Article:

Baldivieso-Monasterios, P.R. and Trodden, P.A. orcid.org/0000-0002-8787-7432 (2021) Coalitional predictive control: consensus-based coalition forming with robust regulation. Automatica, 125. 109380. ISSN 0005-1098

https://doi.org/10.1016/j.automatica.2020.109380

Coalitional predictive control: consensus-based coalition forming with robust regulation

Pablo R. Baldivieso-Monasterios[a], Paul A. Trodden[a]

[a]*Department of Automatic Control & Systems Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK*

Abstract

This paper is concerned with the problem of controlling a system of constrained dynamic subsystems in a way that balances the performance degradation of decentralized control with the practical cost of centralized control. We propose a coalitional control scheme in which controllers of subsystems may, as the need arises, group together into coalitions and operate as a single entity. The scheme employs a robust form of distributed model predictive control for which recursive feasibility and stability are guaranteed, yet—uniquely—the reliance on robust invariant sets is merely implicit, thus enabling applicability to higher-order systems. The robust control algorithm is combined with an algorithm for coalition forming based on consensus theory and potential games; we establish conditions under which controllers reach a consensus on the sets of coalitions. The recursive feasibility and closed-loop stability of the overall time-varying coalitional control scheme are established under a sufficient dwell time, the existence of which is guaranteed.

Keywords: Model predictive control; decentralization; switched systems

1. Introduction

Distributed and decentralized forms of model predictive control (MPC) have attracted significant attention as techniques for controlling large-scale constrained systems. Many proposals have been made, differing according to the nature or source of the coupling between subsystems, and the algorithmic approach taken to coordinate control actions (Christofides et al., 2013; Maestre and Negenborn, 2014; Scattolini, 2009).

A near-ubiquitous assumption in distributed MPC is that the system is initially *partitioned* into subsystems coupled via dynamics, constraints or objectives. A typical approach then assigns an MPC controller to each subsystem, and focuses on what communication is needed, or assumptions necessary, to ensure system-wide constraint satisfaction, stability and optimality. Depending on the partition, however, the degree and strength of coupling may change; indeed, it is well known that the system partition has fundamental implications for many aspects of control system design and operation including, *inter alia*, controllability and observability, dimensionality and complexity, communication, stability and performance (Šiljak, 1991).

A question that naturally arises, therefore, is what is the best choice of system partition for a system that is to be controlled by distributed or decentralized MPC? There is relatively little on this in the literature: early work by Motee and Sayyar-Rodsari (2003) proposed to select an optimal partition online by minimizing an unconstrained open-loop performance index; more recent contributions (Pourkargar et al., 2017; Zheng et al., 2018), considering constraints, have used community detection algorithms to decompose, offline, the system-wide optimal control problem into problems with minimal overlap and weak interactions. Similarly, (Barreiro-Gomez et al., 2019, 2017) focus on finding offline a suitable partition based on different criteria, for example minimization of communication requirements, and topologies that facilitate the computation of game theoretic metrics.

A contemporaneous and relevant development is the emergence of *coalitional control* schemes (Fele et al., 2017). Such schemes aim to design control strategies that optimize (on-line) the trade-off between control performance, complexity, and communication. This is achieved by controllers acting cooperatively or independently at different times. This is the idea pursued in the current paper.

The coalitional control literature first considered an unconstrained linear quadratic (LQ) setting and focused on analysing the benefits of controlling subsystems in coalitions. Tools from cooperative game theory—the Shapley value (Muros et al., 2018), Harsanyi power solutions (Muros et al., 2017a), and the Banzhaf value (Muros et al., 2017b)—have been applied and studied. When constraints are present, however, unconstrained LQ control may be deficient; Fele et al. (2018) therefore proposed a coalitional MPC scheme wherein constraints are handled naturally, and the system partition is determined from bargaining between predictive controllers. Under assumptions of recursive feasibility and weak coupling, input-to-state (ISS) stability was established regardless of when and which coalitions were formed.

The problem of guaranteeing recursive feasibility and

Email addresses: `p.baldivieso@sheffield.ac.uk` (Pablo R. Baldivieso-Monasterios), `p.trodden@sheffield.ac.uk` (Paul A. Trodden)

closed-loop stability in coalitional MPC is actually non-trivial. The main issue is indeed that the time-varying coalitional system is a switched or switching system for which feasibility and stability are not naturally maintained. Additionally, coalitional MPC inherits and exacerbates the fundamental challenges of distributed MPC for dynamically coupled subsystems: in order to *guarantee* feasibility and stability, the control algorithm must either rely on iteration between controllers at each sampling time (*e.g.* (Venkat et al., 2008)), or use techniques from robust MPC (*e.g.* (Farina and Scattolini, 2012; Trodden and Maestre, 2017)), even though the underlying control problem is a nominal one. The latter family of approaches are iteration-free but bring their own challenges in the coalitional setting: firstly, the dynamics of a coalition may be of high order, even if the constituent subsystem dynamics are of low order; secondly, the coalitions will vary, both in size and in membership, over time. Both of these features render impractical a control approach based on robust invariant sets, since these are prohibitively difficult to compute—even offline—for systems of anything other than low order.

With these challenges in mind, the contributions of this paper are three-fold:

- In Section 3, we present a distributed scheme for time-invariant coalitions of subsystems with guarantees of robust recursive feasibility and stability, despite the a-priori unknown disturbances arising from interactions, yet minimal reliance on invariant sets. The proposed scheme, first developed in (Baldivieso Monasterios et al., 2017), employs two MPC controllers for each subsystem. Crucially, the design and formulation of the MPC problems, and the invariance-inducing control law, does not require the explicit characterization of a robust invariant set, but relies only implicitly on the existence of one.
- We propose, in Section 4, a scheme for selecting partitions online using consensus optimization. Subsystems optimize their opinion and reach a consensus on what the system partition should be at the current state. The choice of consensus objective function is shown to be a potential function, and the consensus algorithm inherits strong properties from potential games, including finite-iteration convergence.
- In Section 5, we combine the robust control and consensus-based partition selection algorithms to produce a time-varying coalitional control scheme wherein controller re-design, in response to the system having re-organized into new coalitions, requires the solution of a linear programming problem. We study the properties of the time-varying system, and establish conditions under which it is feasible and stable.

Section 2 defines the problem. The results are illustrated in Section 6. Proofs are given in the Appendix.

Notation and basic definitions: $\mathbb{I}_{\geq 0}$ and $\mathbb{I}_{>0}$ are the sets of non-negative and positive integers; $\mathbb{I}_{a:b}$ is the set of integers between $a < b$. $\mathbb{R}_{\geq 0}$ and $\mathbb{R}_{>0}$ are the sets of non-negative and positive real numbers. $|x|$ denotes the ℓ_2-norm, $\|x\|_2$, of a vector $x \in \mathbb{R}^n$. A C-set is a compact and convex set containing the origin, while a PC-set is a C-set with the origin in its interior. For two sets A and B, the *Minkowski sum* is $A \oplus B = \{a + b : a \in A, b \in B\}$, and the *Pontryagin difference* is $A \ominus B = \{a : a + b \in A, \ \forall b \in B\}$. A set $\mathcal{R}$ is *robust control invariant* (RCI) for a system $x^+ = f(x, u, w)$ and constraints $(\mathbb{X}, \mathbb{U}, \mathbb{W})$ if (i) $\mathcal{R} \subset \mathbb{X}$ and (ii)$\forall x \in \mathcal{R}$, $\exists u = \mu(x) \in \mathbb{U}$ such that $x^+ = f(x, u, w) \in \mathcal{R}$, $\forall w \in \mathbb{W}$; the control law $u = \mu(x)$ is said to be *invariance inducing* over the set $\mathcal{R}$. The origin is locally stable for a discrete-time system $x^+ = f(x)$ if, for any $\epsilon > 0$, there exists a $\delta > 0$ such that $|x(0)| \leq \delta$ implies $\forall k \in \mathbb{I}_{>0}$, $|x(k)| \leq \epsilon$; if, in addition, for any $x(0) \in X$, $x(k) \to 0$ as $k \to \infty$ then the origin is asymptotically stable with region of attraction X; if $|x(k)| \leq c\gamma^k |x(0)|$ for all $x(0) \in X$, where $c > 0$ and $\gamma \in (0, 1)$, then the origin is exponentially stable with region of attraction X. A function $\alpha \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is said to be a $\mathcal{K}$-function if it is continuous, strictly increasing, and has $\alpha(0) = 0$.

2. Problem statement and preliminaries

2.1. The system and its partition into subsystems

We consider the problem of controlling a discrete-time, linear time-invariant system

$$x^+ = Ax + Bu, \tag{1}$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$ are the state and control input, and x^+ is the state at the next instant of time. We consider that a basic *partitioning* of (1) into a number, M, of independently actuated subsystems is known. The dynamics of subsystem $i \in \mathcal{M} \triangleq \{1, \ldots, M\}$ are

$$x_i^+ = A_{ii}x_i + B_iu_i + w_i \text{ where } w_i \triangleq \sum_{j \in \mathcal{M}_i} A_{ij}x_j,$$

and $x_i \in \mathbb{R}^{n_i}$, $u_i \in \mathbb{R}^{m_i}$ are the state and input of $i \in \mathcal{M}$, with $x = (x_1, \ldots, x_M)$, $u = (u_1, \ldots, u_M)$ the aggregate state and input respectively. The set of *neighbours* of subsystem i is $\mathcal{M}_i \triangleq \{j \in \mathcal{M} \setminus \{i\} : A_{ij} \neq 0\}$.

Assumption 1 (Controllability)**.** *For each $i \in \mathcal{M}$ the pair (A_{ii}, B_i) is controllable.*

The system is constrained via local, independent constraints on the states and inputs of each subsystem, *i.e.*, $x_i \in \mathbb{X}_i$, $u_i \in \mathbb{U}_i$ for subsystem i.

Assumption 2 (Constraint sets)**.** *The sets $\mathbb{X}_i \subset \mathbb{R}^{n_i}$ and $\mathbb{U}_i \subset \mathbb{R}^{m_i}$ are PC-sets.*

2.2. Coalitions of subsystems and partitions of the system

The setting of the paper is to consider that subsystems may grouped together into, and controlled as, *coalitions*.

Definition 1 (Coalition of subsystems)**.** *A* coalition *of subsystems is a non-empty subset of* $\mathcal{M}$*.*

The idea is that each coalition of subsystems operates and is controlled as a single entity; a coalitional controller replaces (or coordinates) the local subsystem controllers. Viewed differently, the grouping of the subsystems into coalitions induces an alternative partitioning of the system.

Definition 2 (Partition of the system)**.** *A* partition *of the system is an arrangement of the M subsystems into $C \leq M$ coalitions: formally, the partition of $\mathcal{M} = \{1, \ldots, M\}$ is the set $\mathcal{C}$, satisfying the following properties:*

1. *Coalition $c \in \mathcal{C}$ contains subsystems $c \subseteq \mathcal{M}$; the cardinality of c is M_c.*
2. *Coalitions are non-overlapping: $c \cap d = \emptyset$ for all $c \neq d$ and $c, d \in \mathcal{C}$.*
3. *Coalitions cover the set of subsystems: $\bigcup_{c \in \mathcal{C}} c = \mathcal{M}$.*

These definitions include the trivial cases of (i) a single, grand coalition of all subsystems ($C = 1$, $c_1 = \mathcal{M}$) (the *centralized partition*) and (ii) the basic partitioning of the system, in which each subsystem is a coalition ($C = M, \mathcal{C} = \mathcal{M}$, $c = \{i\}$ for each $i \in \mathcal{M}$) (the *decentralized partition*). The set of all possible partitions is

$$\Pi_{\mathcal{M}} \triangleq \{\mathcal{C} : \mathcal{C} \text{ is a partition of } \mathcal{M}\}.$$

Given a partition $\mathcal{C}$, the state and input of coalition c are, respectively, $x_c = (x_i)_{i \in c}$ and $u_c = (u_i)_{i \in c}$ [1]. The dynamics of coalition c are

$$x_c^+ = A_{cc} x_c + B_c u_c + w_c,$$

where the matrices A_{cc} and B_c contain, as sub-blocks, the matrices of subsystems within the coalition: $A_{cc} = [A_{ij}]_{i,j \in c}$, $B_c = \mathrm{diag}(B_i)_{i \in c}$. Similar to the system basic partition, the coalitions remain coupled via their dynamics: coalition c is coupled with coalition d via the matrices $A_{cd} = \left[A_{ij}\right]_{i \in c, j \in d, d \neq c}$ so that

$$w_c \triangleq \sum_{d \in \mathcal{M}_c} A_{cd} x_d \text{ where } \mathcal{M}_c \triangleq \left\{d \in \mathcal{C} \setminus \{c\} : A_{cd} \neq 0\right\}.$$

Assumption 3. *For any partition $\mathcal{C} \in \Pi_{\mathcal{M}}$, each pair (A_{cc}, B_c), for $c \in \mathcal{C}$, is controllable.*

2.3. Coalitional control problem

The aim is to solve the following optimal control problem: from a state $x(0)$, determine the control policy and coalitional policy that minimizes the bi-criterion cost

$$\underbrace{\sum_{k=0}^{\infty} x^\top(k) Q x(k) + u^\top(k) R u(k)}_{V^\infty(x(0), \mathbf{u}(0))} + J(\mathcal{C}(k), x(k)) \qquad (2)$$

with $Q \triangleq \mathrm{diag}(Q_1, \ldots, Q_M)$, $R \triangleq \mathrm{diag}(R_1, \ldots, R_M)$, while satisfying constraints $x(k) \in \mathbb{X} \triangleq \mathbb{X}_1 \times \cdots \times \mathbb{X}_M$, $u(k) \in \mathbb{U} \triangleq \mathbb{U}_1 \times \cdots \times \mathbb{U}_M$ for $k \in \mathbb{I}_{\geq 0}$. The term $J(\mathcal{C}, x)$ is supposed to measure the *practical* or *operating* cost of controlling subsystems in coalitions: it may include, for instance, costs on communication, computation and complexity.

Assumption 4 (Positive definite stage cost)**.** *Q_i and R_i are, for each $i \in \mathcal{M}$, positive definite matrices.*

The idea is to determine the infinite-horizon control sequence $\mathbf{u}(0) = \left\{u(0), u(1), u(2), \ldots\right\}$ and partition sequence $\left\{\mathcal{C}(0), \mathcal{C}(1), \ldots\right\}$ that minimize this system-wide joint cost on regulation performance and practical operation. (In contrast, the coalitional MPC scheme of Fele et al. (2018) aims to minimize individual subsystem performance costs by using coalitions and game-theoretical measures to allocate payoffs.) It has been shown, via a range of applications, that there is a potential benefit to performance of employing different coalitions over time (Maestre et al., 2014; Muros et al., 2014, 2017c). However, the optimal control problem is generally intractable—even when $J(\mathcal{C})$ is well defined—as it is an infinite-dimensional combinatorial optimization problem. Thus, in the sequel we propose a suboptimal way to solve this problem while achieving guarantees of constraint satisfaction and stability.

3. Robust MPC for time-invariant coalitions

We first consider the scenario where the set of subsystems $\mathcal{M}$ are arranged into a collection of fixed coalitions $\{c_1, c_2, \ldots, c_C\}$. The aim is for each coalition, acting as a single entity, to regulate its combined state to the origin, while respecting constraints. To this end, each coalition is equipped with a model predictive controller.[2] Owing to the presence of dynamic coupling between coalitions, manifested as the disturbance $w_c = \sum_{d \in \mathcal{M}_c} A_{cd} x_d$ for coalition c, consideration needs to be given to handling interactions adequately in order to achieve constraint satisfaction and stability (Scattolini, 2009).

Among the numerous DMPC schemes, algorithms based on robust techniques (Mayne et al., 2005) have the advantage of achieving feasibility and stability guarantees without relying on inter-agent iterations and negotiation (Farina and Scattolini, 2012; Riverso and Ferrari-Trecate, 2012; Trodden and Maestre, 2017). The fundamental ingredient for such schemes is the availability of an RCI set, $\mathcal{R}_c$, for the uncertain dynamics of each coalition that arise when the state interaction is treated as a disturbance (*i.e.*, $x_c^+ = A_{cc} x_c + B_c u_c + w_c$ with $w_c = \sum_{d \in \mathcal{M}_c} A_{cd} x_d$ a-priori unknown) along with its invariance-inducing control law $\tilde{\kappa}_c(\cdot)$. In the simplest implementation, the set $\mathcal{R}_c$ is used to

[1] Our intention is to make the notation as simple as possible by employing a single subscript to denote both a variable of a subsystem and a variable of a coalition.

[2] The MPC problem for a coalition can be solved by a single agent in the coalition (a leader), or distributed among several members, but these details are beyond the scope of this paper.

tighten the constraint sets, *i.e.*, as $\mathbb{X}_c \ominus \mathcal{R}_c$ and $\mathbb{U}_c \ominus \tilde{\kappa}_c(\mathcal{R}_c)$, in a *nominal* MPC problem that employs the disturbance-free prediction model $\bar{x}_c^+ = A_{cc}\bar{x}_c + B_c\bar{u}_c$ involving *nominal* prediction variables $\bar{x}_c$ and $\bar{u}_c$. The composite control law $u_c = \bar{\kappa}_c(\bar{x}_c) + \tilde{\kappa}_c(x_c - \bar{x}_c)$—where the first term is the implicit control law arising from the nominal MPC—bounds the mismatch between true variables (x_c, u_c) and nominal variables $(\bar{x}_c, \bar{u}_c)$, and ensures recursive feasibility of the MPC problems and stability of the closed-loop system.

As explained in the Introduction, however, the features of the coalitional control problem render these approaches impractical. To address this challenge, we adopt therefore the "nested" robust approach initially developed in (Baldivieso Monasterios et al., 2017). This approach replaces the ancillary robust control law—which usually requires knowledge of $\mathcal{R}_c$—with a secondary MPC controller. Constraint restrictions in the primary MPC formulation are achieved via simple scalings of $\mathbb{X}_c$ and $\mathbb{U}_c$ rather than the exact restrictions. We find that the closed-loop properties of the scheme rely on the *implicit* existence of an RCI set, with the implication for design and implementation that there is no need to either explicitly characterize or compute the RCI set, or impose it anywhere in the MPC constraints. Consequently, the dependency on invariant sets is minimized, while stability and feasibility guarantees are retained, making the approach more suitable for higher-order dynamics.

3.1. Primary MPC controller for coalition $c \in \mathcal{C}$

The primary controller, following conventional tube-based MPC, employs a nominal prediction model along with simple constraint restrictions. For coalition $c \in \mathcal{C}$ with (nominal) state $\bar{x}_c$, the optimal control problem is

$$\bar{\mathbb{P}}_c(\bar{x}_c):\ \min_{\bar{\mathbf{u}}_c}\left\{V_c^N(\bar{x}_c, \bar{\mathbf{u}}_c) : \bar{\mathbf{u}}_c \in \bar{\mathcal{U}}_c^N(\bar{x}_c)\right\}$$

where the decision variable $\bar{\mathbf{u}}_c \triangleq \left\{\bar{u}_c(0), \dots, \bar{u}_c(N-1)\right\}$, V_c^N is the finite-horizon regulation cost[3]

$$V_c^N(\bar{x}_c, \bar{\mathbf{u}}_c) = \sum_{j=0}^{N-1} \bar{x}_c^\top(j) Q_c \bar{x}_c(j) + \bar{u}_c^\top(j) R_c \bar{u}_c(j),$$

with $Q_c \triangleq \mathrm{diag}(Q_i)_{i\in c}$, $R_c \triangleq \mathrm{diag}(R_i)_{i\in c}$ and $\bar{\mathcal{U}}_c^N(\bar{x}_c)$ is defined by the following constraints for $j \in \mathbb{I}_{0:N-1}$:

$$\begin{aligned}
\bar{x}_c(0) &= \bar{x}_c,\\
\bar{x}_c(j+1) &= A_{cc}\bar{x}_c(j) + B_c\bar{u}_c(j),\\
\bar{x}_c(j) &\in \alpha_c^x \mathbb{X}_c,\\
\bar{u}_c(j) &\in \alpha_c^u \mathbb{U}_c,\\
\bar{x}_c(N) &= 0,
\end{aligned}$$

[3] $\bar{x}_c(j)$ denotes the prediction of state $\bar{x}_c$ at prediction step j, starting from $\bar{x}_c(0) = \bar{x}_c$, the current measurement of $\bar{x}_c$.

where $\mathbb{X}_c \triangleq \prod_{i\in c} \mathbb{X}_i$ and $\mathbb{U}_c \triangleq \prod_{i\in c} \mathbb{U}_i$. The simple choice of the origin as terminal set is to facilitate the applicability to higher-order dynamics. Selection of the scaling parameters $\alpha_c^x, \alpha_c^u \in (0,1)$ is described later.

Problem $\bar{\mathbb{P}}_c(\bar{x}_c)$ is a finite-horizon approximation to (2) for coalition c, omitting the cost $J(\mathcal{C})$. Solving this problem yields the sequence of nominal control actions $\bar{\mathbf{u}}_c^0(\bar{x}_c) \triangleq \{\bar{u}_c^0(0;\bar{x}_c), \dots, \bar{u}_c^0(N-1;\bar{x}_c)\}$. Taking the first term of the sequence and applying it to coalition c defines the implicit feedback law $\bar{\kappa}_c(\bar{x}_c) = \bar{u}_c^0(0;\bar{x}_c)$. The primary problems are solved in parallel by each coalition, and each then communicates the associated optimized state sequence $\bar{\mathbf{x}}_c^0(\bar{x}_c)$, to its neighbours $d \in \mathcal{M}_c$.

3.2. Secondary MPC controller for coalition $c \in \mathcal{C}$

Having received $\bar{\mathbf{x}}_d^0(\bar{x}_d)$ from neighbouring coalitions $d \in \mathcal{M}_c$, the controller for coalition c solves a secondary MPC problem employing a refined prediction model

$$\hat{x}_c^+ = A_{cc}\hat{x}_c + B_c\hat{u}_c + \bar{w}_c, \tag{3}$$

where $\bar{w}_c = \sum_{d\in\mathcal{M}_c} A_{cd}\bar{x}_d$ is the *planned* disturbance using the primary information obtained from the neighbours; this information forms the $(N+1)$-length sequence of future disturbances, $\bar{\mathbf{w}}_c \triangleq \{\bar{w}_c(0), \bar{w}_c(1), \dots, \bar{w}_c(N)\}$, where $\bar{w}_c(j) = \sum_{d\in\mathcal{M}_c} A_{cd}\bar{x}_d^0(j;\bar{x}_d)$.

The aim of the secondary controller is to design perturbations to the nominal $\bar{\mathbf{u}}_c^0(\bar{x}_c)$ in order to handle the planned interactions. We define the error variables $\bar{e}_c \triangleq \hat{x}_c - \bar{x}_c$ and $\bar{f}_c \triangleq \hat{u}_c - \bar{u}_c$. The secondary problem is then

$$\hat{\mathbb{P}}_c(\bar{e}_c; \bar{\mathbf{w}}_c):\ \min_{\bar{\mathbf{f}}_c}\left\{V_c^H(\bar{e}_c, \bar{\mathbf{f}}_c) : \bar{\mathbf{f}}_c \in \bar{\mathcal{F}}_c^H(\bar{e}_c; \bar{\mathbf{w}}_c)\right\}$$

where the cost function has, for but simplicity not necessity, the same structure as the one in the primary problem (albeit in terms of variables $\bar{e}_c$ and $\bar{\mathbf{f}}_c$ and a horizon H), and the set $\bar{\mathcal{F}}_c^H(\bar{e}_c; \bar{\mathbf{w}}_c)$ is defined by the following constraints for $j \in \mathbb{I}_{0:H-1}$:

$$\begin{aligned}
\bar{e}_c(0) &= \bar{e}_c,\\
\bar{e}_c(j+1) &= A_{cc}\bar{e}_c(j) + B_c\bar{f}_c(j) + \bar{w}_c(j),\\
\bar{e}_c(j) &\in \beta_c^x \mathbb{X}_c,\\
\bar{f}_c(j) &\in \beta_c^u \mathbb{U}_c,\\
\bar{e}_c(H) &= 0.
\end{aligned}$$

Similar to the primary problem, the constraint sets are scaled versions of the original sets, albeit with different scaling factors. The horizon of this problem is H; since $\bar{w}_c(N) = 0$, then setting $H \geq N+1$ will ensure that the disturbance is dealt with during the first N steps of the predictions, with the remaining $H - N$ steps allowing the predicted error to be driven to zero, as required by the terminal constraint. The solution of this problem is the sequence of controls

$$\bar{\mathbf{f}}_c^0(\bar{e}_c; \bar{\mathbf{w}}_c) \triangleq \{\bar{f}_c^0(0; \bar{e}_c, \bar{\mathbf{w}}_c), \dots, \bar{f}_c^0(H-1; \bar{e}_c, \bar{\mathbf{w}}_c)\}.$$

Since $\bar{e}_c = \hat{x}_c - \bar{x}_c$ and $\bar{f}_c = \hat{u}_c - \bar{u}_c$, selecting the first element of $\bar{\mathbf{f}}_c^0(\bar{e}_c; \bar{\mathbf{w}}_c)$ and adding to $\bar{u}_c^0(\bar{x}_c)$ yields the two-term control law

$$\bar{\kappa}_c(\bar{x}_c) + \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}_c) = \bar{u}_c^0(0; \bar{x}_c) + \bar{f}_c^0(0; \bar{e}_c, \bar{\mathbf{w}}_c)$$

that, under suitable conditions, stabilizes $\hat{x}_c = A_{cc}\hat{x}_c + B_c\hat{u}_c + \bar{w}_c$.

3.3. Overall robust controller and algorithm

Closing the loop with $u_c = \bar{\kappa}_c(\bar{x}_c) + \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}_c)$ does not, however, guarantee constraint satisfaction, feasibility and stability for the true coalition dynamics $x_c^+ = A_{cc}x_c + B_c u_c + w_c$, because part of the interaction is still neglected: the true disturbance is $w_c = \sum_{d \in \mathcal{M}_c} A_{cd}x_d$ and not the planned one, $\bar{w}_c = \sum_{d \in \mathcal{M}_c} A_{cd}\bar{x}_d$, used for predictions in the secondary MPC.

The control law is, therefore, completed with a final term $\tilde{\kappa}_c(\hat{e}_c)$—the requirements on which are given in the next section—that acts on the *unplanned* error $\hat{e}_c \triangleq x_c - \hat{x}_c$ that arises from the unplanned, residual disturbance $\hat{w}_c \triangleq w_c - \bar{w}_c$. The total error is $e_c = \bar{e}_c + \hat{e}_c = x_c - \bar{x}_c$. Since $x_c = \bar{x}_c + \bar{e}_c + \hat{e}_c$ and $u_c = \bar{u}_c + \bar{f}_c + \hat{f}_c$, the resulting three-term policy defines a feedback control law on the true state x_c:

$$u_c = \kappa_c(x_c) \triangleq \bar{\kappa}_c(\bar{x}_c) + \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}_c) + \tilde{\kappa}_c(\hat{e}_c). \quad (4)$$

This three-term control law is employed in Algorithm 1.

Algorithm 1 (MPC for coalition c)**.**
Initial data: *Sets $\mathbb{X}_c$, $\mathbb{U}_c$, $\mathcal{M}_c$; matrices A_{cd} for $d \in \mathcal{M}_c$; constants α_c^x, α_c^u, β_c^x, β_c^u; states $\bar{x}_c(0) = x_c$, $\bar{e}_c = 0$, $\bar{\mathbf{w}}_c = \mathbf{0}$, $\hat{V}_c = +\infty$.*
Online Routine:

1. *At time k, controller state $\bar{x}_c$, solve $\bar{\mathbb{P}}_c(\bar{x}_c)$ to obtain $\bar{\mathbf{u}}_c^0$ and $\bar{\mathbf{x}}_c^0$.*
2. *Transmit $\bar{\mathbf{x}}_c^0$ to $d \in \mathcal{M}_c$; having received $\bar{\mathbf{x}}_d^0$ from $d \in \mathcal{M}_c$, compute $\bar{\mathbf{w}}_c^0 = \sum_{d \in \mathcal{M}_c} A_{cd}\bar{\mathbf{x}}_d^0$.*
3. *At controller state $\bar{e}_c$, attempt to solve $\hat{\mathbb{P}}_c(\bar{e}_c; \bar{\mathbf{w}}_c^0)$ to obtain $\bar{\mathbf{f}}_c^0$: if the problem is feasible and $\hat{V}_c^0(\bar{e}_c, \bar{\mathbf{f}}_c^0) \leq \hat{V}_c$, then set $\bar{\mathbf{w}}_c = \bar{\mathbf{w}}_c^0$ and $\hat{V}_c = \hat{V}_c^0(\bar{e}_c, \bar{\mathbf{f}}_c^0)$; otherwise, solve $\hat{\mathbb{P}}_c(\bar{e}_c; \bar{\mathbf{w}}_c)$ for $\bar{\mathbf{f}}_c^0$.*
4. *Measure plant state x_c, calculate $\hat{e}_c = x_c - \bar{x}_c - \bar{e}_c$, and apply $u_c = \bar{u}_c^0(0; \bar{x}_c) + \bar{f}_c^0(0; \bar{e}_c) + \tilde{\kappa}_c(\hat{e}_c)$.*
5. *Update controller states as $\bar{x}_c^+ = A_{cc}\bar{x}_c + B_c\bar{u}_c^0(0; \bar{x}_c)$ and $\bar{e}_c^+ = A_{cc}\bar{e}_c + B_c\bar{f}_c^0(0; \bar{e}_c) + \bar{w}_c$ (where $\bar{w}_c$ is the first element in $\bar{\mathbf{w}}_c$), $\bar{\mathbf{w}}_c^+ = \{\bar{w}_c(1), \ldots, \bar{w}_c(N), 0\}$, and $\hat{V}_c^+ = \hat{V}_c - [\bar{e}_c^\top Q_c\bar{e}_c + \bar{f}_c^{0\top}(0; \bar{e}_c)R_c\bar{f}_c^0(0; \bar{e}_c)]$.*
6. *Wait one time step; set $k = k + 1$, $\bar{x}_c = \bar{x}_c^+$, $\bar{e}_c = \bar{e}_c^+$, $\bar{\mathbf{w}}_c = \bar{\mathbf{w}}_c^+$, $\hat{V}_c = \hat{V}_c^+$, and go to Step 1.*

3.4. Closed-loop properties

Recursive feasibility and stability of the algorithm were established in Baldivieso Monasterios et al. (2017), and are here tailored to the coalitional setting. If the procedure for the design of the final term in the control law and the constraint scaling parameters given in Appendix A terminates, then the scaling factors satisfy the following assumptions:

Assumption 5. *The control law $\hat{f}_c = \tilde{\kappa}_c(\hat{e}_c)$ is invariance inducing over a set $\mathcal{R}_c$ that is RCI for the system $\hat{e}_c^+ = A_{cc}\hat{e}_c + B_c\hat{f}_c + \hat{w}_c$ and constraint set $(\xi_c^x\mathbb{X}_c, \xi_c^u\mathbb{U}_c, \hat{\mathbb{W}}_c)$, for some $\xi_c^x \in [0, 1)$ and $\xi_c^u \in [0, 1)$, and where $\hat{\mathbb{W}}_c \triangleq \bigoplus_{d \in \mathcal{M}_c}(1 - \alpha_d^x)A_{cd}\mathbb{X}_d$.*

Assumption 6. *The constants $(\alpha_c^x, \beta_c^x, \xi_c^x)$ and $(\alpha_c^u, \beta_c^u, \xi_c^u)$ satisfy $\alpha_c^x + \beta_c^x + \xi_c^x \leq 1$ and $\alpha_c^u + \beta_c^u + \xi_c^u \leq 1$.*

To aid the statement of the results, we make the following definitions: $\bar{\mathbb{W}}_c = \bigoplus_{d \in \mathcal{M}_c} \alpha_d^x A_{cd}\mathbb{X}_d$ is the set of disturbances arising from admissible state predictions ($\bar{x}_d \in \alpha_d^x\mathbb{X}_d$) for coalitions $d \in \mathcal{M}_c$; the set $\bar{\mathcal{W}}_c^N \triangleq \bar{\mathbb{W}}_c \times \bar{\mathbb{W}}_c \times \cdots \times \bar{\mathbb{W}}_c \times \{0\}$ corresponds to admissible state *sequences*. Given a disturbance sequence $\bar{\mathbf{w}}_c = \{\bar{w}_c(0), \ldots, \bar{w}_c(N-1), 0\} \in \bar{\mathcal{W}}_c^N$, $\bar{\mathbf{w}}_c^+ = \{\bar{w}_c(1), \ldots, \bar{w}_c(N-1), 0, 0\}$ is the tail of that sequence plus a terminal zero. The domain of $\bar{\mathbb{P}}_c(\bar{x}_c)$ is $\bar{\mathcal{X}}_c^N \triangleq \{\bar{x}_c : \bar{\mathcal{U}}_c^N(\bar{x}_c) \neq \emptyset\}$, while the corresponding domain of the secondary problem $\hat{\mathbb{P}}_c(\bar{e}_c; \bar{\mathbf{w}}_c)$, which depends on the parameter $\bar{\mathbf{w}}_c$, is $\bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c) \triangleq \{\bar{e}_c : \bar{\mathcal{F}}_c^H(\bar{e}_c; \bar{\mathbf{w}}_c) \neq \emptyset\}$. Recursive feasibility for time-invariant coalitions is then established in the following proposition.

Proposition 3.1 (Recursive feasibility)**.** *Suppose that Assumptions 1–6 hold. Then, for each coalition $c \in \mathcal{C}$:*

(i) If $\bar{x}_c \in \bar{\mathcal{X}}_c^N$ then $A_{cc}\bar{x}_c + B_c\bar{\kappa}_c(\bar{x}_c) \in \bar{\mathcal{X}}_c^N$.

(ii) If $\bar{e}_c \in \bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c)$ for some $\bar{\mathbf{w}}_c \in \bar{\mathcal{W}}_c^N$, then $A_{cc}\bar{e}_c + B_c\hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}_c) + \bar{w}_c \in \bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c^+)$, where $\bar{w}_c = \bar{\mathbf{w}}_c(0)$.

(iii) If $\bar{x}_c(0) = x_c(0) \in \bar{\mathcal{X}}_c^N$ then the coalition dynamics $x_c^+ = A_{cc}x_c + B_c u_c + w_c$ under the control law $u_c = \bar{\kappa}_c(\bar{x}_c) + \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}_c) + \tilde{\kappa}_c(\hat{e}_c)$ satisfy $x_c(k) \in \mathbb{X}_c$ and $u_c(k) \in \mathbb{U}_c$ for $k \in \mathbb{I}_{\geq 0}$.

Stability then follows under the following assumption, which ensures that once the state has entered the robust invariant set around the origin, the collection of invariance-inducing control laws then bring the state asymptotically to the origin.

Assumption 7 (Decentralized stabilizability)**.** *The collection of invariance-inducing control laws $\hat{f}_{\mathcal{C}} = \{\tilde{\kappa}_c(x_c)\}_{c \in \mathcal{C}}$, asymptotically stabilize the system $x^+ = Ax + B\hat{f}_{\mathcal{C}}$ in a neighbourhood $\mathcal{R}_{\mathcal{C}} = \prod_{c \in \mathcal{C}} \mathcal{R}_c$ of the origin.*

Theorem 3.1 (Stability)**.** *Suppose that Assumptions 1–7 hold. Then, for each $c \in \mathcal{C}$, the origin is exponentially stable for the nominal coalition system $\bar{x}_c^+ = A_{cc}x_c + B_c\bar{\kappa}_c(\bar{x}_c)$*

and asymptotically stable for the true coalition system $x_c^+ = A_{cc}x_c + B_c\kappa_c(x_c) + \sum_{d\in\mathcal{M}_c} A_{cd}x_d$. *The region of attraction for* $(\bar{x}_c, x_c)$ *is* $\bar{\mathcal{X}}_c^N \times \bar{\mathcal{X}}_c^N$.

Finally, we note some consequences of these results.

Corollary 3.1.1. *For each* $c \in \mathcal{C}$, *the sets* $\bar{\mathcal{X}}_c^N$ *and* $\bar{\mathcal{X}}_c^{N-1}$ *are positively invariant for the nominal dynamics* $\bar{x}_c^+ = A_{cc}\bar{x}_c + B_c\bar{\kappa}_c(\bar{x}_c)$.

Corollary 3.1.2. *For each* $c \in \mathcal{C}$, *starting from* $e_c(0) = 0$ *the error dynamics* $e_c^+ = A_{cc}e_c + B_c\left(\hat{\kappa}_c(\bar{e}_c;\bar{\mathbf{w}}_c) + \tilde{\kappa}_c(\hat{e}_c)\right) + w_c$ *evolve in a robust positively invariant set* $\mathcal{ER}_c \subseteq (\beta_c^x + \xi_c^x)\mathbb{X}_c$.

The set $\mathcal{ER}_c$ is difficult to characterize, given its dependence on the feasibility set of the secondary MPC controller, $\bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c)$, with the latter parameterized by $\bar{\mathbf{w}}_c$. Nevertheless, the basic principle of tube-based robust MPC—that the trajectory of the uncertain system is contained within a tube around the trajectory of the nominal system—holds.

4. Selection of the system partition

In this section we consider the problem of choosing a suitable partition for the system. The overall cost in (2) measures the performance and practical costs of using different partitions over time. However, the infinite-horizon combinatorial optimization problem implied by minimizing (2) is intractable. Our approach is therefore to decouple the problems of regulation and partition selection: the previous section presented a regulation algorithm for fixed coalitions; in this section, we develop a partition selection algorithm, using consensus optimization, to select the system partition at a fixed state; in Section 5, the partition selection and regulation algorithms are combined to produce the overall approach, which varies the system partition in time while regulating its states.

4.1. Consensus optimization problem for partition selection

The solution we propose to selecting the system partition is as follows: each subsystem $i \in \mathcal{M}$ has an initial opinion on the system partition, then an iterative process begins where subsystems exchange information until a consensus on the system partition is reached. With the system at a state x, we define the following consensus optimization problem for subsystem $i \in \mathcal{M}$:

$$\min\left\{J_i(\mathcal{C}_{[i]};\mathcal{C}_{[-i]},x)\colon \mathcal{C}_{[i]} \in \Pi_{\mathcal{M}}\right\} \tag{5}$$

where $J_i(\mathcal{C}_{[i]};\mathcal{C}_{[-i]},x) = J_i^{\text{consensus}} + \rho J_i^{\text{power}}$ and

$$J_i^{\text{consensus}} \triangleq \sum_{j\in\mathcal{M}_i} w_{ij}(x_i,x_j)\left|\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[j]})\right|$$

$$J_i^{\text{power}} \triangleq \sum_{j\in\mathcal{M}_i} (w_{ij}(x_i,x_j)+\epsilon)\sigma_{ij}\left(1-\delta_{ij}(\mathcal{C}_{[i]})\right)$$

In this problem, the decision variable $\mathcal{C}_{[i]}$ is subsystem i's opinion on the system partition $\mathcal{C}$; $\mathcal{C}_{[-i]} \triangleq \left\{\mathcal{C}_{[j]}\right\}_{j\in\mathcal{M}_i}$ is the collection of neighbour's opinions, assumed fixed at the point of solving this problem. The collection of all opinions, written $(\mathcal{C}_{[1]},\ldots,\mathcal{C}_{[M]})$—or equivalently $(\mathcal{C}_{[i]},\mathcal{C}_{[-i]})$ for some i—is called a *profile*.

The objective function comprises two terms: the first term, as common in consensus, penalizes differences between the opinion of subsystem i and the opinions of its neighbours, via the following indicator function. For $i \in \mathcal{M}$,

$$\delta_{ij}(\mathcal{C}) = \delta_{ji}(\mathcal{C}) = \begin{cases} 0 & \text{if } (i,j) \in c \text{ for some } c \in \mathcal{C}, \\ 1 & \text{otherwise.} \end{cases}$$

Thus, $J_i^{\text{consensus}} = 0$ if subsystem i and its neighbours $j \in \mathcal{M}_i$ agree on being within, or *not* being within, the same coalition.

The second term in the objective penalizes a weighted function of *link power* of coalitions in subsystem i's opinion. The scalar $\sigma_{ij} = \sigma_{ji} > 0$ is the *power* associated with the *link* between subsystems i and j that being in a coalition together implies. In the simplest case, $\sigma_{ij} = 1$ so that $\sum_{j\in\mathcal{M}_i}\sigma_{ij}\left(1-\delta_{ij}(\mathcal{C}_{[i]})\right)$ merely counts the number of neighbours of subsystem i contained within the same coalition. More sophisticated approaches have used game-theoretic measures with the aim of the link power accurately capturing the cost of using each link versus the benefit it brings to closed-loop performance (Muros et al., 2017b).

Both terms are weighted by a state-dependent function that measures the coupling strength between a pair of subsystems:

$$w_{ij}(x_i,x_j) = w_{ji}(x_j,x_i) = \frac{1}{2}\left(\|A_{ij}\||x_j| + \|A_{ji}\||x_i|\right).$$

The overall effect is as follows: the minimization of the consensus term in the objective promotes agreement on the system partition, with higher preference for putting more tightly coupled subsystems into the same coalition. The second term, penalizing weighted link power (where $\rho, \epsilon > 0$), is included as a regularization term. The absence of this may lead to pathological cases: in particular, a solution profile $(\mathcal{C}_{[1]}^e,\ldots,\mathcal{C}_{[M]}^e)$ for which $J_i^{\text{consensus}} = 0$ does not necessarily imply a consensus, but something weaker in view of the state-dependent and possibly incomplete coupling structure. Some examples are illuminating in this regard.

Example 1 (Dependence on coupling structure)**.** *For*

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}^+ = \begin{bmatrix} 1 & a_{12} & 0 \\ a_{12} & 1 & a_{23} \\ 0 & a_{23} & 1 \end{bmatrix}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + Bu$$

with $x_i \neq 0$, $J_i^{\text{consensus}} = 0$ *implies consensus at any* $\mathcal{C}^e \in \Pi_{\mathcal{M}}$ *if* $a_{12} \neq 0$ *and* $a_{23} \neq 0$. *If, however,* $a_{23} = 0$ *then the same implies* $\mathcal{C}_{[1]}^e = \mathcal{C}_{[2]}^e$ *but not necessarily* $\mathcal{C}_{[2]}^e = \mathcal{C}_{[3]}^e$.

Example 2 (Dependence on state)**.** *For the same system, but now with* $x_1 = x_2 = x_3 = 0$*,* $J_i^{\text{consensus}} = 0$ *for* all *opinion profiles* $\{\mathcal{C}_{[1]}, \dots, \mathcal{C}_{[M]}\}$*.*

Example 3 (Regularization)**.** *For the scenario in Example 2, the profile* $\mathcal{C}_{[1]} = \mathcal{C}_{[2]} = \mathcal{C}_{[3]} = \mathcal{M}$ *(the decentralized partition) attains the minimum value of* $J_i\big(\mathcal{C}_{[i]}, \mathcal{C}_{[-i]}\big)$*, for all* i*, when all* $\epsilon\sigma_{ij} > 0$*.*

4.2. The consensus algorithm and its convergence

As a prerequisite to the consensus-based algorithm for partition selection, we establish useful properties of the *game* between subsystems that arises from the definition of the consensus optimization problems. In particular, this forms a *potential game* (Monderer and Shapley, 1996), for which strong results apply to the equilibrium actions of the game and convergence of algorithms to these solutions.

Definition 3 (Finite exact potential game (FEPG))**.** *The game defined by a set of players* $\mathcal{M} = \{1, \dots, M\}$ *with finite action sets* $\mathcal{A} = \{\mathcal{A}_i\}_{i \in \mathcal{M}}$ *and objective functions* $\{J_i : \mathcal{A}_i \times \mathcal{A}_{-i} \to \mathbb{R}\}_{i \in \mathcal{M}}$*, is a* finite exact potential game *(FEPG) if there is a* potential function $\phi \colon \mathcal{A} \to \mathbb{R}$ *such that, for all* $i \in \mathcal{M}$*,* $a_{-i} \in \mathcal{A}_{-i}$ *and* $a_i', a_i'' \in \mathcal{A}_i$

$$J_i(a_i', a_{-i}) - J_i(a_i'', a_{-i}) = \phi(a_i', a_{-i}) - \phi(a_i'', a_{-i}).$$

Definition 4 (Nash equilibrium)**.** *An action profile* $a^* = (a_1^*, \dots, a_M^*)$ *is said to be a* Nash equilibrium *of the game* $\big(\mathcal{M}, \{\mathcal{A}_i\}_i, \{J_i\}_i\big)$ *if, for all* $i \in \mathcal{M}$*,*

$$J_i(a_i^*, a_{-i}^*) = \min_{a_i \in \mathcal{A}_i} J_i(a_i, a_{-i}^*).$$

The next result then follows immediately from the choice of objective function in (5).

Theorem 4.1. *The game* $(\mathcal{M}, \{\Pi_{\mathcal{M}}\}_i, \{J_i\}_i)$ *is a finite exact potential game with potential function*

$$\phi(\mathcal{C}_{[1]}, \dots, \mathcal{C}_{[M]}) = \frac{1}{2} \sum_{i \in \mathcal{M}} J_i^{\text{consensus}} + \rho J_i^{\text{power}} \tag{6}$$

Moreover, the game admits at least one Nash equilibrium; the set of these equilibria coincides with the set of Nash equilibria for the game $(\mathcal{M}, \{\Pi_{\mathcal{M}}\}_i, \{\phi\}_i)$*.*

With regards to playing this game, we define the *constrained best response* of player i as

$$\mathcal{C}_{[i]}^0 = \arg \min_{\bar{\mathcal{C}}_{[i]} \in \mathbb{C}_\Delta(\mathcal{C}_{[i]})} \big\{ J_i(\bar{\mathcal{C}}_{[i]}; \mathcal{C}_{[-i]}, x) \big\} \tag{7}$$

where $\mathcal{C}_{[-i]}$ denotes the (fixed) opinions of i's neighbours, $\mathcal{C}_{[i]}$ is the *previous* opinion of player i, and the domain of the problem, $\mathbb{C}_\Delta(\mathcal{C}_{[i]}^-)$, is the set of *chains* of length Δ containing the previous opinion $\mathcal{C}_{[i]}$.

A chain in the poset $\Pi_{\mathcal{M}}$ is any pair of elements that are comparable under the *refinement* order relation, denoted by '$\preceq$', and defined as follows: given $\mathcal{C}, \mathcal{D} \in \Pi_{\mathcal{M}}$, the partition $\mathcal{D} \preceq \mathcal{C}$ ($\mathcal{D}$ *refines* $\mathcal{C}$, or $\mathcal{C}$ *coarsens* $\mathcal{D}$) if every member of $\mathcal{D}$ is contained in some member of $\mathcal{C}$. For example, the partition $\big\{\{1,2\},\{3,4\}\big\}$ refines $\big\{\{1,2,3,4\}\big\}$. Because the ordering is partial, however, not all pairs of partitions are comparable: an *anti-chain* of $\Pi_{\mathcal{M}}$ is a set of incomparable elements. The Hasse diagram in Figure 1 illustrates some chains and anti-chains in the partition set for $\mathcal{M} = \{1,2,3,4\}$.

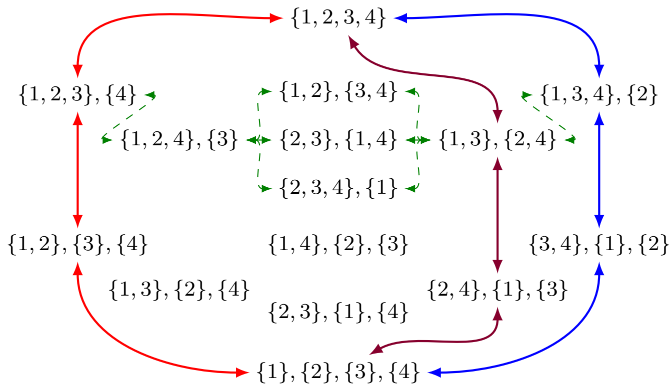

Figure 1: Refinement relation defined over $\Pi_{\{1,2,3,4\}}$. Each path of different colour points towards the elements that are related to each other, through refinement. Solid lines represent elements of chains whereas dashed lines represent members of anti-chains.

Thus, in the constrained optimization problem (7), subsystem i seeks to determine an optimal opinion $\mathcal{C}_{[i]}^0$ from a *subset* of the set of partitions $\Pi_{\mathcal{M}}$. The aim of this restriction is to reduce the cardinality of the decision space, and hence complexity of the optimization problem, since the full partition set $\Pi_{\mathcal{M}}$ grows combinatorially with the number of subsystems.

We consider the algorithm for playing the game in which players take turns to act: at iteration p, a single player i_p who stands to gain from playing its constrained best response (7) does so. Upon reaching an opinion profile where no player can unilaterally lower its cost any further, all players determine their *unconstrained* best responses, *i.e.*,

$$\mathcal{C}_{[i]}^0 = \arg \min_{\bar{\mathcal{C}}_{[i]} \in \Pi_{\mathcal{M}}} \big\{ J_i(\bar{\mathcal{C}}_{[i]}; \mathcal{C}_{[-i]}, x) \big\} \tag{8}$$

If any player can lower its cost via this unconstrained solution, then one such player adopts this solution. This additional step prevents the possibility of the algorithm stalling prior to reaching a Nash equilibrium. The algorithm then restarts with subsequent iterations employing the constrained best response (7). The algorithm terminates when no player may benefit further from playing its unconstrained best response.

Executing this algorithm defines a *path* of opinions

$$\mathfrak{C} = \big\{ (\mathcal{C}_{[i_0]}^0, \mathcal{C}_{[-i_0]}), (\mathcal{C}_{[i_1]}^0, \mathcal{C}_{[-i_1]}), \dots, (\mathcal{C}_{[i_p]}^0, \mathcal{C}_{[-i_p]}), \dots \big\}.$$

A path $\mathfrak{C}$ is called an *improvement path* if, for all $p \geq 0$, $\phi(\mathcal{C}_{[i_p]}^0, \mathcal{C}_{[-i_p]}\big) < \phi\big(\mathcal{C}_{[i_p]}, \mathcal{C}_{[-i_p]}\big)$, *i.e.*, if the player at iteration p sees a reduction in its cost. It is well known that for an FEPG the path of best responses constitutes a finite improvement path, and terminates in a Nash equilibrium; however, because here we have a path of constrained and unconstrained best responses, wherein at times the domain of each player's optimization problem (7) is restricted from $\Pi_{\mathcal{M}}$ to $\mathbb{C}_{\Delta}(\mathcal{C}_{[i]})$, the algorithm is no longer a straightforward application of an FEPG and these facts need to be established separately.

Theorem 4.2 (Finite improvement path)**.** *If $\Delta \geq 1$, then every best response path $\mathfrak{C}$ is a finite improvement path, terminating in a Nash equilibrium $(\mathcal{C}^e_{[1]}, \ldots, \mathcal{C}^e_{[M]})$ that is a minimum of the potential function* (6)*.*

The next and final result of this section gives some interesting implications about the nature of the attained equilibrium, depending on the relative weightings of the power term and consensus term in the overall cost. With the regularization term added, a minimum of (6) does not necessarily imply consensus. The following result therefore establishes that conditions on the weights in the regularization term such that consensus implies a minimum of (6), and therefore equilibrium at a consensus.

Theorem 4.3 (When consensus implies equilibrium)**.** *Suppose $(\mathcal{C}^e_{[1]}, \ldots, \mathcal{C}^e_{[M]})$ is such that $J_i^{\text{consensus}} = 0$ for all $i \in \mathcal{M}$. If $\rho\sigma_{ij} < 1$ for all i, j and $\epsilon \geq 0$ is sufficiently small, then $(\mathcal{C}^e_{[1]}, \ldots, \mathcal{C}^e_{[M]})$ is a Nash equilibrium.*

Remark 1. *We note that $\epsilon = 0$ will always be sufficiently small, but such a choice will remove the regularizing behaviour of this additional term, as Examples 2 and 3 show. It is interesting to note that Theorem 4.3 confirms that when all $x_i = 0$, any $\epsilon > 0$ and all $\sigma_{ij} > 0$ will result in a potential decrease from a point $(\mathcal{C}^e_{[1]}, \ldots, \mathcal{C}^e_{[M]})$ with $J_i^{\text{consensus}} = 0$ until the decentralized partition is reached.*

5. Coalitional MPC with time-varying coalitions

The final part of the development is to unite the coalitional regulation scheme of Section 3 with the partition selection algorithm of Section 4, and study the properties of the overall approach. We first note that the timescales of the regulation and partition selection algorithms are not coupled. Even though the regulation algorithm computes a new control input to apply to the system at every time k, there is no assumption or requirement on the rate of the partition selection algorithm: it may, for example, propose a new partition at every sampling instance or less frequently. In any case, significant challenges arise: recursive feasibility is the most basic requirement for any MPC controller, since guaranteeing the stability of the controlled system relies on the continued feasibility of the underlying optimal control problem. While recursive feasibility for each time-invariant coalition is established by Proposition 3.1, this does not continue to hold for a time-varying system partition. Even assuming this can be established, stability does not necessarily follow: the system of time-varying coalitions is akin to a system that switches between independently stable modes; it is well known that the stability of such systems can be lost through excessively fast switching. These aspects of the approach are discussed in detail in this section. Our aim is to determine conditions under which a new system partition proposed by the consensus algorithm is feasible to adopt and maintains closed-loop stability.

5.1. Recursive feasibility for time-varying coalitions

5.1.1. Definitions

The set $\bar{\mathcal{X}}_{\mathcal{C}}^N$ is defined as the product of the individual feasibility sets of the coalitions:

$$\bar{\mathcal{X}}_{\mathcal{C}}^N \triangleq \prod_{c \in \mathcal{C}} \bar{\mathcal{X}}_c^N,$$

with a similar definition for $\mathcal{ER}_{\mathcal{C}}$. Note that the condition $x \in \bar{\mathcal{X}}_{\mathcal{C}}^N$ is equivalent to the condition $x_c \in \bar{\mathcal{X}}_c^N$ for all $c \in \mathcal{C}$, but utilizes, for convenience, a more compact notation.

We also introduce notions of *feasibility* and *strong feasibility* with respect to a partition. In the context of adopting a new system partition, strong feasibility has the advantage of permitting the simple initialization of coalitional controller states that is proposed in Algorithm 1, avoiding an otherwise iterative and coupled design process.

Definition 5 (Feasible and strongly feasible partition)**.** *A partition is said to be* feasible *at a state x if $x \in \bar{\mathcal{X}}_{\mathcal{C}}^N \oplus \mathcal{ER}_{\mathcal{C}}$, and* strongly feasible *at a state x if $x \in \bar{\mathcal{X}}_{\mathcal{C}}^N$.*

5.1.2. Basic results on how partition coarsening and refinement affect feasibility

Any change in system partition proposed by the selection algorithm is either a coarsening or a refinement. The next set of results presented are therefore relevant and useful, having fundamental implications for allowing a change in system partition online. In the following, the RCI sets considered are those obtained using the design procedure of Appendix A, which are minimal (in the sense of the metric defined in the synthesis procedure) with respect to the disturbance. The first result is a consequence of the fact that, with coarsening, the disturbance set that each coalition sees diminishes, leading to a smaller RCI set. A similar relation holds with partition refinement: the disturbance set grows leading to a larger RCI set.

Theorem 5.1 (Nesting of RCI sets)**.** *Suppose that Assumptions 1–6 hold. If $\mathcal{C} \succeq \mathcal{D}$, then $\mathcal{R}_{\mathcal{C}} \subseteq \mathcal{R}_{\mathcal{D}}$.*

This, in turn, implies less restriction of the constraints in the primary MPC problems under coarsening, and more restriction with refinement.

Corollary 5.1.1 (Nesting of nominal feasibility regions)**.** *If $\mathcal{C} \succeq \mathcal{D}$, then $\bar{\mathcal{X}}_{\mathcal{C}}^i \supseteq \bar{\mathcal{X}}_{\mathcal{D}}^i$ for $i = 0 \ldots N$.*

This might seem to suggest that feasibility is trivially maintained with partition coarsening. The reality is, however, not so simple, owing to the following corollary and the fact that $x \in \mathcal{X}_{\mathcal{C}} \oplus \mathcal{ER}_{\mathcal{C}}$.

Corollary 5.1.2 (Counter-nesting of error feasibility regions). *If $\mathcal{C} \succeq \mathcal{D}$, then $\mathcal{ER}_{\mathcal{C}} \subseteq \mathcal{ER}_{\mathcal{D}}$.*

Proposition 5.1 (Feasibility is not necessarily maintained with coarsening). *Suppose $\mathcal{D} \in \Pi_{\mathcal{M}}$ is feasible at x. Then $\mathcal{C} \succeq \mathcal{D}$ is not necessarily feasible at x.*

That is, even though $\bar{\mathcal{X}}_{\mathcal{C}}^N \supseteq \bar{\mathcal{X}}_{\mathcal{D}}^N$ when $\mathcal{C} \succeq \mathcal{D}$, there is no clear relation between the sets $\mathcal{X}_{\mathcal{D}} \oplus \mathcal{ER}_{\mathcal{D}}$ and $\mathcal{X}_{\mathcal{C}} \oplus \mathcal{ER}_{\mathcal{C}}$; the example in Section 6.3 illustrates (and hence proves) this. On the other hand, *strong* feasibility is guaranteed under the same assumptions, as a consequence of Corollary 5.1.1.

Proposition 5.2 (Strong feasibility is maintained with coarsening). *Suppose $\mathcal{D} \in \Pi_{\mathcal{M}}$ is strongly feasible at a state x. Then $\mathcal{C} \succeq \mathcal{D}$ is strongly feasible at x.*

The situation is more challenging in the case of partition refinement, since a counterpart to Proposition 5.2 for a movement from $\mathcal{C}$ to $\mathcal{D} \preceq \mathcal{C}$ does not hold.

Proposition 5.3 (Strong feasibility does not imply feasibility after refinement). *Suppose $\mathcal{C} \in \Pi_{\mathcal{M}}$ is strongly feasible at a state x. Then $\mathcal{D} \preceq \mathcal{C}$ is not necessarily feasible at x.*

This raises at least two questions: firstly, *when* is the hypothesis of Proposition 5.2, for partition coarsening, met? Even though the initial state $x(0) \in \bar{\mathcal{X}}_{\mathcal{C}}^N$ implies that the nominal state $\bar{x}(k) \in \bar{\mathcal{X}}_{\mathcal{C}}^N$ for all $k \in \mathbb{I}_{>0}$, it *does not* imply that the true state $x(k) \in \bar{\mathcal{X}}_{\mathcal{C}}^N$. Secondly, when is strong feasibility achieved under partition refinement? In the next subsection, we present and discuss answers to these questions.

5.1.3. Schemes for feasible partition switching

We outline three schemes for enabling feasible switching between partitions over time. Our intention is not to develop any scheme into a comprehensive proposal, but to explore the range of options and illustrate the comparative ease or difficulty of implementing each. In the following, we assume that when a change from one partition to another is proposed, the new partition is viable in the sense that the design procedure given in Appendix A has successfully terminated, meeting the formal assumptions stated earlier in the paper. We further note that that executing this design procedure requires solving two LPs per coalition plus some elementary operations.

A quest for feasibility by design. With the system in a partition $\mathcal{C}$ that is feasible at a state x, the nominal state $\bar{x} \in \bar{\mathcal{X}}_{\mathcal{C}}^N$. Under Algorithm 1, the successor nominal state $\bar{x}^+ \in \bar{\mathcal{X}}_{\mathcal{C}}^{N-1}$ and the true state $x^+ \in \bar{\mathcal{X}}_{\mathcal{C}}^{N-1} \oplus \mathcal{ER}_{\mathcal{C}} \subset \bar{\mathcal{X}}_{\mathcal{C}}^N \oplus \mathcal{ER}_{\mathcal{C}}$. This motivates the following proposition concerning switching between partitions, offering two possible ways to ensure strong feasibility.

Proposition 5.4. *Suppose partition $\mathcal{C} \in \Pi_{\mathcal{M}}$ is feasible at a state x. Partition $\mathcal{D} \in \Pi_{\mathcal{M}}$ is strongly feasible for the successor state x^+ if $\bar{\mathcal{X}}_{\mathcal{C}}^{N-1} \oplus \mathcal{ER}_{\mathcal{C}} \subseteq \bar{\mathcal{X}}_{\mathcal{D}}^N$, which is satisfied if*

$$\bar{\mathcal{X}}_{\mathcal{C}}^{N-1} \oplus \prod_{c \in \mathcal{C}} (\beta_c^x + \xi_c^x)\mathbb{X}_c \subseteq \bar{\mathcal{X}}_{\mathcal{D}}^N.$$

With respect to the *usefulness* of this result, the first inclusion is not straightforward to verify or enforce, in view of the difficulty of characterizing $\mathcal{ER}_{\mathcal{C}}$. The second inclusion offers a more practical—since the scalars β_c^x and ξ_c^x are generated from the design procedure in Appendix A—yet more conservative way to guarantee strong feasibility. However, the condition is still problematic to impose, and perhaps impossible to meet, as a design constraint because of the fundamental relations governing the relations between sets under refinement and coarsening. For example, for the condition to be met under the refinement $\mathcal{D} \preceq \mathcal{C}$, it is necessary that $\bar{\mathcal{X}}_{\mathcal{C}}^{N-1} \subset \text{interior}(\bar{\mathcal{X}}_{\mathcal{D}}^N)$ even though $\bar{\mathcal{X}}_{\mathcal{C}}^N \supseteq \bar{\mathcal{X}}_{\mathcal{D}}^N$; possibilility of satisfaction would be highly problem specific and, even if possible, would require careful design. Even under coarsening, for which $\bar{\mathcal{X}}_{\mathcal{C}}^N \subseteq \bar{\mathcal{X}}_{\mathcal{D}}^N$ already, the condition is not trivially met and relies on weak coupling for satisfaction—the size of the coalitional disturbance set $\mathbb{W}_c = \bigoplus_{d \in \mathcal{M}_c} A_{cd}\mathbb{X}_d$ must be sufficiently small. More constructive alternatives are therefore discussed next.

Use of a feasibility dwell time. An attractive option, well established in the switched systems literature, and more recently in the context of MPC for switched systems (Hernandez Vicente and Trodden, 2019; Müller et al., 2012; Zhang et al., 2016), is the use of a *dwell time* to ensure the state lies within the feasibility region for the new partition at the moment of switching. The next result, which follows directly from the stability of each coalition in the time-invariant setting (Theorem 3.1), enables this.

Proposition 5.5 (Feasibility becomes and remains strong feasibility). *Suppose the system is in a partition $\mathcal{C} \in \Pi_{\mathcal{M}}$ that is feasible at x, and controlled by Algorithm 1. The same partition $\mathcal{C}$ becomes, and remains, strongly feasible a finite number of timesteps thereafter. Moreover, if $\mathcal{ER}_{\mathcal{C}} \subset \bar{\mathcal{X}}_{\mathcal{C}}$, then this happens exponentially fast.*

The hypothesis $\mathcal{ER}_{\mathcal{C}} \subset \bar{\mathcal{X}}_{\mathcal{C}}$ is satisfied if the constraint scaling factors follow $\beta_c^x + \xi_c^x < \alpha_c^x$ for all $c \in \mathcal{C}$; note that this is, again, a weak coupling requirement.

Once strong feasibility is established for all subsequent times, a similar result establishes that a switch from partition $\mathcal{C}$ to partition $\mathcal{D}$ is, if the coupling is sufficiently weak, possible after a finite number of steps.

Proposition 5.6 (Strong feasibility dwell time). *Suppose the system is in a partition $\mathcal{C} \in \Pi_{\mathcal{M}}$ that is feasible at x. A partition $\mathcal{D} \neq \mathcal{C}$ becomes strongly feasible a finite number of time steps thereafter. Moreover, if $\mathcal{R}_{\mathcal{C}} \subset \mathcal{X}_{\mathcal{D}}^N$, then this happens exponentially fast.*

It is not necessary to implement a candidate partition. A key observation is that, unlike in the case of switched systems where mode-to-mode switches are not necessarily something that can be controlled, the choice of system partition at each time is a controllable degree of freedom. It follows, then, that if the system partition is $\mathcal{C}$ and, subsequently, a new partition $\mathcal{D}$ is selected by the decision-making process, it is not *necessary* to adopt the new partition. Indeed, if the new partition $\mathcal{D}$ is not (strongly) feasible, then Proposition 3.1 already ensures that the current partition $\mathcal{C}$ is.

5.2. Closed-loop stability for time-varying coalitions

Our final development is to consider the impact of changing the system partition on closed-loop stability. We recall the following result, concerning the multiple Lyapunov-like function approach to stability for switched systems.

Lemma 5.1 (Zhang and Braatz (2013, Lemma 1))**.** *Consider the switched system $x(k+1) = f_{\sigma(k)}\left(x(k)\right)$, where $f_\sigma(\cdot)$ is globally Lipschitz continuous with $f_\sigma(0) = 0$, and $\sigma\colon \mathbb{I}_{\geq 0} \to \Sigma$ is the switching signal that takes values in a finite set Σ. If there exists a family of continuous positive-definite functions $V_m : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$, $\sigma(k) = m \in \Sigma$, satisfying for all $k \in \mathbb{I}_{\geq 0}$ and all $m \in \Sigma$*

$$\alpha_1\left(|x(k)|\right) \leq V_m\left(x(k)\right) \leq \alpha_2\left(|x(k)|\right)$$
$$V_m\left(x(k)\right) \leq \mu V_{\sigma(0)}\left(x(0)\right)$$

for two $\mathcal{K}$-functions $\alpha_1(\cdot)$ and $\alpha_2(\cdot)$ and a positive scalar $\mu > 0$, then the origin is locally stable.

Establishing closed-loop stability of the time-varying coalitional system then amounts to showing that the value functions of the coalitional MPC controllers satisfy the conditions of this lemma. In the following result, $\kappa_{\mathcal{C}}(\cdot)$ refers to the collection of control laws when the system is partitioned into $\mathcal{C} = \{1, \dots, C\} \in \Pi_{\mathcal{M}}$; more precisely, $\kappa_{\mathcal{C}}(x) = \left(\kappa_c(x_c)\right)_{c\in\mathcal{C}}$, where $\kappa_c(x_c) = \bar{\kappa}_c(\bar{x}_c) + \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}_c) + \tilde{\kappa}_c(\hat{e}_c)$ is the three-term control policy for coalition c defined in (4). The *switching signal* $\sigma\colon \mathbb{I}_{\geq 0} \mapsto \Pi_{\mathcal{M}}$ is implicitly defined by the partition selection algorithm; at time k, the algorithm has the system partitioned into a set of coalitions $\sigma(k) = \mathcal{C}(k) \in \Pi_{\mathcal{M}}$.

Theorem 5.2 (Local stability)**.** *Suppose that Assumptions 1–6 hold, that $x(0) \in \bar{\mathcal{X}}^N_{\mathcal{C}(0)}$ and that the switching signal $\sigma(\cdot)$ is such that, at each moment of switching, strong feasibility of the next partition is attained. Then the origin is locally stable for $x^+ = Ax + B\kappa_\sigma(x)$.*

Finally, the stronger result of asymptotic stability is obtained under the further assumption that the system does not change its partition after a finite time $k_f > 0$

Assumption 8. *The signal $\sigma(k) = \bar{\mathcal{C}} \in \Pi_{\mathcal{M}}$ for $k \geq k_f$ and some finite $k_f \geq 0$ such that $x(k_f) \in \bar{\mathcal{X}}^N_{\bar{\mathcal{C}}}$.*

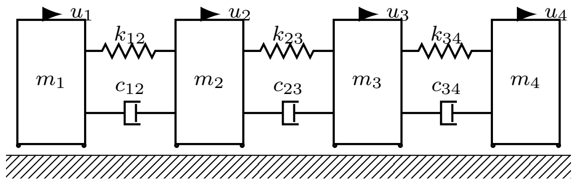

Figure 2: Coupled mass–spring–damper system

Theorem 5.3 (Asymptotic stability)**.** *Under the same hypotheses as Theorem 5.2, plus Assumptions 7 and 8, the origin is asymptotically stable for $x^+ = Ax + B\kappa_\sigma(x)$. The region of attraction is $\bar{\mathcal{X}}^N_{\mathcal{C}(0)}$.*

Remark 2. *The further characterization of dwell times, or bounds on these, is severely complicated by the two-layer MPC framework employed in our approach. The characterization of minimum dwell times has been investigated in connection with a different, but related, problem in Hernández Vicente and Trodden (2019), which considered a switching system with each model controlled by tube-based MPC.*

6. Illustrative examples

We illustrate and explore the results via an example system of a planar chain of four coupled mass–spring–dampers, shown in Figure 2. Each mass corresponds to a subsystem, and the state of mass i comprises its position (relative to some datum) and velocity, $x_i = (r_i, v_i)$. The continuous-time dynamics are

$$\begin{bmatrix}\dot{r}_i\\ \dot{v}_i\end{bmatrix} = \underbrace{\begin{bmatrix}0 & 1\\ -\frac{1}{m_i}\sum_{j\in\mathcal{M}_i} k_{ij} & -\frac{1}{m_i}\sum_{j\in\mathcal{M}_i} c_{ij}\end{bmatrix}}_{A_{ii}}\begin{bmatrix}r_i\\ v_i\end{bmatrix} + \begin{bmatrix}0\\ 100\end{bmatrix} u_i + w_i$$

Table 1: Closed-loop performance costs and practical costs—where $\sigma_{ij} = 1$ for all i, j—for different partitions of the system.

	Partition	$V^\infty(x, \mathbf{u})$	$(1/2)\sum_i J_i^{\text{power}}$
$\mathcal{C}^1$	$\{\{1,2,3,4\}\}$	73.1327	1.2
$\mathcal{C}^2$	$\{\{1,2,3\},\{4\}\}$	73.6061	0.6
$\mathcal{C}^3$	$\{\{1,2,4\},\{3\}\}$	73.8282	0.6
$\mathcal{C}^4$	$\{\{1,2\},\{4,3\}\}$	73.8376	0.4
$\mathcal{C}^5$	$\{\{1,2\},\{4\},\{3\}\}$	73.8270	0.2
$\mathcal{C}^6$	$\{1,3,4\},\{2\}$	73.3287	0.6
$\mathcal{C}^7$	$\{\{1,3\}\{2,4\}\}$	—	0.4
$\mathcal{C}^8$	$\{\{1,3\},\{2\},\{4\}\}$	—	0.2
$\mathcal{C}^9$	$\{\{1,4\},\{2,3\}\}$	73.2006	0.4
$\mathcal{C}^{10}$	$\{\{1\},\{2,3,4\}\}$	73.2008	0.6
$\mathcal{C}^{11}$	$\{\{1\},\{2,3\},\{4\}\}$	73.1890	0.2
$\mathcal{C}^{12}$	$\{\{1,4\},\{2\},\{3\}\}$	73.2027	0.2
$\mathcal{C}^{13}$	$\{\{1\},\{2,4\},\{3\}\}$	73.1910	0.2
$\mathcal{C}^{14}$	$\{\{1\},\{2\},\{3,4\}\}$	73.2016	0.2
$\mathcal{C}^{15}$	$\{\{1\},\{2\},\{3\},\{4\}\}$	73.1910	0.0

$$w_i = \sum_{j\in\mathcal{M}_i} \underbrace{\begin{bmatrix} 0 & 0 \\ \frac{1}{m_i}\sum_{j\in\mathcal{M}_i} k_{ij} & \frac{1}{m_i}\sum_{j\in\mathcal{M}_i} c_{ij} \end{bmatrix}}_{A_{ij}} \begin{bmatrix} r_j \\ v_j \end{bmatrix}.$$

Here, u_i is the control input (acceleration) to mass i. The disturbance w_i arises via the coupling between masses: mass 1 ($m_1 = 3\,\mathrm{kg}$) is coupled to mass 2 ($m_2 = 2\,\mathrm{kg}$) via a spring (stiffness $k_{12} = 0.5\,\mathrm{N\,m^{-1}}$) and damper ($c_{12} = 0.2\,\mathrm{N\,m^{-1}\,s^{-1}}$). Likewise, mass 3 ($m_3 = 3\,\mathrm{kg}$) is coupled to mass 4 ($m_4 = 6\,\mathrm{kg}$) via $k_{34} = 1\,\mathrm{N\,m^{-1}}$ and $c_{34} = 0.3\,\mathrm{N\,m^{-1}\,s^{-1}}$. Moreover, masses 2 and 3 are also coupled, via $k_{23} = 0.75\,\mathrm{N\,m^{-1}}$ and $c_{23} = 0.25\,\mathrm{N\,m^{-1}\,s^{-1}}$.

The system is subject to constraints on the states, $\mathbb{X}_i = \{(r_i, v_i) : -2 \le r_i \le 2,\ -8 \le v_i \le 8\}$, and control inputs, $\mathbb{U}_i = \{u_i : -4 \le u_i \le 4\}$. The initial conditions are $x_1(0) = (1.8, 0)$, $x_2(0) = (-0.5, 0)$, $x_3(0) = (1, 0)$, and $x_4(0) = (-1, 0)$. The parameters for the MPC controllers are a horizon $N = H - 1 = 25$ and matrices $Q_i = I$, $R_i = 1$.

6.1. Performance and practical costs with time-invariant partitions

The initial exploration is the performance of the system under different, albeit static, partitions. For $\mathcal{M} = \{1, 2, 3, 4\}$, there are 15 possible partitions in the partition set $\Pi_\mathcal{M} = \{\mathcal{C}^1, \mathcal{C}^2, \ldots, \mathcal{C}^{14}, \mathcal{C}^{15}\}$, where $\mathcal{C}^1$ corresponds to the centralized partition $\mathcal{C}^{\text{cen}}$ and $\mathcal{C}^{15}$ corresponds to the decentralized partition $\mathcal{C}^{\text{dec}}$. For each partition, the primary and secondary MPC controllers were designed for each coalition using the procedure in Appendix A. The design succeeded for all partitions except $\mathcal{C}^7 = \{\{1, 3\}, \{2, 4\}\}$ and $\mathcal{C}^8 = \{\{1, 3\}, \{2\}, \{4\}\}$. Inspection revealed that $\mathcal{C}^7$ and $\mathcal{C}^8$ place subsystems with no physical coupling into coalitions; consequently, the interactions between coalitions are too strong in order for the design procedure to succeed, and therefore there do not exist suitable scaling factors to build the optimal control problems in these cases.

The closed-loop state trajectories for mass 1 are shown in Figure 3. The whole-system closed-loop costs are in Table 1. The centralized partition achieves the lowest regulation cost but has the highest practical cost at each step. The opposite outcome is observed for the decentralized partition. The performance of $\mathcal{C}^3$, $\mathcal{C}^4$, and $\mathcal{C}^5$ exhibit the worst regulation performance, as a consequence of these coalitions leading to the highest constraint tightening margins, but with lower (higher) practical cost than for the centralized (decentralized). On the other hand, some surprising results emerge: performance does not necessarily always deteriorate with refinement (compare $\mathcal{C}^{11}$ with $\mathcal{C}^2$), and neither does it necessarily improve with coarsening.

Table 2: Cost function values for different algorithms.

	$V^\infty(x, \mathbf{u})$	$J^\infty(x, \mathcal{C})$	$J(\mathcal{C}(k))$
$\mathcal{C}^{\text{cen}}$	176.86	120.00	1.20
$\mathcal{C}^{\text{dec}}$	178.19	27.30	0.00
$\mathcal{C}(k)$	177.67	27.65	0.15

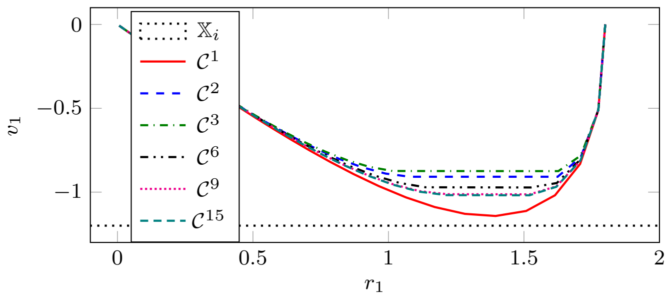

Figure 3: Closed-loop state trajectories for different system partitions. $\mathcal{C}^4$ and $\mathcal{C}^5$ gave responses visually indistinguishable from that of $\mathcal{C}^3$; likewise, $\mathcal{C}^{10}$–$\mathcal{C}^{14}$ gave responses in between those of $\mathcal{C}^9$ and $\mathcal{C}^{15}$.

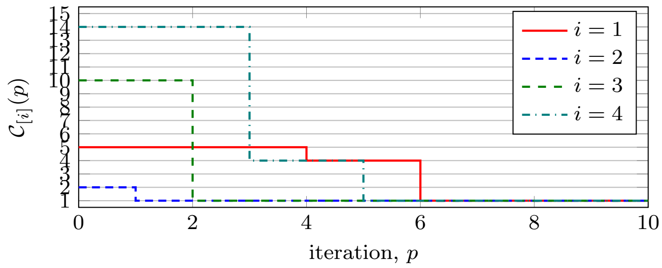

Figure 4: Consensus algorithm iterations at the initial time step $k = 0$. At each iteration one subsystem can improve its choice of partition following the proposed best reply map.

Overall, the results suggest a complicated relationship between partition and closed-loop performance that demands further research.

6.2. Time-varying partitions

To increase coupling strength, we change the initial conditions to $x_1(0) = (1.0, -7)$, $x_2(0) = (-0.51, 4)$, $x_3(0) = (-1.71, 0)$, and $x_4(0) = (1.8, -4)$. The aim is to investigate the performance of the overall scheme, including the consensus-based partition selection algorithm. To this end, the partition opinions for each subsystem are initialized, at iteration $p = 0$ and time $k = 0$, as $\mathcal{C}_{[1]}(0) = \mathcal{C}^2$, $\mathcal{C}_{[2]}(0) = \mathcal{C}^2$, $\mathcal{C}_{[3]}(0) = \mathcal{C}^{10}$, and $\mathcal{C}_{[4]}(0) = \mathcal{C}^{14}$.

The execution of the partition selection algorithm at time $k = 0$ is illustrated in Figure 4; the outcome, after six iterations, is a consensus among subsystems to adopt the centralized partition, $\mathcal{C}^1$, for the initial time $k = 0$. Figure 5 shows the system partition selected and employed at each subsequent time step during the simulation. Following the use of initial centralized partition, the subsystems agree on the partition $\mathcal{C}^{10}$—which groups subsystem 2, 3, and 4 together, and subsystem 1 separately—between times $k = 1$ and $k = 7$. At time $k = 8$, the partition changes to $\mathcal{C}^{14}$, which groups only subsystems 3 and 4. At $k = 19$ the subsystems agree to disband into the decentralized partition $\mathcal{C}^{15}$, as the effect of the dynamic coupling weakens, *i.e.*, $w_{ij}(x_i, x_j) \to 0$.

Finally, we compare the performance of the time-varying scheme against that of a fixed centralized partition and a

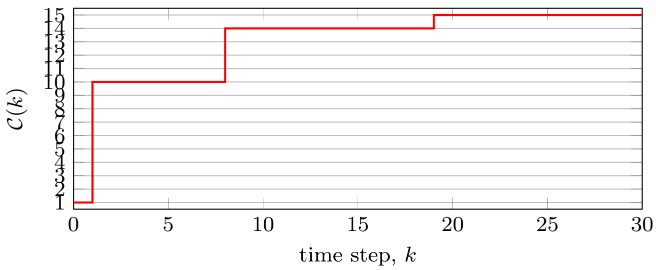

Figure 5: Partition selected at each time step.

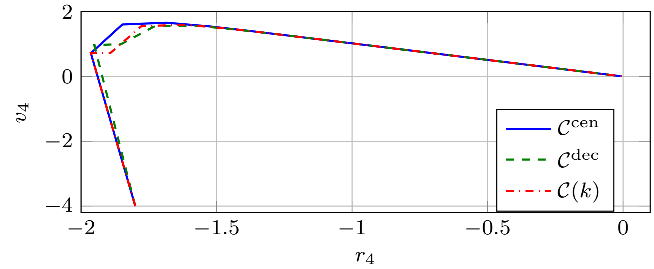

Figure 6: State trajectory for subsystem 4 under control by the centralized partition, decentralized partition and the proposed time-varying scheme.

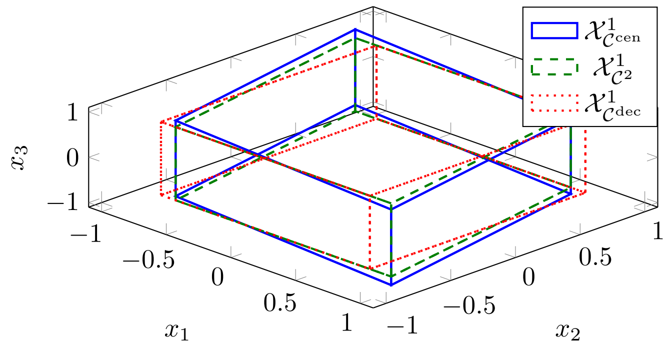

Figure 7: Feasibility regions for the partitions $\mathcal{C}^{\mathrm{cen}}$, $\mathcal{C}^2$ and $\mathcal{C}^{\mathrm{dec}}$.

fixed decentralized partition. Figure 6 shows the closed-loop state trajectory for subsystem 4, for which differences can be seen during the transient. Table 2 summarizes the closed-loop costs. The first column gives the closed-loop performance costs, the second column shows the summation (over the simulation) of the consensus optimization cost *i.e.*, $J^\infty(x, \mathcal{C}) = (1/2)\sum_k \sum_{i\in\mathcal{M}} J_i(\mathcal{C}_{[i](k)}; \mathcal{C}_{[-i]}(k), x(k))$, and the third column gives the value of $(1/2)\sum_{i\in\mathcal{M}} J_i^{\mathrm{power}}$, where all $\sigma_{ij} = 1$, averaged across the simulation. We observe that the proposed scheme balances the trade-off between closed-loop performance and practical cost.

6.3. Feasibility regions under different partitions

A smaller-scale example is now considered to illustrate the feasibility results reported in Section 5.1. Let

$$x_i^+ = 0.6x_i + u_i + w_i \text{ for } i = 1, 2, 3, \tag{9}$$

where $w_1 = 0.1x_2$, $w_2 = 0$ and $w_3 = 0.1x_1$. The subsystems $i = 1, 2, 3$ are subject to constraints $|x_i| \leq 2$ and $|u_i| \leq 0.5$.

There are five possible partitions of $\mathcal{M}$. We consider the chain $\mathbb{C} = \{\mathcal{C}^{\mathrm{cen}}, \mathcal{C}^2, \mathcal{C}^{\mathrm{dec}}\}$, where $\mathcal{C}^2 = \{\{1,2\},\{3\}\}$. The one-step feasibility regions for the *true* coalition dynamics, *i.e.*, the product set of $\mathcal{X}_c^1 \triangleq \bar{\mathcal{X}}_c^1 \oplus \mathcal{ER}_c^1$ over $c \in \mathcal{C}$, for each $\mathcal{C} \in \mathbb{C}$, is displayed in Figure 7.

Illustrating Proposition 5.1 and the subsequent discussion, there is no nesting of the sets $\mathcal{X}^1_{\mathcal{C}^{\mathrm{cen}}}$, $\mathcal{X}^1_{\mathcal{C}^2}$, and $\mathcal{X}^1_{\mathcal{C}^{\mathrm{dec}}}$. Indeed, there exist states $x \in \mathcal{X}^1_{\mathcal{C}^2}$ such that $x \notin \mathcal{X}^1_{\mathcal{C}^{\mathrm{cen}}}$ and/or $x \notin \mathcal{X}^1_{\mathcal{C}^{\mathrm{dec}}}$; for example, the state $x_1 = 0.9722$, $x_2 = -0.8333$, $x_3 = 0.8074$. Applying the proposed algorithm from this state, we find that the system may only be feasibly controlled if partitioned as $\mathcal{C}^2$. At later times, once the state has been steered into $\bar{\mathcal{X}}^1_{\mathcal{C}^{\mathrm{dec}}}$, the algorithm switches to the decentralized partition.

7. Conclusions

A coalitional MPC approach for dynamically coupled subsystems was presented. The approach controls a constrained system of subsystems in a way that balances the performance degradation of decentralized control with the practical cost of centralized control, by allowing subsystem controllers to form coalitions over time, thus reconfiguring the partition of the system dynamics. For regulation, the approach employs a robust distributed model predictive control with implicit, rather than explicit, reliance on robust invariant sets. Re-partitioning of the system into different coalitions is achieved via a consensus-based algorithm. Recursive feasibility and stability of the overall time-varying coalitional control scheme are established under dwell-time assumptions. An interesting consequence of the developments is the finding that a coalitional controller may reach states that are otherwise infeasible for a centralized controller.

References

Baldivieso Monasterios, P.R., Hernandez, B., Trodden, P.A., 2017. Nested distributed MPC. IFAC-PapersOnLine 50, 11822–11828. doi:10.1016/j.ifacol.2017.08.1994. 20th IFAC World Congress.

Barreiro-Gomez, J., Ocampo-Martinez, C., Quijano, N., 2019. Time-varying partitioning for predictive control design: Density-games approach. J. Process Control 75, 1–14.

Barreiro-Gomez, J., Ocampo-Martinez, C., Quijano, N., Maestre, J.M., 2017. Non-centralized control for flow-based distribution networks: A game-theoretical insight. J. Franklin Inst. 354, 5771–5796.

Christofides, P.D., Scattolini, R., Muñoz del la Peña, D., Liu, J., 2013. Distributed model predictive control: A tutorial review and future research directions. Computers & Chemical Engineering 51, 21–41.

Farina, M., Scattolini, R., 2012. Distributed predictive control: A non-cooperative algorithm with neighbor-to-neighbor communication for linear systems. Automatica 48, 1088–1096. doi:10.1016/j.automatica.2012.03.020.

Fele, F., Debada, E., Maestre, J.M., Camacho, E.F., 2018. Coalitional control for self-organizing agents. IEEE Transactions on Automatic Control 63, 2883–2897. doi:10.1109/TAC.2018.2792301.

Fele, F., Maestre, J.M., Camacho, E.F., 2017. Coalitional control: Cooperative game theory and control. IEEE Control Systems Magazine 37, 53–69. doi:10.1109/MCS.2016.2621465.

Hernandez Vicente, B., Trodden, P., 2019. Switching tube-based MPC: Characterization of minimum dwell-time for feasible and robustly stable switching. IEEE Transactions on Automatic Control 64, 4345–4352.

Hernández Vicente, B.A., Trodden, P.A., 2019. Switching tube-based MPC: Characterization of minimum dwell-time for feasible and robustly stable switching. IEEE Transactions on Automatic Control 64, 4345–4352. doi:10.1109/TAC.2019.2897551.

Maestre, J.M., Muñoz de la Peña, D., Jiménez Losada, A., Algaba, E., Camacho, E.F., 2014. A coalitional control scheme with applications to cooperative game theory. Optimal Control Applications and Methods 35, 592–608.

Maestre, J.M., Negenborn, R.R. (Eds.), 2014. Distributed Model Predictive Control Made Easy. Springer.

Mayne, D.Q., Seron, M.M., Raković, S.V., 2005. Robust model predictive control of constrained linear systems with bounded disturbances. Automatica 41, 219–224. doi:10.1016/j.automatica.2004.08.019.

Monderer, D., Shapley, L.S., 1996. Potential games. Games and Economic Behavior 14, 124–143.

Motee, N., Sayyar-Rodsari, B., 2003. Optimal partitioning in distributed model predictive control, in: Proceedings of the American Control Conference, pp. 5300–5305.

Müller, M.A., Martius, P., Allgöwer, F., 2012. Model predictive control of switched nonlinear systems under average dwell-time. J. Process Control 22, 1702–1710. doi:10.1016/j.jprocont.2012.07.004.

Muros, F., Algaba, E., Maestre, J., Camacho, E., 2017a. Harsanyi Power Solutions in Coalitional Control Systems. IEEE Trans. Automat. Contr. 62, 3369–3381. doi:10.1109/TAC.2017.2651642.

Muros, F., Algaba, E., Maestre, J., Camacho, E., 2017b. The Banzhaf value as a design tool in coalitional control. Syst. Control Lett. 104, 21–30. doi:10.1016/j.sysconle.2017.03.007.

Muros, F., Maestre, J., Algaba, E., Alamo, T., Camacho, E., 2014. An Iterative Design Method for Coalitional Control Networks with Constraints on the Shapley Value. IFAC Proc. Vol. 47, 1188–1193. doi:10.3182/20140824-6-ZA-1003.02616.

Muros, F., Maestre, J., Algaba, E., Alamo, T., Camacho, E., 2017c. Networked control design for coalitional schemes using game-theoretic methods. Automatica 78, 320–332. doi:10.1016/j.automatica.2016.12.010.

Muros, F.J., Maestre, J.M., Ocampo-Martinez, C., Algaba, E., Camacho, E.F., 2018. A Game Theoretical Randomized Method for Large-Scale Systems Partitioning. IEEE Access 6, 42245–42263.

Pourkargar, D.B., Almansoori, A., Daoutidis, P., 2017. Impact of Decomposition on Distributed Model Predictive Control: A Process Network Case Study. Ind. Eng. Chem. Res. 56, 9606–9616.

Raković, S.V., Kerrigan, E.C., Mayne, D.Q., Kouramas, K.I., 2007. Optimized robust control invariance for linear discrete-time systems: Theoretical foundations. Automatica 43, 831–841. doi:10.1016/j.automatica.2006.11.006.

Raković, S.V., Mayne, D.Q., 2005. Set robust control invariance for linear discrete time systems, in: Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference.

Rawlings, J.B., Mayne, D.Q., 2009. Model Predictive Control: Theory and Design. Nob Hill Publishing.

Riverso, S., Ferrari-Trecate, G., 2012. Tube-based distributed control of linear constrained systems. Automatica 48, 2860–2865. doi:10.1016/j.automatica.2012.08.024.

Scattolini, R., 2009. Architectures for distributed and hierarchical Model Predictive Control – A review. J. Process Control 19, 723–731.

Trodden, P.A., Maestre, J.M., 2017. Distributed predictive control with minimization of mutual disturbances. Automatica 77, 31–43. doi:10.1016/j.automatica.2016.11.023.

Venkat, A.N., Hiskens, I.A., Rawlings, J.B., Wright, S.J., 2008. Distributed MPC strategies with application to power system automatic generation control. IEEE Trans. Control Syst. Technol. 16, 1192–1206.

Šiljak, D.D., 1991. Decentralized control of complex systems. Academic Press, Inc., Boston.

Zhang, L., Braatz, R.D., 2013. On switched MPC of a class of switched linear systems with modal dwell time, in: 52nd IEEE Conf. Decis. Control, IEEE. pp. 91–96. doi:10.1109/CDC.2013.6759864.

Zhang, L., Zhuang, S., Braatz, R.D., 2016. Switched model predictive control of switched linear systems: Feasibility, stability and robustness. Automatica 67, 8–21. doi:10.1016/j.automatica.2016.01.010.

Zheng, Y., Wei, Y., Li, S., 2018. Coupling Degree Clustering-Based Distributed Model Predictive Control Network Design. IEEE Trans. Autom. Sci. Eng. 15, 1749–1758.

Appendix A. Controller design

The design of the primary and secondary MPC controllers for each coalition, via the selection of the scaling factors that restrict the constraints in the optimal control problems, is based on the theory of optimized robust control invariance (Raković et al., 2007); therefore, we begin with an overview of the main concepts.

Appendix A.1. Optimized Robust Control Invariance

The optimized RCI approach (Raković et al., 2007) proposed a novel characterization of an RCI set for a system $x^+ = Ax + Bu + w$ and constraint set $(\mathbb{X}, \mathbb{U}, \mathbb{W})$ as

$$\mathcal{R}_h(\mathbf{M}_h) = \bigoplus_{l=0}^{h-1} D_l(\mathbf{M}_h)\mathbb{W} \text{ with } \mu(\mathcal{R}_h(\mathbf{M}_h)) = \bigoplus_{l=0}^{h-1} M_l\mathbb{W}.$$

The set $\mu(\mathcal{R}_h(\mathbf{M}_h))$ is the set of invariance-inducing control actions, defined as $\mu(\mathcal{R}_h) \triangleq \{\mu(x) : x \in \mathcal{R}_h\} = \{u \in \mathbb{U} : x^+ \in \mathcal{R}_h, \forall w \in \mathbb{W}\}$. The matrices $D_l(\mathbf{M}_h), l = 0 \ldots h$ are

$$D_l(\mathbf{M}_h) = \begin{cases} I & l = 0 \\ A^l + \sum_{j=0}^{l-1} A^{l-1-j} B M_j & l \geq 1, \end{cases}$$

with $M_j \in \mathbb{R}^{m\times n}$, and $\mathbf{M}_h \triangleq (M_0, M_1, \ldots, M_{h-1})$, such that $D_h(\mathbf{M}_h) = 0$; the latter is ensured by setting h greater than or equal to the controllability index of (A, B). The set of matrices that satisfy these conditions is given by $\mathbb{M}_h \triangleq \{\mathbf{M}_h : D_h(\mathbf{M}_h) = 0\}$. Constraint satisfaction is guaranteed if $\mathcal{R}_h(\mathbf{M}_h) \subseteq \eta\mathbb{X}$ and $\mu(\mathcal{R}_h(\mathbf{M}_h)) \subseteq \theta\mathbb{U}$, with $(\eta, \theta) \in [0, 1]\times[0, 1]$. The linear programming (LP) problem to compute these matrices is

$$\min\{\delta : \gamma \in \Gamma\}, \tag{A.1}$$

where $\gamma = (\mathbf{M}_h, \eta, \theta, \delta)$, and the set $\Gamma = \{\gamma : \mathbf{M}_h \in \mathbb{M}_h, \mathcal{R}_h(\mathbf{M}_h) \subseteq \eta\mathbb{X}, \mu(\mathcal{R}_h(\mathbf{M}_h)) \subseteq \theta\mathbb{U}, (\eta, \theta) \in [0, 1] \times [0, 1],\ q_\eta\eta + q_\theta\theta \leq \delta\}$; q_η and q_θ are weights to express a preference for the relative contraction of state and input constraint sets. Feasibility of this problem is linked to the controllability of the system and the relative sizes of constraint and disturbance sets: if (A, B) is controllable, then $\mathbf{M}_h$ exists if h is as least as large as the controllability index of the system; consequently, an $\mathcal{R}_h(\mathbf{M}_h)$ exists satisfying the invariance properties. However, Problem (A.1) is feasible if and only if $\mathcal{R}_h(\mathbf{M}_h)$ is constraint admissible (Raković et al., 2007). The resulting RCI set is minimal with respect to the criteria δ, and is the only set satisfying the inclusion constraints for state and control.

Appendix A.2. Design algorithm

The RCI LP problem is useful in the current context because solving it provides an invariance-inducing robust control law—a suitable candidate for the third term in the overall control law—plus scaling constants that outer-bound (with respect to the state and input constraint sets) the size of the RCI set and its corresponding set of control actions. Therefore, we employ the RCI LP problem as the key ingredient in the following design procedure, for each coalition $c \in \mathcal{C}$. The design starts with determining an RCI set for the overall error, e_c, and the overall disturbance set $\mathbb{W}_c$, because the latter is known. The real aim is to determine an RCI control law for the unplanned error, $\hat{e}_c$, and unplanned disturbance set $\hat{\mathbb{W}}_c$; however, the latter is not known until the scaling constants α_c^x for each coalition have been determined.

1. The problem (A.1) associated with the dynamics $e_c^+ = A_{cc}e_c + B_c f_c + w_c$ and known constraint set $(\mathbb{X}_c, \mathbb{U}_c, \mathbb{W}_c)$ is solved to yield $\gamma_{c,h} = (\mathbf{M}_{c,h}, \eta_c, \theta_c, \delta_c)$, where η_c and θ_c are scalings of $\mathbb{X}_c$ and $\mathbb{U}_c$ such that $\mathcal{R}_{c,h} \subset \eta_c\mathbb{X}_c$ and $\mu_c(\mathcal{R}_{c,h}) \subset \theta_c\mathbb{U}_c$ respectively.

2. Given that, under the RCI control law $f_c = \mu_c(e_c)$, $e_c \in \mathcal{R}_{c,h} \subset \eta_c\mathbb{X}_c$ and $f_c \in \mu(\mathcal{R}_{c,h}) \subset \theta_c\mathbb{U}_c$, we select
$$\alpha_c^x = 1 - \eta_c; \quad \alpha_c^u = 1 - \theta_c,$$
for the scaling factors in the main MPC problem. Then $x_c = \bar{x}_c + e_c \in \alpha_c^x\mathbb{X}_c \oplus \eta_c\mathbb{X}_c = \mathbb{X}_c$, with a similar expression for u_c. The scaling factor α_c^x is transmitted to neighbouring coalitions.

3. Given α_d^x for $d \in \mathcal{M}_c$, the set $\hat{\mathbb{W}}_c = \bigoplus_{d\in\mathcal{M}_c}(1 - \alpha_c^x)A_{cd}\mathbb{X}_d$ is computed[4] and the RCI problem (A.1), now associated with $\hat{e}_c = A_{cc}\hat{e}_c + B_c\hat{f}_c + \hat{w}_c$ and $(\mathbb{X}_c, \mathbb{U}_c, \hat{\mathbb{W}}_c)$, is re-solved for $\tilde{\gamma}_{(c,h)} = (\mathbf{M}_{c,h}, \tilde{\eta}_c, \tilde{\theta}_c, \tilde{\delta}_c)$, yielding the scaling factors
$$\xi_c^x = \tilde{\eta}_c; \quad \xi_c^u = \tilde{\theta}_c.$$
These scaling factors are such that $\mathcal{R}_{c,h} \subset \xi_c^x\mathbb{X}_c$ and $\mu_c(\mathcal{R}_{c,h}) \subset \xi_c^u\mathbb{U}_c$; that is, the regions of the constraint sets that the third-term robust control law occupies in response to the unplanned error and unplanned disturbance.

4. The selection of the constants β_c^x and β_c^u for the secondary MPC problem is made as
$$\beta_c^x = 1 - \alpha_c^x - \xi_c^x; \quad \beta_c^u = 1 - \alpha_c^u - \xi_c^u.$$
Then $x_c = \bar{x}_c + \bar{e}_c + \hat{e}_c \in \alpha_c^x\mathbb{X}_c \oplus \beta_c^x\mathbb{X}_c \oplus \xi_c^x\mathbb{X}_c = \mathbb{X}_c$, as required, with a similar expression for u_c.

5. The control law $\hat{f}_c = \tilde{\kappa}_c(\hat{e}_c) = \mu_c(\hat{e}_c)$ is computed from the matrices $\mathbf{M}_{c,h}$, using the minimal selection map procedure described in (Raković and Mayne, 2005).

[4] Note that an outer-approximation to $\hat{\mathbb{W}}_c$ is easily computed as $t_c\mathbb{W}_c$ where $t_c = \left(\max_{d\in\mathcal{M}_c}(1 - \alpha_d^x)\right)$.

Appendix B. Proofs

Appendix B.1. Proof of Proposition 3.1

For part (i), because the nominal model is linear, $\alpha_c^x\mathbb{X}_c$ and $\alpha_c^u\mathbb{U}_c$ are PC-sets, and the terminal constraint is control invariant, the set $\bar{\mathcal{X}}_c^N$ is compact, contains the origin and satisfies $\bar{\mathcal{X}}_c^N \supseteq \bar{\mathcal{X}}_c^{N-1} \supseteq \cdots \supseteq \bar{\mathcal{X}}_c^0 = \{0\}$. Moreover, $\bar{\mathcal{X}}_c^N$ is positively invariant for $\bar{x}_c^+ = A_{cc}\bar{x}_c + B_c\bar{\kappa}_c(\bar{x}_c)$, which is sufficient to prove the claim. (For a detailed proof, see (Rawlings and Mayne, 2009, Proposition 2.11).) The same arguments applied to $\bar{\mathcal{E}}_c^H(\bar{\mathbf{w}})$ establish part (ii).

For (iii), suppose that at time k, $\bar{x}_c \in \bar{\mathcal{X}}_c^N$, $\bar{e}_c \in \bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c)$ with $\bar{\mathbf{w}}_c \in \bar{\mathcal{W}}_c^N$, and $\hat{e}_c \in \mathcal{R}_c$. Then $x_c \in \bar{\mathcal{X}}_c^N \oplus \bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c) \oplus \mathcal{R}_c \subseteq \alpha_c^x\mathbb{X}_c \oplus \beta_c^x\mathbb{X}_c \oplus \xi_c^x\mathbb{X}_c = (\alpha_c^x + \beta_c^x + \xi_c^x)\mathbb{X}_c \subseteq \mathbb{X}_c$. The applied control is $u_c = \bar{u}_c^0(0;\bar{x}_c) + \bar{f}_c^0(0;\bar{e}_c,\bar{\mathbf{w}}_c) + \mu_c(\hat{e}_c) \in \alpha_c^u\mathbb{U}_c \oplus \beta_c^u\mathbb{U} \oplus \xi_c^u\mathbb{U}_c \subseteq \mathbb{U}_c$. Then, because of parts (i) and (ii), $x_c^+ = A_{cc}x_c + B_cu_c + w_c \in \bar{\mathcal{X}}_c^N \oplus \bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c^+) \oplus \mathcal{R}_c$. To complete the proof, however, we must consider the possibility that the disturbance sequence at the successor state is $\bar{\mathbf{w}}_c^0 \neq \bar{\mathbf{w}}_c^+$: in that case, if $\hat{\mathbb{P}}_c(\bar{e}_c^+;\bar{\mathbf{w}}_c^0)$ is feasible then $x_c^+ \in \bar{\mathcal{X}}_c^N \oplus \bar{\mathcal{E}}_c^H(\bar{\mathbf{w}}_c^0) \oplus \mathcal{R}_c$, which is still within $\mathbb{X}_c$ by construction, and $u_c = \bar{u}_c^0(0;\bar{x}_c^+) + \bar{f}_c^0(0;\bar{e}_c^+,\bar{\mathbf{w}}_c^0) + \mu_c(\hat{e}_c^+) \subseteq \mathbb{U}_c$. If $\hat{\mathbb{P}}_c(\bar{e}_c^+;\bar{\mathbf{w}}_c^0)$ is not feasible, then $\hat{\mathbb{P}}_c(\bar{e}_c^+;\bar{\mathbf{w}}_c^+)$ *is* feasible (by the tail), and $u_c = \bar{u}_c^0(0;\bar{x}_c^+) + \bar{f}_c^0(0;\bar{e}_c^+,\bar{\mathbf{w}}_c^+) + \mu_c(\hat{e}_c^+) \subseteq \mathbb{U}_c$. This establishes recursive feasibility of the algorithm.

Finally, if, at time 0, $\bar{x}_c = x_c \in \bar{\mathcal{X}}_c^N$ then $\bar{e}_c = 0$. Moreover, if $\bar{\mathbf{w}}_c = 0$, then—trivially—$\bar{e}_c \in \bar{\mathcal{E}}_c^H(0)$ and both the primary and secondary problems are feasible. By recursion, feasibility is retained at the next step, and the proof is complete. □

Appendix B.2. Proof of Theorem 3.1

By Proposition 3.1, $\bar{x}_c \in \bar{\mathcal{X}}_c^N$ implies $\bar{x}_c^+ = A_{cc}\bar{x}_c + B_c\bar{\kappa}_c(\bar{x}_c) \in \bar{\mathcal{X}}_c^N$, with $\bar{V}_c^0(\bar{x}_c^+) \leq \bar{V}_c^0(\bar{x}_c) - \ell_c(\bar{x}_c, \bar{\kappa}_c(\bar{x}_c))$ where $\ell_c(x_c, u_c) \triangleq x_c^\top Q_c x_c + u_c^\top R_c u_c$. By Assumption 4, there exists a constant $a_c > 0$ such that
$$\bar{V}_c^0(\bar{x}_c) \geq \ell_c(x_c, \bar{\kappa}_c(\bar{x}_c)) \geq a_c|\bar{x}_c|^2$$
for all $\bar{x}_c \in \bar{\mathcal{X}}_c^N$ and, moreover, Assumption 1 ensures the existence of a constant $b_c > a_c > 0$ such that $\bar{V}_c^0(\bar{x}_c) \leq b_c|\bar{x}_c|^2$ over the same domain. Then $\bar{V}_c^0(\bar{x}_c^+) \leq \gamma_c\bar{V}_c^0(\bar{x}_c)$ where $\gamma_c \triangleq (1 - a_c/b_c) \in (0,1)$. If $\bar{x}_c(0) \in \bar{\mathcal{X}}_c^N$ then $\bar{V}_c^0(\bar{x}_c(k)) \leq \gamma_c^k\bar{V}_c^0(\bar{x}_c(0))$ and $\left|\bar{x}_c(k)\right| \leq d_c\delta_c^k\left|\bar{x}_c(0)\right|$ where $\delta_c \triangleq \sqrt{\gamma_c}$ and $d_c \triangleq \sqrt{b_c/a_c}$. This establishes exponential stability for the nominal system.

Now consider the true trajectory $\{x_c(k)\}_k$. We have $x_c(0) = \bar{x}_c(0) \in \bar{\mathcal{X}}_c^N$, so $|x_c(0)| = |\bar{x}_c(0)|$. Consider some $x_c = \bar{x}_c + e_c$; by Proposition 3.1, $e_c = \bar{e}_c + \hat{e}_c \in \bar{\mathcal{E}}_c^H\left(\bar{\mathbf{w}}_c\right) \oplus \hat{\mathcal{R}}_c$ implies $e_c^+ = \bar{e}_c^+ + \hat{e}_c^+ \in \bar{\mathcal{E}}_c^H\left(\bar{\mathbf{w}}_c^+\right) \oplus \hat{\mathcal{R}}_c$, with
$$V^H(\bar{e}_c^+, \bar{\mathbf{f}}_c^+; \bar{\mathbf{w}}_c^+) \leq \hat{V}_c^0(\bar{e}_c; \bar{\mathbf{w}}_c) - \ell_c(\bar{e}_c, \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}})).$$
At $\bar{e}_c^+$, the actual planned disturbance is $\tilde{\mathbf{w}}_c$, which may differ from $\bar{\mathbf{w}}_c^+$; this $\tilde{\mathbf{w}}_c$ is adopted if and only if $\bar{e}_c^+ \in \bar{\mathcal{E}}^H(\tilde{\mathbf{w}}_c)$ and $\hat{V}^0(\bar{e}_c^+; \tilde{\mathbf{w}}_c) \leq V^H(\bar{e}_c^+, \bar{\mathbf{f}}_c^+; \bar{\mathbf{w}}_c^+)$, implying that
$$\hat{V}^0(\bar{e}_c^+; \tilde{\mathbf{w}}_c) \leq \hat{V}_c^0(\bar{e}_c; \bar{\mathbf{w}}_c) - \ell_c(\bar{e}_c, \hat{\kappa}_c(\bar{e}_c; \bar{\mathbf{w}}))$$

Thus, the sequence $\{\hat{V}_c^0(\bar{e}_c(k);\bar{\mathbf{w}}_c(k))\}_{k\in\mathbb{I}_{\geq 0}}$ is bounded below by zero and monotonically decreasing, so the limit $\lim_{k\to\infty}\hat{V}_c^0(k)=0$, which in turn implies $\bar{e}_c(k)\to 0$. Finally, we consider the dynamics of $\hat{e}_c$; since $\hat{e}_c = e_c - \bar{e}_c = x_c - \bar{x}_c - \bar{e}_c$ then

$$\hat{e}_c^+ = A_{cc}\hat{e}_c + B_c\tilde{\kappa}_c(\hat{e}_c) + \sum_{d\in\mathcal{M}_c} A_{cd}\hat{e}_d.$$

Thus, $\hat{e}^+ = A\hat{e} + B\tilde{\kappa}(\hat{e})$, where $\tilde{\kappa}(\cdot)$ denotes the diagonal collection of $\tilde{\kappa}_c(\cdot)$, for which $\hat{e}(k)\to 0$ as $k\to\infty$, in view of Assumption 7. Finally, since $x_c = \bar{x}_c + \bar{e}_c + \hat{e}_c$, and all three terms decay asympotically to zero, then $x_c(k)\to 0$ as $k\to\infty$. □

Appendix B.3. Proof of Theorem 4.2

If $\Delta \geq 1$ then, for all $i\in\mathcal{M}$, $\mathbb{C}_\Delta(\mathcal{C}_{[i]})$ is a non-singleton finite set for all $\mathcal{C}_{[i]}\in\Pi_{\mathcal{M}}$; thus, for any profile, there always exists a feasible opinion for each player *different* to the one it is currently playing. Consider an profile $\mathcal{C} = (\mathcal{C}_{[1]},\ldots,\mathcal{C}_{[M]})$. If there does not exist any player i for which $J_i(\mathcal{C}_{[i]}^0,\mathcal{C}_{[-i]}) < J_i(\mathcal{C}_{[i]},\mathcal{C}_{[-i]})$, where $\mathcal{C}_{[i]}^0$ is the constrained best response defined by (7), then following the proposed algorithm, each player implements an unconstrained version of (7). In this setting, there are only two possible outcomes: the first is that no player may gain from this unconstrained optimisation, and so the game is by definition in a Nash equilibrium. The second is that there exists at least one $i\in\mathcal{M}$ for which $J_i(\mathcal{C}_{[i]}^0,\mathcal{C}_{[-i]}) < J_i(\mathcal{C}_{[i]},\mathcal{C}_{[-i]})$ and hence $\phi(\mathcal{C}_{[i]}^0,\mathcal{C}_{[-i]}) < \phi(\mathcal{C}_{[i]},\mathcal{C}_{[-i]})$. In this latter case, therefore, one player deviates from the constrained equilibrium, and algorithm moves to a new iteration of constrained best response. Applying this reasoning recursively, the overall algorithm follows an improvement path: at iteration p, either the game is in a Nash equilibrium, hence has terminated, or at least one player can lower the potential function further (by either constrained or unconstrained optimization). Thus, the path $\mathfrak{C}$ is a better response path. It must be finite because $\Pi_{\mathcal{M}}$ is a finite set. Thus, it terminates in a finite number of iterations to a Nash equilibrium, which is a minimum of the potential function. □

Appendix B.4. Proof of Theorem 4.3

The potential function satisfies, for any $\mathcal{C}_{[i]}\in\mathbb{C}_\Delta(\mathcal{C}_{[i]}^e)$

$$\begin{aligned}\Delta\phi &\triangleq \phi(\mathcal{C}_{[1]}^e,\ldots,\mathcal{C}_{[i]},\ldots,\mathcal{C}_{[M]}^e) - \phi(\mathcal{C}_{[1]}^e,\ldots,\mathcal{C}_{[M]}^e)\\ &= \sum_{j\in\mathcal{M}_i}\Big(w_{ij}(x_i,x_j)\big|\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[j]}^e)\big| \\ &\quad + \rho\big(w_{ij}(x_i,x_j)+\epsilon\big)\sigma_{ij}\big[\delta_{ij}(\mathcal{C}_{[i]}^e) - \delta_{ij}(\mathcal{C}_{[i]})\big]\Big),\end{aligned}$$

using the fact that $J_i^{\text{consensus}} = 0$ at $(\mathcal{C}_{[1]}^e,\ldots,\mathcal{C}_{[M]}^e)$. Subsystem i may either refine $\mathcal{C}_{[i]}^e$ or coarsen it. If the latter, then $\delta_{ij}(\mathcal{C}_{[i]}^e)\geq\delta_{ij}(\mathcal{C}_{[i]})$ for all $j\in\mathcal{M}_i$ and all i, and, since the first term in $\Delta\phi$ is always non-negative, the potential function can not decrease. If the former, however, then second term in $\Delta\phi$ may be negative: we need to show that, if the provided bound holds and ϵ is sufficiently small, then the overall difference in potential is still positive.

Note that $w_{ij}(x_i,x_j)\big|\delta_{ij}(\mathcal{C}_{[i]}^e) - \delta_{ij}(\mathcal{C}_{[j]}^e)\big| = 0$ for all $j\in\mathcal{M}_i$, so, for all pairs (i,j), either $w_{ij}(x_i,x_j)=0$ or $\big|\delta_{ij}(\mathcal{C}_{[i]}^e) - \delta_{ij}(\mathcal{C}_{[j]}^e)\big| = 0$. Let $\tilde{\mathcal{M}}_i$ denote the subset of $\mathcal{M}_i$ for which $w_{ij}(x_i,x_j)>0$. The potential difference is

$$\begin{aligned}\Delta\phi = &\sum_{j\in\tilde{\mathcal{M}}_i} w_{ij}(x_i,x_j)\Big(\big|\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[j]}^e)\big| \\ &-\rho\sigma_{ij}\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big]\Big) \\ &-\sum_{j\in\mathcal{M}_i}\rho\epsilon\sigma_{ij}\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big].\end{aligned}$$

For the first term, since $\big|\delta_{ij}(\mathcal{C}_{[i]}^e) - \delta_{ij}(\mathcal{C}_{[j]}^e)\big| = 0$ then $\delta_{ij}(\mathcal{C}_{[i]}^e) = \delta_{ij}(\mathcal{C}_{[j]}^e)$ for all $j\in\tilde{\mathcal{M}}_i$. It follows—also using the fact that $\delta_{ij}(\mathcal{C}_{[i]}^e)\leq\delta_{ij}(\mathcal{C}_{[i]})$ under refinement—that $\big|\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[j]}^e)\big| = \big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big]$ and so this first term may be written

$$\sum_{j\in\tilde{\mathcal{M}}_i} w_{ij}(x_i,x_j)(1-\rho\sigma_{ij})\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big],$$

which is non-negative if $\rho\sigma_{ij}<1$ for all $j\in\mathcal{M}_i$. We now consider the subtraction of the second term from this one. Suppose that $\rho\sigma_{ij}\leq\gamma<1$ for all i and j. Then

$$\begin{aligned}\Delta\phi \geq (1-\gamma)&\sum_{j\in\tilde{\mathcal{M}}_i} w_{ij}(x_i,x_j)\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big] \\ &- \gamma\epsilon\sum_{j\in\mathcal{M}_i}\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big].\end{aligned}$$

It follows that $\Delta\phi\geq 0$ if, for all $i\in\mathcal{M}$,

$$\epsilon \leq \frac{1-\gamma}{\gamma}\frac{\sum_{j\in\tilde{\mathcal{M}}_i} w_{ij}(x_i,x_j)\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big]}{\sum_{j\in\mathcal{M}_i}\big[\delta_{ij}(\mathcal{C}_{[i]}) - \delta_{ij}(\mathcal{C}_{[i]}^e)\big]}.$$

□

Appendix B.5. Proof of Theorem 5.1

The first part of the proof is to show that if $\mathcal{C}\succeq\mathcal{D}$, then

$$\mathbb{W}_{\mathcal{C}} = \prod_{c\in\mathcal{C}}\bigoplus_{d\in\mathcal{M}_c} A_{cd}\mathbb{X}_d \subseteq \prod_{c\in\mathcal{D}}\bigoplus_{d\in\mathcal{M}_c} A_{cd}\mathbb{X}_d = \mathbb{W}_{\mathcal{D}}.$$

Consider an arbitrary partition $\mathcal{C} = \{c_1,\ldots,c_C\}\in\Pi_{\mathcal{M}}$, and a refinement $\mathcal{D} = \{c_1,\ldots,c_{C-1},c',c^*\}$; that is, coalition c_C in the first partition is split into two coalitions, containing subsystems c' and c^* such that $c'\cup c^* = c_C$. If the set of neighbours for coalition c_C is $\mathcal{M}_{c_C} = \{d\in\mathcal{C} : A_{Cd}\neq 0\}$ then the sets of neighbours for the new coalitions are

$$\begin{aligned}\mathcal{M}_{c'} &= \big\{d\in\{c_1,\ldots,c_{C-1},c^*\} : A_{c'd}\neq 0\big\}\\ \mathcal{M}_{c^*} &= \big\{d\in\{c_1,\ldots,c_{C-1},c'\} : A_{c^*d}\neq 0\big\}\end{aligned}$$

such that $\mathcal{M}_{c'} \setminus \mathcal{M}_{c_C} = \{c^*\}$ and $\mathcal{M}_{c^*} \setminus \mathcal{M}_{c_C} = \{c'\}$. Then

$$\mathbb{W}_{c'} = \bigoplus_{d \in \mathcal{M}_{c'}} A_{c'd} \mathbb{X}_d = \bigoplus_{j \in \mathcal{H}_{c'}} \prod_{i \in c'} A_{ij} \mathbb{X}_j$$

where $\mathcal{H}_{c'} = \bigcup_{d \in \mathcal{M}'_c} d$, and with a similar expression for $\mathbb{W}_{c^*}$. Note also that $\mathbb{W}_{c_C} = \bigoplus_{j \in \mathcal{H}_i} \prod_{i \in c_C} A_{ij} \mathbb{X}_j$ with $\mathcal{H}_{c_C} = \bigcup_{d \in \mathcal{M}_{c_C}} d$ and then

$$\mathbb{W}_{c'} \times \mathbb{W}_{c^*} = \bigoplus_{j \in \tilde{\mathcal{H}}_i} \prod_{i \in c_C} A_{ij} \mathbb{X}_j \supseteq \mathbb{W}_{c_C}$$

where $\tilde{\mathcal{H}}_{c_C} = \bigcup_{d \in (\mathcal{M}_{c'} \cup \mathcal{M}_{c^*})} d$ and the latter inclusion follows from the fact that $\mathcal{M}_{c'} \cup \mathcal{M}_{c^*} \supseteq \mathcal{M}_{c_C}$. Since $\mathbb{W}_{\mathcal{C}} = \prod_{c \in \mathcal{C}} \mathbb{W}_c$, it follows that $\mathbb{W}_{\mathcal{C}} \subseteq \mathbb{W}_{\mathcal{D}}$; since $\mathcal{C}$ and $\mathcal{D}$ were arbitrary, the result holds in general.

The second and final part is to show that $\mathbb{W}_{\mathcal{C}} \subseteq \mathbb{W}_{\mathcal{D}}$ implies $\mathcal{R}_{\mathcal{C}} \subseteq \mathcal{R}_{\mathcal{D}}$. Consider again the refinement $\mathcal{D}$ and suppose the RCI sets for the c^* coalitions are $\mathcal{R}_{c_1}, \ldots, \mathcal{R}_{c_{C-1}}$, $\mathcal{R}_{c'}, \mathcal{R}_{c^*}$. The product of the latter two coalitional sets is $\mathcal{R}_{c'} \times \mathcal{R}_{c^*}$, associated with the disturbance set $\mathbb{W}_{c'} \times \mathbb{W}_{c^*}$. Since $\mathbb{W}_{c'} \times \mathbb{W}_{c^*} \supseteq \mathbb{W}_{c_C}$, then there exist two scalars $0 < a < b \leq 1$ such that

$$a(\mathbb{W}_{c'} \times \mathbb{W}_{c^*}) \subseteq \mathbb{W}_{c_C} \subseteq b(\mathbb{W}_{c'} \times \mathbb{W}_{c^*}).$$

Then consider the set

$$\bigoplus_{l=0}^{h-1} D_l(\mathbf{M}_h) b(\mathbb{W}_{c'} \times \mathbb{W}_{c^*}) = b(\mathcal{R}_{c'} \times \mathcal{R}_{c^*}).$$

Since $b(\mathcal{R}_{c'} \times \mathcal{R}_{c^*}) \subseteq (\mathcal{R}_{c'} \times \mathcal{R}_{c^*})$ is RCI for $b(\mathbb{W}_{c'} \times \mathbb{W}_{c^*})$, it must be RCI for any subset of $b(\mathbb{W}_{c'} \times \mathbb{W}_{c^*})$, including $\mathbb{W}_{c_C}$. Thus, $b(\mathcal{R}_{c'} \times \mathcal{R}_{c^*})$ outer-bounds $\mathcal{R}_{c_C}$. $\square$

Appendix B.6. Proof of Theorems 5.2 and 5.3

Let $x(0) \in \bar{\mathcal{X}}^N_{\mathcal{C}(0)}$, where the initial partition is $\mathcal{C}(0)$. Since $\bar{\mathcal{X}}^N_{\mathcal{C}(0)} \subset \mathbb{R}^n$ is compact, there exists an $r > 0$ such that $x(0) \in \mathcal{B}_r \triangleq \{x : |x| \leq r\}$. Take $\zeta_{\mathcal{C}(0)} = \min_{\bar{x}} \left\{ \bar{V}^0_{\mathcal{C}(0)}(\bar{x}) : |\bar{x}| = r \right\}$, where $\bar{V}^0_{\mathcal{C}(0)}(\bar{x})$ denotes the collective value functions of the primary MPC controllers, *i.e.*, the sum of $\bar{V}^0_c(\bar{x}_c)$ over $c \in \mathcal{C}(0)$. The associated sub-level set is $S(\zeta_{\mathcal{C}(0)}) = \left\{ \bar{x} : \bar{V}^0_{\mathcal{C}(0)}(\bar{x}) \leq \zeta_{\mathcal{C}(0)} \right\} \subseteq \mathcal{B}_r$. In line with the hypothesis on the switching signal, the system remains in the partition $\mathcal{C}(0)$ for a number of time steps; call this number k_s. By Theorem 3.1, the value function at time $k = k_s$ is bounded as $V_{\mathcal{C}(0)}\left(\bar{x}(k_s)\right) \leq \gamma^{k_s}_{\mathcal{C}(0)} \zeta_{\mathcal{C}(0)}$ where $\gamma_{\mathcal{C}(0)} = \max_{c \in \mathcal{C}(0)} \gamma_c$, and $\gamma_c \in (0, 1)$ is the decay constant for the primary controller for coalition c from Theorem 3.1.

Suppose that, at $k = k_s$, the strongly feasible partition $\mathcal{C}(k_s) \in \Pi_{\mathcal{M}}$ is selected, such that $x(k_s) \in \bar{\mathcal{X}}^N_{\mathcal{C}(k_s)}$. The value functions $\{\bar{V}^0_c(\cdot)\}_{c \in \mathcal{C}(k_s)}$ are bounded on the individual sets $\bar{\mathcal{X}}^N_c$ such that $\theta_{\mathcal{C}} \triangleq \sup \left\{ \bar{V}^0_{\mathcal{C}(k)}(\bar{x}) : \bar{x} \in \bar{\mathcal{X}}^N_{\mathcal{C}} \right\}$ is attained as the maximum, and is finite, for any $\mathcal{C} \in \Pi_{\mathcal{M}}$; in addition, define $\mu_{\mathcal{C}} = \zeta^{-1}_{\mathcal{C}}(\theta_{\mathcal{C}} + \epsilon_{\mathcal{C}})$ for an $\epsilon_{\mathcal{C}} > 0$, and $\mu = \max \{\mu_{\mathcal{C}} : \mathcal{C} \in \Pi_{\mathcal{M}}\}$. The value function for $\mathcal{C}(k_s)$ satisfies $V_{\mathcal{C}(k_s)}\left(\bar{x}(k_s)\right) \leq \theta_{\mathcal{C}(k_s)} < \theta_{\mathcal{C}(k_s)} + \epsilon_{\mathcal{C}(k_s)} \leq \mu \zeta_{\mathcal{C}(0)} = \mu V_{\mathcal{C}(0)}\left(\bar{x}(0)\right)$. Thus, the conditions of Lemma 5.1 are met, and the system is locally stable. The final part of the proof establishes attractivity of the origin, as sufficient to prove Theorem 5.3. Owing to the previous result, the switched value function remains bounded, and the system remains feasible by hypothesis, for all $k \in [0, k_f]$. For $k \geq k_f$, it is assumed that $\sigma(k) = \bar{\mathcal{C}}$

$\square$