

A Study On Distributed Model Predictive Consensus

Tamás Keviczky* and Karl Henrik Johansson**

* Delft Center for Systems and Control, Delft University of Technology,
2628 CD, Delft, The Netherlands (e-mail: t.keviczky@tudelft.nl)

** ACCESS Linneaus Center, School of Electrical Engineering, Royal
Institute of Technology (KTH), 100 44 Stockholm, Sweden, (e-mail:
kallej@ee.kth.se)

Abstract: We investigate convergence properties of a proposed distributed model predictive control (DMPC) scheme, where agents negotiate to compute an optimal consensus point using an incremental subgradient method based on primal decomposition as described in Johansson et al. [2006, 2007]. The objective of the distributed control strategy is to agree upon and achieve an optimal common output value for a group of agents in the presence of constraints on the agent dynamics using local predictive controllers. Stability analysis using a receding horizon implementation of the distributed optimal consensus scheme is performed. Conditions are given under which convergence can be obtained even if the negotiations do not reach full consensus.

1. INTRODUCTION

Engineered systems are becoming increasingly complex and larger in size, which presents a need for the distribution of decision-making processes that interact with or are part of these large-scale technologies and applications. An important problem that arises among such distributed decision-making systems (often called agents), is related to consensus-seeking and rendezvous, which has received a high level of interest in the recent literature [Olfati-Saber et al., 2007]. The consensus-seeking and rendezvous problem consists of designing distributed control strategies such that the state or output of a group of agents asymptotically converge to a common value, a consensus point, which is agreed upon either a priori or on-the-fly using some negotiation scheme. In this paper, we assume that a consensus point is not fixed in advance, but is rather determined by an optimal control problem. We focus on the combination of model predictive controllers and subgradient-based negotiation of optimal consensus (along the lines of the work in Johansson et al. [2007]), and investigate conditions for asymptotic convergence of such distributed control schemes. We propose an algorithm for distributed model predictive consensus, which guarantees convergence under reasonable assumptions given a sufficient number of subgradient iterations can be performed without interruption.

We will model agents as constrained linear dynamical systems and build on the decentralized negotiation algorithm described in Johansson et al. [2007] to compute exactly or at least approach the optimal consensus point. This negotiation algorithm relies on primal decomposition of the optimal consensus and control problem and makes use of distributed implementation of an incremental subgradient method. Each agent performs individual planning of its trajectory and negotiates with neighbors to find an optimal or near optimal consensus point, before applying a control signal.

The paper is structured as follows. Section 2 introduces the optimal consensus problem and some basic notation and assumptions. The decentralized negotiation scheme of Johansson et al. [2007] is summarized in Section 3 along with a decentralized receding horizon implementation of the optimal consensus problem. Stability of the proposed decentralized negotiation and control scheme is studied in Section 4 for both converged and interrupted negotiations. Finally, Section 5 presents our conclusions.

2. PROBLEM FORMULATION

Consider $N > 1$ dynamic agents whose dynamics are described by the following discrete-time state equations

$$\begin{aligned}x_{t+1}^i &= A^i x_t^i + B^i u_t^i, \\y_t^i &= C^i x_t^i,\end{aligned}\tag{1}$$

for $i = 1, \dots, N$, where $A^i \in \mathbb{R}^{n^i \times n^i}$, $B^i \in \mathbb{R}^{n^i \times m^i}$ and $C^i \in \mathbb{R}^{p \times n^i}$. We assume that the states and inputs of each agent are constrained to lie in polyhedral sets

$$x_t^i \in \mathcal{X}^i, \quad u_t^i \in \mathcal{U}^i, \quad t \geq 0.\tag{2}$$

Definition 1. [Johansson et al., 2007] The dynamic agents described by (1) reach *consensus* at time T if

$$\begin{aligned}y_{T+k}^i &= \theta, \quad \forall k \geq 0, \quad i = 1, \dots, N, \\u_{T+k}^i &= u_T^i, \quad \forall k \geq 0, \quad i = 1, \dots, N,\end{aligned}\tag{3}$$

where θ lies in a compact and convex set $\Theta \subset \mathbb{R}^p$.

In this paper, the consensus point θ is a vector that specifies, for example, the position and velocity the agents shall converge to.

Our objective is to find a consensus point $\theta \in \Theta \subset \mathbb{R}^p$ and a sequence of inputs u_0^i, \dots, u_{T-1}^i , with $i = 1, \dots, N$ and $u_t^i \in \mathcal{U}^i$ for all $t = 1, \dots, T-1$, such that all agent outputs are equal at time T :

$$y_T^i = \theta, \quad i = 1, \dots, N.\tag{4}$$

We will also require each agent to be at an equilibrium at time T and denote the state and control equilibrium pairs of the i -th agent corresponding to a θ value with $(x_e^i(\theta), u_e^i(\theta))$. The set of equilibria for each agent $i = 1, \dots, N$ thus will be a function of θ on the domain Θ :

$$\begin{aligned} \mathcal{E}^i(\theta) &= (x_e^i(\theta), u_e^i(\theta)) \\ &= \left\{ x \in \mathbb{R}^{n^i}, u \in \mathbb{R}^{m^i} \mid x = A^i x + B^i u, C^i x = \theta \right\}. \end{aligned} \quad (5)$$

We assume that the following cost function is associated with the i -th system:

$$\begin{aligned} V^i(x_k^i, u_k^i, \theta) &= (x_k^i - x_e^i(\theta))^\top Q^i (x_k^i - x_e^i(\theta)) \\ &\quad + (u_k^i - u_e^i(\theta))^\top R^i (u_k^i - u_e^i(\theta)), \end{aligned} \quad (6)$$

where $Q^i \in \mathbb{R}^{n^i \times n^i}$ and $R^i \in \mathbb{R}^{m^i \times m^i}$ are positive definite symmetric matrices (i.e., we penalize deviations from the equilibrium states corresponding to the consensus point and the use of control effort).

Assumption 1. Each agent dynamics (A^i, B^i) is controllable and systems $(A^i, (Q^i)^{\frac{1}{2}})$ are observable.

We then formulate the following finite-time optimal control problem at time t based on Johansson et al. [2007]:

Problem 1. Let $T > 0$ be fixed. Determine control vectors $u_{k,t}^i, k = 0, \dots, T-1$, for all $i = 1, \dots, N$ and the consensus point θ_t , which solve the following optimization problem:

$$\begin{aligned} \min_{U_t, \theta_t} \quad & \sum_{i=1}^N \sum_{k=0}^{T-1} V^i(x_{k,t}^i, u_{k,t}^i, \theta_t) \\ \text{subj. to} \quad & x_{k+1,t}^i = A^i x_{k,t}^i + B^i u_{k,t}^i, \quad (7a) \\ & y_{k,t}^i = C^i x_{k,t}^i, \\ & x_{k,t}^i \in \mathcal{X}^i, \quad k = 1, \dots, T, \quad (7b) \\ & u_{k,t}^i \in \mathcal{U}^i, \quad k = 0, \dots, T-1, \quad (7c) \\ & y_{T,t}^i = \theta_t, \quad (7d) \\ & x_{T,t}^i = x_e^i(\theta_t), \quad (7e) \\ & x_{0,t}^i = x_t^i, \quad (7f) \\ & i = 1, \dots, N, \\ & \theta_t \in \Theta, \quad (7g) \end{aligned}$$

where $U_t \triangleq [u_{0,t}, \dots, u_{T-1,t}] \in \mathbb{R}^T \sum_i m^i$ with $u_{k,t} \triangleq [u_{k,t}^1, \dots, u_{k,t}^N]$, denotes part of the optimization vector containing control inputs, $x_{k,t}^i$ denotes the state vector of the i -th agent predicted at time $t+k$ obtained by starting from the state x_t^i and applying to system (1) the input sequence $u_{0,t}^i, \dots, u_{k-1,t}^i$. The full optimization vector consists of the vector U_t defined above and the consensus variable θ_t . The subscript t will be significant later in Section 3, when this problem will be solved repeatedly in a receding horizon fashion.

By implementing the solution to Problem 1, agents reach consensus at time T in the sense of Definition 1. We will make the following assumptions on the feasibility of reaching the consensus point by all agents:

Assumption 2. The rendezvous time horizon T is large enough so that all θ_t in the set Θ are feasible, i.e., reachable consensus equilibrium points for all agents.

Assumption 3. For all $\theta_t \in \Theta$ and $i = 1, \dots, N$, there exists a sequence u_0^i, \dots, u_{T-1}^i in the relative interior of \mathcal{U}^i such that $y_T^i = \theta$.

This means that it should be possible to reach θ_t without saturating the control signal (not necessarily in an optimal way).

The solution of Problem 1 was distributed among the agents in Johansson et al. [2007] by using primal decomposition in combination with an incremental subgradient method [Bertsekas et al., 2003]. First, a multiparametric solution of the individual optimization problems was defined as

$$\begin{aligned} q^i(x_t^i, \theta_t) &= \min_{U_t} \sum_{k=0}^{T-1} V^i(x_{k,t}^i, u_{k,t}^i, \theta_t) \\ \text{subj. to} \quad & (7a) - (7g), \quad k = 1, \dots, T-1. \end{aligned} \quad (8)$$

The optimal consensus problem in (7) can then be written as

$$\begin{aligned} q^*(x_t) &= \min_{\theta_t} \sum_{i=1}^N q^i(x_t^i, \theta_t) \\ \text{subj. to} \quad & \theta_t \in \Theta. \end{aligned} \quad (9)$$

The set of optimal consensus points is defined as

$$\Theta_t^* = \left\{ \theta_t \in \Theta \mid \sum_{i=1}^N q^i(x_t^i, \theta_t) = q^*(x_t) \right\}. \quad (10)$$

It can be established that the cost function $q^i(\cdot)$ defined in (8) is a convex function and a subgradient g^i for $q^i(\cdot)$ at θ_t is given by the Lagrange multipliers corresponding to the terminal point constraint.

A principal method for solving problem (8) is the subgradient method

$$\theta_t(k+1) = \mathcal{P}_\Theta \left[\theta_t(k) - \alpha(k) \sum_{i=1}^N g^i(k) \right] \quad (11)$$

where $g^i(k)$ is a subgradient of q^i at $\theta_t(k)$, $\alpha(k)$ is a positive stepsize, and \mathcal{P}_Θ denotes projection on the set $\Theta \subset \mathbb{R}^p$. In the following, we will consider the incremental subgradient method proposed in Nedić and Bertsekas [2001b]. It is similar to the standard subgradient method (11), the main difference being that at each iteration k , θ_t is changed incrementally, through a sequence of N steps. Each step is a subgradient iteration for a single component function q^i , and there is one step per component function. Thus, an iteration can be viewed as a cycle of N subiterations. If $\theta_t(k)$ is the vector obtained after k cycles, the vector $\theta_t(k+1)$ obtained after one more cycle is

$$\theta_t(k+1) = \vartheta_t^N(k), \quad (12)$$

where $\vartheta_t^N(k)$ is obtained after the N steps

$$\begin{aligned} \vartheta_t^i(k) &= \mathcal{P}_\Theta [\vartheta_t^{i-1}(k) - \alpha(k) g^i(k)], \\ g^i(k) &\in \partial q^i(x_t^i, \vartheta_t^{i-1}(k)), \quad i = 1, \dots, N, \end{aligned} \quad (13)$$

starting with

$$\vartheta_t^0(k) = \theta_t(k), \quad (14)$$

where $\partial q^i(x_t^i, \vartheta_t^{i-1}(k))$ denotes the subdifferential (set of all subgradients) of q^i at the point $\vartheta_t^{i-1}(k)$. The updates described by (13) are referred to as the subiterations of the k -th cycle.

We will make the following assumptions, which will allow us to formulate well-posed problems and characterize the number of subgradient iterations needed for convergence to a certain tolerance.

Assumption 4. (Existence of an Optimal Solution). The optimal solution set Θ_t^* is nonempty.

Assumption 5. (Subgradient Boundedness). There exists a scalar β such that

$$\begin{aligned} \|g^i\| &\leq \beta, \\ \forall g^i &\in \partial q^i(x_t^i, \theta_t(k)) \cup \partial q^i(x_t^i, \vartheta_t^{i-1}(k)), \\ i &= 1, \dots, N, \quad k \geq 0, \end{aligned} \quad (15)$$

where N is the number of subiterations in each cycle.

Since we assume that the set Θ is compact, Assumptions 4 and 5 are automatically satisfied.

Definition 2. We will denote the Euclidean distance from a point z to the set Θ_t^* by $\mathbf{dist}(z, \Theta_t^*)$.

Definition 3. A function $\gamma(\cdot)$, defined on nonnegative reals, is a K function if it is continuous, strictly increasing with $\gamma(0) = 0$.

In the next section, we briefly describe the agreement mechanism of Johansson et al. [2007] and propose a closed-loop feedback control policy, which can be used in a receding horizon fashion, interleaved with subgradient-based negotiation of optimal consensus point updates.

3. DECENTRALIZED NEGOTIATION AND RECEDING HORIZON IMPLEMENTATION SCHEME

The optimal consensus point θ_t^* can be computed in a distributed way using the incremental subgradient method described in (12)-(14). Reference [Johansson et al., 2007] describes an algorithm, where an estimate of the optimal consensus point is passed around between agents. Upon receiving an estimate from its neighbor, an agent solves the optimization problem (8) to evaluate its cost of reaching the suggested consensus point and to compute an associated subgradient (Lagrange multiplier of terminal point constraint). The agent then performs a subiteration by updating the consensus estimate according to (13) and passing the estimate to the next agent. Each agent only computes a subgradient with respect to its own part of the objective function and not the global objective function. The convergence of the incremental subgradient algorithm is guaranteed if the agents can be organized into a cycle graph (for more details see Johansson et al. [2007]).

Remark 1. Besides some technical assumptions given in Johansson et al. [2006], the primal decomposition scheme and convergence to the optimal solution of (7) using sequential local subgradient iterations is possible due to decoupled and independently constrained agent dynamics. Furthermore, the overall objective function is decomposable into a sum of terms that share only a single coupling variable, θ_t . Thus fixing a θ_t value in the cost and constraints separates the optimal control problem into local ones.

The control solution U_t^* corresponding to a negotiated optimal consensus point θ_t^* provides an open-loop control strategy for finite-time optimal consensus. However, this

solution is sensitive to model mismatch and disturbances, which suggests considering a receding horizon implementation and repeated solution of the finite-time optimal consensus problem due to its feedback nature. Our goal in such an approach is to guarantee constraint fulfillment and asymptotic convergence to a consensus point by repeatedly solving optimal consensus problems and implementing the first sample of the control solution.

More formally, let $U_t^* = [u_{0,t}^*, \dots, u_{T-1,t}^*]$ and θ_t^* be an optimal solution of (7) at time t . Then, the first sample of U_t^* is applied to the collection of agents:

$$u_t = u_{0,t}^*. \quad (16)$$

The optimization (7) is repeated at time $t + 1$, based on the new state x_{t+1} .

Remark 2. Stability of such a combination of DMPC and incremental subgradient methods is not a trivial question, especially since the terminal constraint value in the receding horizon scheme based on (7) is an optimization variable as well. The main point of the following investigation is to rule out a scenario where repeatedly solving and implementing the first step of a finite-time optimal control solution with changing terminal constraint value eventually results in divergence or lack of stability. Compared to the work in Johansson et al. [2006], this question arises because we are no longer considering only the open-loop implementation of a control sequence that terminates with the value $u_e(\theta_t^*)$ at time T , but one that is updated every time step (along with θ_t^*), based on new measurements in a receding horizon fashion.

4. STABILITY ANALYSIS

In this section we will be primarily interested in establishing conditions for asymptotic convergence of the combined DMPC and consensus algorithm to the set of equilibria defined as

$$\mathcal{E} = (\mathcal{E}^i(\theta), \dots, \mathcal{E}^N(\theta), \theta), \quad \theta \in \Theta. \quad (17)$$

4.1 Fully Converged Negotiations

For now, we will assume that in each implementation cycle (i.e., at sampling time t), the distributed negotiations on the optimal consensus value θ_t^* have converged before the implementation of the corresponding control actions. In other words, the optimal solution of problem (7) is attained by every agent in each time step by means of the distributed consensus algorithm of Johansson et al. [2006]. This allows us to consider the overall system as a whole for stability analysis, using the following aggregate dynamics

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t, \\ y_t &= Cx_t, \end{aligned} \quad (18)$$

where $A = \mathbf{diag}(A^i) \in \mathbb{R}^{\sum_i n^i \times \sum_i n^i}$, $B = \mathbf{diag}(B^i) \in \mathbb{R}^{\sum_i n^i \times \sum_i m^i}$ and $C = \mathbf{diag}(C^i) \in \mathbb{R}^{pN \times \sum_i n^i}$. The states and inputs of the overall system are constrained by

$$x_t \in \mathcal{X} = \prod_i \mathcal{X}^i, \quad u_t \in \mathcal{U} = \prod_i \mathcal{U}^i, \quad t \geq 0, \quad (19)$$

where the symbol \prod denotes the standard Cartesian product of sets. Note that according to (4), consensus for the aggregate system dynamics means $y_T = Cx_T = 1_N \otimes \theta_t^*$.

Stability analysis in this case pertains to the study of the receding horizon control scheme given in (7) and (16) with a terminal point constraint to one of its optimization variables θ_t . This will be performed next.

The set of states at time k feasible for Problem 1 is given by

$$\mathcal{X}_k = \{x \mid \exists u \in \mathcal{U} \text{ such that } Ax + Bu \in \mathcal{X}_{k+1}\} \cap \mathcal{X},$$

with (20)

$$\mathcal{X}_{T-1} = \{x \mid \exists u \in \mathcal{U} \text{ and } \theta \in \Theta \text{ such that } x = Ax + Bu \text{ and } C(Ax + Bu) = 1_N \otimes \theta\} \cap \mathcal{X}.$$

Denote with

$$c(x_t) = u_{0,t}^*, \quad (21)$$

the control law obtained by applying the receding horizon control policy in (7) and (16) with cost function (6) for each agent, when the current state is $x_t = [x_t^1, \dots, x_t^N]$. Consider the aggregate dynamical model (18) and denote with

$$x_{t+1} = Ax_t + Bc(x_t), \quad (22)$$

the closed-loop dynamics of the entire system. In the following theorem, we state sufficient conditions for the asymptotic convergence of the closed-loop system to the set of equilibria \mathcal{E} .

Theorem 1. Assume that

(A0) $Q^i \succ 0, R^i \succ 0$ for all $i = 1, \dots, N$.

(A1) For all $\theta_t \in \Theta$ there exists a unique equilibrium $x_e^i(\theta_t) \in \mathcal{X}^i, u_e^i(\theta_t) \in \mathcal{U}^i$ for all $i = 1, \dots, N$ such that $x_e^i = A^i x_e^i + B^i u_e^i$ and $C^i x_e^i = \theta_t$.

(A2) The state and input constraint sets \mathcal{X}^i and \mathcal{U}^i contain all x_e^i and u_e^i equilibrium pairs in their interior, respectively, for all $i = 1, \dots, N$.

Then, the closed-loop system (22) asymptotically converges to the set of equilibria \mathcal{E} with domain of attraction \mathcal{X}_0 .

Proof: We introduce the following notation:

$$J^i(x_t^i, U_t^i, \theta_t) = \sum_{k=0}^{T-1} V^i(x_{k,t}^i, u_{k,t}^i, \theta_t) \quad (23)$$

and

$$J(x_t, U_t, \theta_t) = \sum_{i=1}^N J^i(x_t^i, U_t^i, \theta_t). \quad (24)$$

The optimal value function obtained from solving problem (7) at time t will thus be denoted as $J^*(x_t, U_t^*, \theta_t^*)$.

We will show first that the optimal value function $J^*(x_t, U_t^*, \theta_t^*)$ decreases along the closed-loop trajectories of the overall system at each time step $J^*(x_{t+1}, U_{t+1}^*, \theta_{t+1}^*) \leq J^*(x_t, U_t^*, \theta_t^*)$, if the assumptions of the theorem hold.

Let the initial state at time t be $x_t = x_{0,t} \in \mathcal{X}_0$ and let $U_t^* = [u_{0,t}^*, \dots, u_{T-1,t}^*]$ and θ_t^* be the optimizers of problem (7). Denote with $\mathbf{x}_t^* = [x_{0,t}, x_{1,t}^*, \dots, x_{T,t}^*]$ the corresponding optimal state trajectory, with $1_N \otimes \theta_t^* = Cx_{T,t}^*$. Let $x_{t+1} = x_{1,t}^* = Ax_{0,t} + Bu_{0,t}^*$ and consider problem (7) for time $t + 1$. We will construct an upper bound for $J^*(x_{t+1}, U_{t+1}^*, \theta_{t+1}^*)$. Consider the sequence $U_{t+1} = [u_{1,t}^*, \dots, u_{T-1,t}^*, v]$ and the corresponding state trajectory resulting from the initial state $x_{t+1}, \mathbf{x}_{t+1} = [x_{1,t}^*, \dots, x_{T,t}^*, Ax_{T,t}^* + Bv]$. The input U_{t+1} will be feasible

for the problem at $t + 1$ if and only if $v \in \mathcal{U}$ keeps $C(Ax_{T,t}^* + Bv)$ equal to some $1_N \otimes \theta$ with $\theta \in \Theta$ at step T of the prediction, i.e., $C(Ax_{T,t}^* + Bv) = 1_N \otimes \theta$. Such v exists by hypothesis (A1). Since $x_{T,t}^*$ is an equilibrium of the system, this also allows us to choose a feasible v , for which in fact $C(Ax_{T,t}^* + Bv) = 1_N \otimes \theta_t^*$. This is accomplished by noticing that $x_{T,t}^* = x_e(\theta_t^*)$ and selecting

$$v = u_e(\theta_t^*). \quad (25)$$

$J(x_{t+1}, U_{t+1}, \theta_t^*)$ will be an upper bound for the optimal $J^*(x_{t+1}, U_{t+1}^*, \theta_{t+1}^*)$. Since trajectories generated by U_t^* and U_{t+1} overlap (except for the first and last sampling intervals), it is immediate to show that

$$\begin{aligned} & J^*(x_{t+1}, U_{t+1}^*, \theta_{t+1}^*) \\ & \leq J(x_{t+1}, U_{t+1}, \theta_t^*) \\ & = J^*(x_t, U_t^*, \theta_t^*) - (x_{0,t} - x_e(\theta_t^*))^\top Q(x_{0,t} - x_e(\theta_t^*)) \\ & \quad - (u_{0,t}^* - u_e(\theta_t^*))^\top R(u_{0,t}^* - u_e(\theta_t^*)) \\ & \quad + ((Ax_{T,t}^* + Bv) - x_e(\theta_t^*))^\top Q((Ax_{T,t}^* + Bv) - x_e(\theta_t^*)) \\ & \quad + (v - u_e(\theta_t^*))^\top R(v - u_e(\theta_t^*)), \end{aligned} \quad (26)$$

where $Q = \mathbf{diag}(Q^i) \in \mathbb{R}^{\sum_i n^i \times \sum_i n^i}$, $R = \mathbf{diag}(R^i) \in \mathbb{R}^{\sum_i m^i \times \sum_i m^i}$. Choosing the particular v value given in (25) leads to $Ax_{T,t}^* + Bv - x_e(\theta_t^*) = 0$, so equation (26) becomes

$$\begin{aligned} & J^*(x_{t+1}, U_{t+1}^*, \theta_{t+1}^*) - J^*(x_t, U_t^*, \theta_t^*) \\ & \leq - (x_{0,t} - x_e(\theta_t^*))^\top Q(x_{0,t} - x_e(\theta_t^*)) \\ & \quad - (u_{0,t}^* - u_e(\theta_t^*))^\top R(u_{0,t}^* - u_e(\theta_t^*)) \\ & \leq -\gamma(\|(x_t - x_e(\theta), u_t - u_e(\theta))\|), \quad \forall x_t \in \mathcal{X}_t. \end{aligned} \quad (27)$$

where γ is a class K function. This inequality along with hypothesis (A0) on the matrices Q and R ensure that $J^*(x_t, U_t^*, \theta_t^*)$ decreases along the state trajectories of the closed-loop system (22) for any $x_t \in \mathcal{X}_t$. Since $J^*(x_t, U_t^*, \theta_t^*) \geq 0$ for all x_t, U_t^*, θ_t^* , it follows that $J^*(x_t, U_t^*, \theta_t^*) \rightarrow J^*$ as $t \rightarrow \infty$, where J^* is a nonnegative constant. We conclude that $J^*(x_{t+1}, U_{t+1}^*, \theta_{t+1}^*) - J^*(x_t, U_t^*, \theta_t^*) \rightarrow 0$ as $t \rightarrow \infty$ and this implies that $\gamma(\|(x_t - x_e(\theta), u_t - u_e(\theta))\|) \rightarrow 0$. From $\gamma(\cdot)$ being a K function, it follows that $x_t - x_e(\theta), u_t - u_e(\theta) \rightarrow 0$ as $t \rightarrow \infty$. \square

4.2 Interrupted Negotiations

In case the distributed negotiation process is interrupted (e.g., due to execution time constraints) or otherwise allowed to run only for a finite number of iterations before the control inputs are implemented, the θ_t^i values do not converge to a common optimal value θ_t^* . This means that individual agents will issue control commands that will guide them to possibly close but different terminal consensus points. In such a situation, we desire to find conditions under which repeated negotiation and implementation of intermediate consensus results will still allow asymptotic convergence to a *common* consensus point for each agent.

We propose an algorithm that fulfills the above objective if the subgradient iterations in subsequent time steps approach the optimal consensus point to an increasingly more accurate level and at the same time the local MPC solutions satisfy an improvement property along the closed-loop evolution of the agents' dynamics. The first requirement ensures that the mismatch between different

interrupted θ_t^i values diminishes as $t \rightarrow \infty$. The second requirement is analogous to the standard suboptimal MPC scheme in Scokaert et al. [1999], where it is established that feasibility of such an improvement constraint implies stability of the receding horizon control scheme.

In the following, we will denote the last (i.e., implemented) final consensus point reached by agent i in the subgradient negotiation process of time instance t by θ_t^i . This intermediate consensus point is not optimal for the global optimization problem (7), but due to Assumptions 2-3 it is certainly feasible for the following local problem:

$$\min_{\theta_t^i} q^i(x_t^i, \theta_t^i) \quad (28a)$$

$$\text{subj. to } \theta_t^i \in \Theta. \quad (28b)$$

Distinguishing between the local θ_t^i variables allows the original global optimization problem (9) to be restated as

$$\min_{\theta_t} \sum_{i=1}^N q^i(x_t^i, \theta_t^i) \quad (29)$$

$$\text{subj. to } \theta_t^1 = \dots = \theta_t^N \in \Theta.$$

As opposed to the fully converged subgradient scheme in Section 4.1, the θ_t^i variables do not converge to the globally optimal one, thus we cannot rely on optimality of the MPC scheme to prove global convergence. Instead, an improvement property as shown in Scokaert et al. [1999], which is required for asymptotic convergence to the set of equilibria will be formulated as

$$\sum_{i=1}^N (J^i(x_{t+1}^i, U_{t+1}^i, \theta_{t+1}^i) - J^i(x_t^i, U_t^i, \theta_t^i)) \quad (30)$$

$$\leq -\gamma(\|(x_t - x_e(\theta), u_t - u_e(\theta))\|),$$

where γ is a class K function. A feasible sequence for such a constraint always exists based on Assumptions 2-3 and the earlier developments in Section 4.1.

The value function improvement property in (30) is not sufficient for convergence to the global optimum, since the *common* terminal point constraint is missing and the local θ_t^{i*} values are in general different. Thus, if the agents' initial states are close to their local $x_e^i(\theta_t^{i*})$ equilibria, which are significantly different from each other, then any subgradient-based or other adjustment of the local terminal point constraints towards the globally optimal θ_t^* value would necessarily result in both local and global cost increase.

This suggests that an additional requirement besides the cost improvement property is needed, which ensures that the θ_t^i values will also converge over time to a common θ_t . This can be accomplished by requiring that in each iteration the subgradient-based negotiation scheme is executed at least until

$$\|\theta_t^i - \theta_t^*\| \leq \varepsilon_t \quad \forall i = 1, \dots, N, \quad (31)$$

where the approximation bound is updated for instance according to

$$\varepsilon_t \leq \varepsilon \frac{1}{t}, \quad \varepsilon > 0. \quad (32)$$

In order to have some information about the required number of incremental subgradient iterations that guarantee fulfillment of constraints (30) and (31), we will make use

of the following result from Nedić and Bertsekas [2001a]. It can be shown that under a strong convexity type assumption, the incremental subgradient method defined earlier in (12)-(14) with an appropriately chosen stepsize $\alpha(k)$ has a sublinear convergence rate:

Proposition 2. [Nedić and Bertsekas, 2001a] Let Assumptions 4 and 5 hold, and assume that there exists a positive scalar μ such that

$$q(x_t, \theta_t) - q^*(x_t) \geq \mu (\mathbf{dist}(\theta_t, \Theta^*))^2, \quad \forall \theta_t \in \Theta. \quad (33)$$

Then for the sequence $\{\theta_t(k)\}$ generated by the incremental subgradient method with the stepsize of $\alpha(k) = \frac{1}{2\mu} \frac{1}{k+1}$, $\mu > 0$, we have

$$(\mathbf{dist}(\theta_t(k+1), \Theta^*))^2 \leq \frac{1 + \ln(k+1)}{k+1} \frac{N^2 \beta^2}{4\mu^2}. \quad (34)$$

In the following, we describe a scheme, which allows the two conditions (30) and (31) to be tested based on the cyclic communication scheme underlying the subgradient-based negotiation.

In Algorithm 1, agents perform cyclic iterations of the subgradient (SG) method (12)-(14). They execute at least the number of iterations dictated by the optimal θ_t^* approximation requirement in (31). Satisfaction of the test (31) is signaled by a flag f_{SG} . If needed, agents continue with subgradient iterations until the global cost improvement property in (30) is satisfied. This is signaled by flag f_{DMPC} .

In order to accomplish this, agents pass along besides their current subiterate $\vartheta_t^i(k)$ of the consensus point in iteration k at sampling time t , the two binary variables (flags) f_{DMPC} and f_{SG} corresponding to tests (30) and (31), and two vectors of dimension N : J_{curr} and J_{prev} . These vectors contain the individual cost values associated with the current and previous sampling time, respectively. J_{prev} has values corresponding to the cost of using the final $\vartheta_{t-1}^i(k)$ consensus points for implementation during the previous sampling time $t-1$. The current cost J_{curr} gets filled up cyclically using the most recent subiterate $\vartheta_t^i(k)$ for each agent.

When an agent computes its own consensus point subiterate, it calculates the corresponding local cost value and checks the sum of previous and current cost values for each agent to decide whether the improvement property (30) is satisfied. If it is, then it sets a flag f_{DMPC} , which indicates that the improvement property (30) is fulfilled and every other agent should enter in an implementation phase, provided that condition (31) is also satisfied. The message reaches all other agents eventually as they pass along this information in a cyclic pattern. If property (30) is not satisfied, then it puts its current cost value entry in the vector J_{curr} and passes it on to the next agent.

Theorem 3. Under the assumptions of Section 2, Algorithm 1 converges asymptotically to the set of equilibria \mathcal{E} .

Proof: The main idea of the proof follows along the lines of Theorem 1, except for two crucial points. A feasible sequence for the improvement constraint (30) always exists based on Assumptions 2-3 and the developments in Section 4.1. This improvement property guarantees that even with interrupted negotiations, the distributed MPC problem converges asymptotically to some set of different

terminal points (since θ_t^i are different in this case). However, these terminal points are guaranteed to form a single consensus point, attained asymptotically by the repeated application of the iterative subgradient method, due to (31) and the compactness of Θ . \square

Algorithm 1: Cyclic incremental DMPC algorithm

```

1 Initialize  $\beta, \mu, \epsilon, \theta_0^i$ ;
2  $f_{\text{DMPC}}, f_{\text{SG}} \leftarrow \text{false}$ ;
3  $k, t \leftarrow 0$ ;
4  $J_{\text{curr}} \in \mathbb{R}^N \leftarrow \mathbf{0}$ ;
5  $J_{\text{prev}} \in \mathbb{R}^N \leftarrow -M \cdot \mathbf{1}$ ; /* M is large number */
6 loop
7   Measure states  $x_t^i$ ;
8   repeat
9      $\alpha(k) \leftarrow \frac{1}{2\mu} \frac{1}{k+1}$ ;
10     $\vartheta_t^0(k) \leftarrow \theta_t(k)$ ;
11    for  $i = 1$  to  $N$  do
12      if  $f_{\text{DMPC}} \wedge f_{\text{SG}}$  then
13         $J_{\text{prev}}^i \leftarrow J_{\text{curr}}^i(x_t^i, U_t^i, \vartheta_t^i(k-1))$ ;
14        Implement  $u_{0,t}^{i*}(\vartheta_t^i(k-1))$ ;
15         $\vartheta_t^i(k) \leftarrow \vartheta_t^i(k-1)$ 
16      else
17        Compute a  $g^i(k) \in \partial q^i(x_t^i, \vartheta_t^{i-1}(k))$ ;
18         $\vartheta_t^i(k) \leftarrow \mathcal{P}_\Theta[\vartheta_t^{i-1}(k) - \alpha(k)g^i(k)]$ ;
19         $J_{\text{curr}}^i \leftarrow J_{\text{curr}}^i(x_t^i, U_t^i, \vartheta_t^i(k))$ ;
20        if  $\sum_{i=1}^N (J_{\text{curr}}^i - J_{\text{prev}}^i) \leq \mathbf{0}$  then
21          | Set  $f_{\text{DMPC}}$  true;
22        else
23          | Set  $f_{\text{DMPC}}$  false;
24        end
25      end
26    end
27     $\theta_t(k+1) \leftarrow \vartheta_t^N(k)$ ;
28    if  $\frac{1+\ln(k+1)}{k+1} \frac{N^2\beta^2}{4\mu^2} \leq \frac{\epsilon}{t}$  then
29      | Set  $f_{\text{SG}}$  true;
30    else
31      | Set  $f_{\text{SG}}$  false;
32    end
33     $k \leftarrow k+1$ ;
34  until new measurement is available ;
35   $t \leftarrow t+1$ ;
36   $k \leftarrow 0$ ;
37 end loop

```

Remark 3. Although Algorithm 1 guarantees global convergence, it requires an increasing number of subgradient iterations in subsequent time steps in order to approach the optimal value with a decreasing tolerance. The requirements (30) and (31) are only sufficient conditions and thus might be somewhat conservative. Decreasing the initial stepsize of the subgradient iterations may solve this problem. The increasing number of subgradient iterations can also be alleviated in practice in the following way: Once ϵ_t gets small enough or another condition indicating closeness to the global consensus point is satisfied, the θ consensus point can be fixed for all agents and the scheme could proceed with a pure decentralized MPC scheme. This would ensure convergence due to the result shown in Section 4.1.

5. CONCLUSIONS

We have introduced a distributed model predictive control (DMPC) framework, where the control objective is to agree upon and achieve an optimal consensus point for constrained dynamic agents. The negotiation scheme makes use of the cyclic incremental subgradient algorithm described in Johansson et al. [2007]. Convergence properties of the combined DMPC / incremental subgradient approach were analyzed and a sufficient minimum number of subgradient iterations were established. An algorithm was proposed that ensures convergence of the decentralized scheme. A numerical implementation example representing an aerial refueling scenario can be found in Keviczky and Johansson [2008]. Other applications include distributed “synchronization”, where agents with constrained dynamics have to agree upon and achieve simultaneously an “optimal” consensus value, which is not known a priori. Our current work considers schemes that relax the cyclic, sequential communication requirement and rely on parallel, localized iterations.

ACKNOWLEDGEMENTS

The authors wish to thank Richard Murray for helpful discussions. The work of Tamás Keviczky was partially funded by the Boeing Corporation while he was a post-doctoral scholar at the Control and Dynamical Systems, California Institute of Technology. The work by Karl H. Johansson was partially funded by the Swedish Research Council, the Swedish Strategic Research Foundation, and the European Commission through the NoE HYCON. It was partially done during a stimulating sabbatical visit at Caltech, which is gratefully acknowledged.

REFERENCES

- D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.
- B. Johansson, A. Speranzon, M. Johansson, and K. H. Johansson. Distributed model predictive consensus. In *Proc. 17th Int'l Symposium on Mathematical Theory of Networks and Systems (MTNS)*, pages 2438–2444, Kyoto, Japan, July 2006.
- B. Johansson, A. Speranzon, M. Johansson, and K. H. Johansson. On decentralized negotiation of optimal consensus. *Automatica*, 2007. doi: 10.1016/j.automatica.2007.09.003.
- T. Keviczky and K. H. Johansson. A study on distributed model predictive consensus. Technical report, February 2008. <http://arxiv.org/abs/0802.4450v1>.
- A. Nedić and D. Bertsekas. Convergence rate of incremental subgradient algorithms. In S. Uryasev and P. M. Pardalos, editors, *Stochastic Optimization: Algorithms and Applications*, pages 223–264. Kluwer Academic Publishers, Dordrecht, Netherlands, 2001a.
- A. Nedić and D. Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM J. on Optimization*, 12:109–138, March 2001b.
- R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- P. O. M. Scokaert, D. Q. Mayne, and J. B. Rawlings. Suboptimal model predictive control (feasibility implies stability). *IEEE Trans. on Automatic Control*, 44(3): 648–654, March 1999.