness of Solutions

In this section we will examine the evolution of the sequence

$$ z^i(k) = \begin{pmatrix} x^i(k) \\ \mu^i(k) \end{pmatrix} \quad (14) $$

Timestep	Actions
k	Each agent receives a transmission from the cloud and then takes 1 step in its own copy of the problem to update (only) its own state using Equation (12a). At the same time, the cloud computes updated μ values using (11a).
$k+1$	Each agent sends it state to the cloud. Equation (12b) is used by the agents and Equation (11b) is used by the cloud so that no further computations are carried out during this timestep.
$k+2$	The cloud receives the agents' transmissions from time $k+1$ and stores them in x^c. It then sends $\boldsymbol{y}^i$ to agent i, along with $\mu^c(k+1)$, the most recently computed vector of KT multipliers (computed in timestep $k+1$). As in timestep $k+1$, Equations (12b) and (11b) are used so that no further computations take place across the network.
$k+3$	This step is identical to step k. Agent i receives $\boldsymbol{z}^i$ and then takes 1 step in its own copy of the problem to update (only) its own state using Equation (12a). At the same time, the cloud computes updated μ values using (11a).

Table 1: A sample schedule for one communications cycle used by the agents and cloud to exchange information.

for an arbitrary $i \in A$ in order to show that each agent's local copy of the problem converges to an ϵ-ball about the point $\hat{z} = (\hat{x}, \hat{\mu})$.

Specifically, the goal here is two-fold: to prove that the state of each agent's optimization problem enters a ball of radius ϵ about the saddle point $\hat{z}$ in finite time and to show that it does not leave that ball thereafter. Our approach will differ from that of [24] because we use the notion of ultimate boundedness, published after Uzawa, to simplify certain components of proof. We restate the definition of ultimate boundedness here for general discrete-time systems of the form

$$w(k) = f(w(k-1)). \tag{15}$$

Lemma 1: Let $G \subseteq \mathbb{R}^N$ and let $V(w)$ be a Lyapunov candidate function for the system in Equation (15) defined on G such that for all $w \in G$

$$\Delta V(w) = V(f(w)) - V(w) \le a \tag{16}$$

for some $a \ge 0$. Let $\bar{G}$ denote the closure of G and let S be the set

$$S = \{w \in \bar{G} : \Delta V(w) \ge 0\}. \tag{17}$$

Let $b = \sup\{V(w) : w \in S\}$ and define the set A by

$$A = \{z \in \bar{G} : V(w) \le a + b\}. \tag{18}$$

Then any solution $\{w(k)\}$ to Equation (15) which remains in G for all time and enters A at some point is contained in A for all time thereafter.
Proof: See [11], Theorem 5. ∎

We also state a corollary to this result which will be used below.

Corollary 1: Let the conditions of Lemma 1 hold. Suppose that

$$\sup\{-\Delta V(w) : w \in \bar{G} \backslash A\} > 0 \tag{19}$$

and that G is of the form

$$G = \{w : V(w) \le r\}. \tag{20}$$

Then every solution $\{w(k)\}$ of Equation (15) which starts in G remains in G for all time and enters A in a finite number of steps.
Proof: See [11], Corollaries 3 and 4. ∎

Below, we will combine Lemma 1 and Corollary 1 to show that each agent's state trajectory enters a ball of radius ϵ about $\hat{z}$, denoted $B_\epsilon(\hat{z})$, and stays within that ball. Before proving the main convergence result, we prove the following lemma which establishes a positive upper bound on the stepsizes that can be used. We will proceed in the vein of [24] and consider the Lyapunov function

$$V(x, \mu) = \|x - \hat{x}\|^2 + \|\mu - \hat{\mu}\|^2. \tag{21}$$

Lemma 2: Let L denote the Lagrangian in Equation (3) and set $R = \max\{\epsilon, V(x(0), \mu(0))\}$. Define the constants γ_1 and γ_2 by

$$\gamma_1 = \min_{(x,\mu)} \left\{ \sqrt{\frac{\epsilon/2}{\left\|\frac{\partial L}{\partial x}\right\|^2 + \left\|\frac{\partial L}{\partial \mu}\right\|^2}} \,\middle|\, V(x,\mu) \le \frac{\epsilon}{2} \right\} \tag{22}$$

and

$$\gamma_2 = \min_{(x,\mu)} \left\{ \frac{-(\hat{x}-x)^T \frac{\partial L}{\partial x} + (\hat{\mu}-\mu)^T \frac{\partial L}{\partial \mu}}{\left\|\frac{\partial L}{\partial x}\right\|^2 + \left\|\frac{\partial L}{\partial \mu}\right\|^2} \,\middle|\, \frac{\epsilon}{2} \le V(x,\mu) \le R \right\} \tag{23}$$

where $\frac{\partial L}{\partial x}$ and $\frac{\partial L}{\partial \mu}$ above are (implicitly) functions of any x and μ satisfying the conditions on V pertaining to each set. Then setting

$$\rho_{max} = \min\{\gamma_1, \gamma_2\} \tag{24}$$

provides $\rho_{max} > 0$.
Proof: It suffices to show that γ_1 and γ_2 are both positive. The denominator of γ_1 is always positive and tends to zero as $(x, \mu) \to (\hat{x}, \hat{\mu})$ so that the minimum defining γ_1 does not go to zero at $(\hat{x}, \hat{\mu})$. The numerator of γ_1 is positive by insepction and γ_1 itself is therefore the square root of a positive real number.

For γ_2, we note that L is convex in x and concave in μ. The term $-(\hat{x} - x)^T \frac{\partial L}{\partial x}$ is the negation of the directional derivative of $L(\cdot, \mu)$ with respect to x pointing toward its minimizer, and the term $(\hat{\mu} - \mu)^T \frac{\partial L}{\partial \mu}$ is the directional derivative of $L(x, \cdot)$ with respect to μ pointing toward its maximizer. Both terms are therefore non-negative and because the definition of γ_2 precludes $(x, \mu) = (\hat{x}, \hat{\mu})$, the sum of these two terms is strictly positive. The denominator in the definition of γ_2 is positive as well so that γ_2 itself is. ∎

The above Lemmata and Corollaries are stated in terms of the Lagrangian defined in Equation (3). While each agent in the network stores and manipulates its own state vector and thus has its own (unique) Lagrangian, after each gradient

descent step is taken and all states and KT multipliers are shared across the network, every agent ends up with the same information before taking its next step. In addition, every agent and the cloud use the same stepsize, ρ. Then despite the distribution of information and computation throughout the network, the effective outcome of each cycle of communication and computation as described in Section II-B is one step in each of Equations (5) and (6) performed simultaneously.

Therefore, the analysis of the algorithm can be carried out for Equations (5) and (6), and for simplicity we choose to use Equations (5) and (6) in the forthcoming analysis with the understanding that it applies equally well to all agents. Due to the centrality of the convergence of Uzawa's algorithm to this paper and in order to make use of results published after the algorithm's original publication, we now present the main result on the ultimate boundedness of solutions to the problem at hand.

Theorem 1: Let every agent use a strictly convex objective function $f_i : \mathbb{R} \to \mathbb{R}$, $f_i \in C^2$ and let the global constraints, $g : \mathbb{R}^N \to \mathbb{R}^m$, be convex with $g_j \in C^2$ for each j. Then for any stepsize ρ such that $0 < \rho \le \rho_{max}$ used by all agents and the cloud, each agent's local copy of the problem enters an ϵ-ball about $\hat{z}$ in a finite number of steps and stays within that ball for all time thereafter.

Proof: In addition to using V as defined in Equation (21), we equivalently use that $z(k) = (x(k), \mu(k))^T$ to write

$$V(z(k)) = \|z(k) - \hat{z}\|^2. \tag{25}$$

When it is convenient, we will also use the more concise notation $V(k) = V(z(k))$.

We further define

$$\Delta V(k) = \|z(k+1) - \hat{z}\|^2 - \|z(k) - \hat{z}\|^2 \tag{26}$$
$$= \left(\|\mu(k+1) - \hat{\mu}\|^2 - \|\mu(k) - \hat{\mu}\|^2 \right) \tag{27}$$
$$+ \left(\|x(k+1) - \hat{x}\|^2 - \|x(k) - \hat{x}\|^2 \right). \tag{28}$$

As in Lemma 2, we define

$$R = \max \left\{ \epsilon, V\big(x(0), \mu(0)\big) \right\}. \tag{29}$$

Let gradient steps be taken at timestep k so that Equation (12a) is used by all agents to update their states and Equation (11a) is used by the cloud to update μ. From Equation (5) we see that

$$\|x(k+1)\|^2 = \|x(k)\|^2$$
$$- 2\rho x(k)^T \frac{\partial L}{\partial x}(x(k), \mu(k)) + \rho^2 \left\| \frac{\partial L}{\partial x}(x(k), \mu(k)) \right\|^2, \tag{30}$$

and multiplying both sides of Equation (5) by $-2\hat{x}^T$ gives

$$- 2\hat{x}^T x(k+1) = -2\hat{x}^T x(k) - 2\hat{x}^T \frac{\partial L}{\partial x}(x(k), \mu(k)). \tag{31}$$

Using Equations (30) and (31) we see that

$$\|x(k+1) - \hat{x}\|^2 = \|x(k) - \hat{x}\|^2$$
$$+ 2\rho(\hat{x} - x(k))^T \frac{\partial L}{\partial x}(x(k), \mu(k)) + \rho^2 \left\| \frac{\partial L}{\partial x}(x(k), \mu(k)) \right\|^2. \tag{32}$$

Carrying out the same steps for μ using Equation (6) gives

$$\|\mu(k+1) - \hat{\mu}\|^2 = \|\mu(k) - \hat{\mu}\|^2$$
$$- 2\rho(\hat{\mu} - \mu(k))^T \frac{\partial L}{\partial \mu}(x(k), \mu(k)) + \rho^2 \left\| \frac{\partial L}{\partial \mu}(x(k), \mu(k)) \right\|^2. \tag{33}$$

Summing Equations (32) and (33) gives

$$\|x(k+1) - \hat{x}\|^2 + \|\mu(k+1) - \hat{\mu}\|^2$$
$$= \|x(k) - \hat{x}\|^2 + \|\mu(k) - \hat{\mu}\|^2$$
$$- \rho \left[2 \left(-(\hat{x} - x(k))^T \frac{\partial L}{\partial x} + (\hat{\mu} - \mu(k))^T \frac{\partial L}{\partial \mu} \right) \right.$$
$$\left. - \rho \left(\left\| \frac{\partial L}{\partial x} \right\|^2 + \left\| \frac{\partial L}{\partial \mu} \right\|^2 \right) \right], \tag{34}$$

and hence

$$\Delta V(k) =$$
$$- \rho \left[2 \left(-(\hat{x} - x(k))^T \frac{\partial L}{\partial x} + (\hat{\mu} - \mu(k))^T \frac{\partial L}{\partial \mu} \right) \right.$$
$$\left. - \rho \left(\left\| \frac{\partial L}{\partial x} \right\|^2 + \left\| \frac{\partial L}{\partial \mu} \right\|^2 \right) \right]. \tag{35}$$

Suppose now that $\frac{\epsilon}{2} \le V(k) \le R$. Then using the fact that $\rho \le \gamma_2$ we see that

$$\Delta V(k) \le \rho \left[(\hat{x} - x(k))^T \frac{\partial L}{\partial x} - (\hat{\mu} - \mu(k))^T \frac{\partial L}{\partial \mu} \right] < 0 \tag{36}$$

where the right-hand side is negative because ρ is positive and the term inside brackets is negative. The negativity of the term in the brackets is established by observing that it is the numerator of the term defining γ_2 multiplied by -1 and, because the numerator of the fraction defining γ_2 was shown to be positive, we see here that this term is negative. In fact, the term in brackets is bounded above by some negative constant, i.e., there exists $\delta > 0$ such that

$$(\hat{x} - x(k))^T \frac{\partial L}{\partial x} - (\hat{\mu} - \mu(k))^T \frac{\partial L}{\partial \mu} \le -\delta < 0 \tag{37}$$

which is seen to be true because the additive inverse of this term was shown to be bounded below by a positive constant when γ_2 was defined. Then for any k satisfying $\frac{\epsilon}{2} \le V(k) \le R$, we see that $\Delta V(k) \le -\rho\delta$ for some $\delta > 0$.

Now suppose that $V(k) \le \frac{\epsilon}{2}$. Then using Equation (35) and the fact that $\rho \le \gamma_1$ we see that

$$\Delta V(k) \le \frac{\epsilon}{2}. \tag{38}$$

Here we see that $\Delta V(k) \ge 0$ only for $z(k) \in B_{\epsilon/2}(\hat{z})$ and that $\Delta V(k) \le \frac{\epsilon}{2}$ in the set $B_{\epsilon/2}(\hat{z})$. Then the conditions of Lemma 1 are satisfied with $a = b = \frac{\epsilon}{2}$ and $A = B_\epsilon(\hat{z})$. In addition, for Corollary 1 we see that for any $z(k)$ satisfying $\frac{\epsilon}{2} \le V(z(k)) \le R$, there is some $\delta > 0$ such that $\sup\{-\Delta V(k)\} \ge \rho\delta > 0$. Moreover, the set G takes the form $\{z : V(z) \le R\}$. Then the conditions of Corollary
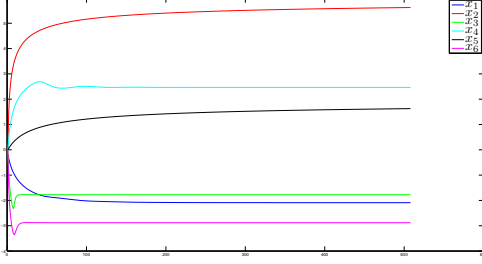
Fig. 1. The states onboard agent 1 over time. Since this is a gradient-based method with a fixed stepsize, we see larger changes in earlier iterations, followed by smaller steps taken at later iterations.
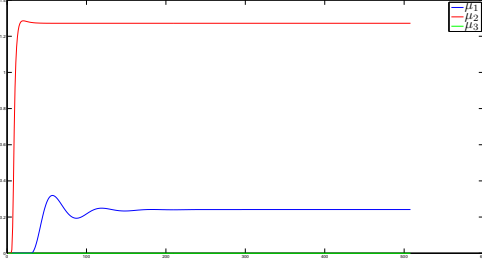

Fig. 2. The Kuhn-Tucker multipliers onboard agent 1 over time. As with the states, we see larger changes generally coming earlier in the time-evolution of the problem because they are computed using a gradient-based method using fixed stepsizes.

1 are satisfied as well. Then $z(k)$ enters $B_\epsilon(\hat{z})$ in a finite number of steps and does not ever leave thereafter. ∎

To summarize, a radially unbounded, discrete-time Lyapunov function was constructed. The Lyapunov function was shown to satisfy the conditions needed for ultimate boundedness and the system's trajectory was shown to come within ϵ of the Lagrangian's saddle point in finite time and never to be more than ϵ away thereafter.

B. Extension to Private Optimization

While above only each agent's objective function is assumed to be private, we can extend this problem to the case where individual states are kept private. To do this, we modify the initialization of the system. When the cloud sends to agent i the function $\frac{\partial g}{\partial x_i^c}$, rather than initializing agent i with $\frac{\partial g}{\partial x_i^c}$ as a function of, e.g., (x_2, x_6, x_7), it can instead initialize agent i with $\frac{\partial g}{\partial x_i^c}$ as a function of $\eta^i = (\eta_1, \eta_2, \eta_3)$, where, unbeknownst to agent i, $\eta_1 = x_2$, $\eta_2 = x_6$, and $\eta_3 = x_7$. By hiding the labels of each state which will be later sent to agent i, these states are kept private in the sense that agent i does not know which agent they belong to.

IV. SIMULATION RESULTS

A numerical implementation of the above cloud architecture was run for a particular choice of simulation example. The problem simulated was chosen to use $N = 6$ agents, each associated with a scalar state as above. The objective
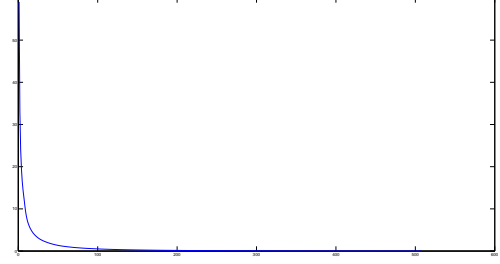

Fig. 3. The value of $V(x^1(k), \mu^1(k))$ over time. As was proven in the Lyapunov analysis in Section III, the Lyapunov function is non-increasing over time.

function of each agent was chosen to be $f_i(x_i) = (x_i - \tilde{x}_i)^4$, where

$$\tilde{x} = \begin{pmatrix} -3.0 \\ 6.0 \\ -5.0 \\ 4.0 \\ 2.0 \\ -6.0 \end{pmatrix}. \tag{39}$$

The constraints in this problem were chosen to be

$$g(x) = \begin{pmatrix} 3x_1^2 + x_4^4 - 50 \\ x_3^6 + x_6^4 - 100 \\ 9x_2 + x_5^6 - 100 \end{pmatrix} \le 0. \tag{40}$$

The Lagrangian of the full problem is

$$L(x, \mu) = \sum_{i=1}^{6} f_i(x_i) + \mu^T g(x) \tag{41}$$

where $\mu \in \mathbb{R}_+^3$.

For this example, γ_1 was found to be approximately 0.003799 and γ_2 was found to be approximately 0.001968. Accordingly, the stepsize used was $\rho = 0.0017$. The gradient descent algorithm described above was initialized with all agents and the cloud having all states set to 0. All agents and the cloud had all Kuhn-Tucker multipliers initialized to 0 as well. Here the value $\epsilon = 0.3$ was chosen.

For the purposes of analyzing and verifying the algorithm presented here, the points $\hat{x}$ and $\hat{\mu}$ were computed ahead of time to be

$$\hat{x} = \begin{pmatrix} -2.1278 \\ 5.7178 \\ -1.7745 \\ 2.4566 \\ 1.6395 \\ -2.8798 \end{pmatrix} \tag{42}$$

and

$$\hat{\mu} = \begin{pmatrix} 0.2462 \\ 1.2718 \\ 0 \end{pmatrix}. \tag{43}$$

The cloud algorithm was run for $50,000$ total iterations. It took $1,524$ iterations to enter a ball of radius ϵ about $\hat{z}$,

of which 508 were spent taking gradient descent steps and $1,016$ were spent communicating values across the network.

The value of $x^c(50,000)$ was

$$x^c = \begin{pmatrix} -2.0887 \\ 5.6219 \\ -1.7744 \\ 2.4649 \\ 1.6271 \\ -2.8799 \end{pmatrix} \qquad (44)$$

and the final value of $\mu^c(50,000)$ was

$$\mu^c = \begin{pmatrix} 0.24158 \\ 1.27176 \\ 0.00000 \end{pmatrix}. \qquad (45)$$

The final value of V in the cloud was $V(x^c(50,000), \mu^c(50,000)) = 0.0110$. Based on the definition of V, this means that the square of the Euclidean distance from $(x^c(50,000), \mu^c(50,000))$ to $(\hat{x}, \hat{\mu})$ is just 0.0110. This result confirms both that $z^c(k)$ comes within ϵ of $\hat{z}$ in finite time and that it does not go more than ϵ away from $\hat{z}$ after that.

To further illustrate the convergence of this problem, the histories of the states, Kuhn-Tucker multipliers, and value of V over time onboard agent 1 for all $50,000$ timesteps are shown in Figures 1, 2, and 3, respectively. That V is non-increasing in time was verified numerically in the MATLAB implementation and is evident in graph shown in Figure 3.

V. CONCLUSION

We presented a cloud architecture for coordinating a team of mobile agents in a distributed optimization task. Each agent has direct knowledge only of its own local objective function and its own influence upon the global constraint functions but receives occasional updates from the cloud computer containing values of each other agent's state and updated Kuhn-Tucker multipliers. Using this architecture, inequality constrained multi-agent optimization problems were proven to come within ϵ of the constrained minimum in finite time and to never be more than ϵ away thereafter. Simulation results were provided to attest to the viability of this approach.

REFERENCES

[1] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1989.

[2] L. Carlone, V. Srivastava, F. Bullo, and G. C. Calafiore. Distributed random convex programming via constraints consensus. 52(1):629–662, 2014.

[3] S. Caron and G. Kesidis. Incentive-based energy consumption scheduling algorithms for the smart grid. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 391–396, Oct 2010.

[4] Benoit Chachuat. Nonlinear and dynamic optimization: From theory to practice. Technical report, Automatic Control Laboratory, EPFL, Switzerland, 2007.

[5] Mung Chiang, S.H. Low, A.R. Calderbank, and J.C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE*, 95(1):255–312, Jan 2007.

[6] Greg Droge and Magnus Egerstedt. Proportional integral distributed optimization for dynamic network topologies. In *IEEE American Control Conference (ACC), 2014*, June 2014.

[7] Diego Feijer and Fernando Paganini. Stability of primal-dual gradient dynamics and applications to network optimization. *Automatica*, 46(12):1974–1981, December 2010.

[8] B. Gharesifard and J. Cortes. Distributed continuous-time convex optimization on weight-balanced digraphs. *Automatic Control, IEEE Transactions on*, 59(3):781–786, March 2014.

[9] Ken Goldberg and Ben Kehoe. Cloud robotics and automation: A survey of related work. Technical Report UCB/EECS-2013-5, EECS Department, University of California, Berkeley, Jan 2013.

[10] Bruce Hendrickson and Tamara G. Kolda. Graph partitioning models for parallel computing. *Parallel Comput.*, 26(12):1519–1534, November 2000.

[11] James Hurt. Some stability theorems for ordinary difference equations. *SIAM Journal on Numerical Analysis*, 4(4):582–596, 1967.

[12] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, volume 49, 1998.

[13] M. Khan, G. Pandurangan, and V.S.A. Kumar. Distributed algorithms for constructing approximate minimum spanning trees in wireless sensor networks. *Parallel and Distributed Systems, IEEE Transactions on*, 20(1):124–139, Jan 2009.

[14] A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *Automatic Control, IEEE Transactions on*, 54(1):48–61, Jan 2009.

[15] Angelia Nedic and Alex Olshevsky. Distributed optimization over time-varying directed graphs. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 6855–6860, Dec 2013.

[16] Angelia Nedic and Asuman Ozdaglar. On the rate of convergence of distributed subgradient methods for multi-agent optimization. In *Proceedings of IEEE CDC*, pages 4711–4716, 2007.

[17] Angelia Nedić and Asuman Ozdaglar. Convergence rate for consensus with delays. *Journal of Global Optimization*, 47(3):437–456, 2010.

[18] Angelia Nedic, Asuman Ozdaglar, and Pablo A Parrilo. Constrained consensus and optimization in multi-agent networks. *Automatic Control, IEEE Transactions on*, 55(4):922–938, 2010.

[19] Sikandar Samar, Stephen Boyd, and Dimitry Gorinevsky. Distributed estimation via dual decomposition. In *Proc. European Control Conference*, pages 1511–1519, 2007.

[20] H.D. Simon. Partitioning of unstructured problems for parallel processing. *Computing Systems in Engineering*, 2(23):135 – 148, 1991. Parallel Methods on Large-scale Structural Analysis and Physics Applications.

[21] Daniel E Soltero, Mac Schwager, and Daniela Rus. Decentralized path planning for coverage tasks using gradient descent adaptive control. *The International Journal of Robotics Research*, 2013.

[22] Andre Teixeira, Euhanna Ghadimi, Iman Shames, Henrik Sandberg, and Mikael Johansson. Optimal scaling of the admm algorithm for distributed quadratic programming. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 6868–6873, Dec 2013.

[23] Niki Trigoni and Bhaskar Krishnamachari. Sensor network algorithms and applications Introduction. *Philosophical Transactions of the Royal Scoeity A - Mathematical, Physical, and Engineering Sciences*, 370(1958, SI):5–10, JAN 13 2012.

[24] H. Uzawa. Iterative methods in concave programming. *Studies in Linear and Non-Linear Programming*, 1958.

[25] H. Uzawa. The kuhn-tucker theorem in concave programming. *Studies in Linear and Non-Linear Programming*, 1958.

[26] Minyi Zhong and C.G. Cassandras. Asynchronous distributed optimization with event-driven communication. *Automatic Control, IEEE Transactions on*, 55(12):2735–2750, Dec 2010.