



This is a repository copy of *Life lessons from and for distributed MPC – Part 2: Choice of decision makers*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/139572/>

Version: Published Version

Proceedings Paper:

Olaru, S., McNamara, P., Canizares, J.C. et al. (4 more authors) (2018) Life lessons from and for distributed MPC – Part 2: Choice of decision makers. In: Kopacek, P., Wien, T.U. and Ibrahimov, B., (eds.) IFAC-PapersOnLine. 18th IFAC Conference on Technology, Culture and International Stability (TECIS18), 13-15 Sep 2018, Baku, Azerbaijan. International Federation of Automatic Control , pp. 107-111.

<https://doi.org/10.1016/j.ifacol.2018.11.257>

© 2018 IFAC. Reproduced in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Life lessons from and for distributed MPC

– Part 2: Choice of decision makers

S. Olaru, P. McNamara, J.C. Cañizares, M. Farina, J.M. Maestre,
P. Trodden, R.R. Negenborn.

Abstract: This paper and an accompanying paper (McNamara et al., 2018) revisit the Distributed Predictive Control (DMPC) literature and seek to establish links with the social behaviour, focusing in particular on ways in which DMPC could be used to provide insights into the mechanisms of group regulation in social systems. It will be noted that there are major differences between the way in which DMPC algorithms and Social Human Participants (SHPs) form their respective decisions.

Whereas in a first paper (McNamara et al., 2018) we concentrated on the dynamics of the cooperation and the weightings in the agents' decision, the present paper extends the discussion to the arrangements in the group of decision makers. This paper concludes with some caveats as regards further analyses of social system using the methods proposed in these two papers.

© 2018, IFAC (International Federation of Automatic Control) Hosting by Elsevier Ltd. All rights reserved.

Keywords: Distributed Model Predictive Control, Social Systems.

1. INTRODUCTION

Social systems are composed of dynamically interacting individuals and groups with interdependent goals. These individuals and groups are capable of taking actions and sensing their environments, but are limited by all sorts of constraints. There are several issues that can arise when these parties seek to make decisions together. Examples of such issues are an inability to reach consensus on decisions either efficiently or at all, displaying a lack of empathy for the other decisions makers, bad leadership, a lack of transparency in decision making, etc. The area of Distributed Model Predictive Control (DMPC) provides a number of novel methods for automated decision making in group scenarios in which individual decision makers, called agents, communicate with each other in some fashion in order to control a number of interconnected subsystems.

In this article, it is proposed that DMPC techniques could be useful for the general analysis of agent systems such as social systems. Moreover, the straightforward analysis of such systems using abstract mathematical concepts can be problematic. We provide several caveats in this regard at the end of the article. Thus this article seeks to see if there's the potential for "life lessons" to be learned by the life sciences from results in the DMPC literature, and equally posit a number of ways that the DMPC literature could learn from previous experiences gained in the past using other mathematical tools to analyse the life sciences.

Society is composed of groups or individuals, which henceforth are referred to as Social Human Participants (SHPs). The actions of SHPs in trying to achieve some goal, typically have consequences not only for the environment of the SHP responsible for the action but also for other SHPs who are connected in some way to the SHP responsible for the original action.

Then SHPs will have a range of goals that they wish to fulfil and will seek to achieve these goals using the mental models that they have of the particular system with which

they are engaged. However, as SHPs must share resources and dynamically interact with other SHPs, some degree of collaboration with other SHPs is necessary in order to achieve these goals. Thus, SHPs must consider the actions of other SHPs in order to reach their objectives. Equally the models SHPs have of external SHPs will typically be based on experience. For example, over time people will have developed an understanding of what is acceptable social behaviour in various situations and will have an idea of the likely consequences of their actions in various circumstances.

The paper is organized as follows. Section 2 recalls the Distributed MPC framework in order to establish a link with the companion paper and to offer the necessary elements for the main analysis. Section 3 presents a series of insights on the arrangements of the decision makers and performs the analysis in parallel between the DMPC and SHPs. The paper is completed by two sections, one dedicated to the caveats of the present study and the second to the conclusions and outlook.

2. RECALL OF TYPICAL DMPC FORMULATION

In order to provide a self-contained material, this section recalls the basic description of DMPC problems. The notation and formulation are those introduced in the companion paper (McNamara et al., 2018). The prediction mechanism builds on a discrete-time, linear, time-invariant state-space model for each subsystem i ,

$$\mathbf{x}_i(k+1) = \mathbf{A}_i \mathbf{x}_i(k) + \mathbf{B}_i \mathbf{u}_i(k) + \mathbf{V}_i \mathbf{v}_i(k) \quad (1)$$

$$\mathbf{y}_i(k) = \mathbf{C}_i \mathbf{x}_i(k), \quad (2)$$

where $\mathbf{x}_i(k)$, $\mathbf{u}_i(k)$, and $\mathbf{y}_i(k)$ are the states, inputs, and outputs of the i^{th} subsystem at sample step k , respectively, and $\mathbf{v}_i(k)$ are external inputs from other subsystems that influence subsystem i at sample step k .

Using this discrete-time model, the i^{th} subsystem's trajectory over H sample steps into the future can be mathematically described using the information of the current

states and the future control action, where H is called the prediction horizon. Using a centralised predictive control approach, a finite-time optimal control problem is solved at each sample step for a system of N interconnected subsystems

$$\min_{\tilde{\mathbf{u}}_1(k), \dots, \tilde{\mathbf{u}}_N(k)} \sum_{i=1}^N w_i J_i^{\text{local}}(k), \quad (3)$$

subject to constraints, where the cost function $J_i^{\text{local}}(k)$ accounts the control goals of area i , $\tilde{\mathbf{u}}_i(k)$ are the values of \mathbf{u}_i over the H predicted samples steps, and the weight w_i determines the relative importance of minimising $J_i^{\text{local}}(k)$ in the cost function. Each agent then applies only $\mathbf{u}_i(k)$ to the system and repeats this process each sample step by receding the prediction horizon. It should be noted that tuning of the weights w_i can have a significant effect on how the system operates. A discussion of some of the implications of weight tuning in the context of distributed MPC is given in Section 4 of the companion paper.

As mentioned in (McNamara et al., 2018), some distributed controllers are built with the capability to solve (3) in a non-centralised iterative fashion, where the i^{th} agent solves for $\tilde{\mathbf{u}}_i$, and the result is Pareto optimal Venkat (2006). This implies that each agent has access to the global system model that a centralised controller does, and all agents can communicate with each other. As in cooperative optimization routines, found in game theory, these agents seek to solve the global system goal of (3) based on access to a global system model. These algorithms are also called Cooperative DMPC algorithms.

Often agents only have access to local variables and then may be only capable of communication with agents with whom they share an interconnecting variable. Given that these algorithms are based on local cost functions they are referred to as Non-Cooperative distributed MPC algorithms. In the simplest form of such algorithms, agents will solve:

$$\min_{\mathbf{u}_i(k)} J_i^{\text{local}}(k), \quad (4)$$

subject to some constraints.

If agents are allowed to build on an inter-agent communication it is typically then possible for them to achieve performance ranging from that achievable using (4) to that using (3). Many of these solutions will take a form which explicitly accounts the local performance index and the interconnection cost:

$$\min_{\tilde{\mathbf{u}}_i(k), \tilde{\boldsymbol{\theta}}_i(k)} w_i J_i^{\text{local}}(k) + J_i^{\text{inter}}(k), \quad (5)$$

subject to constraints, where the $J_i^{\text{inter}}(k)$ cost is designed to allow agent i deal with interconnecting constraints. The vector $\tilde{\boldsymbol{\theta}}_i$ represents a collection of variables used to coordinate the actions of the i^{th} agent with other agents with whom the i^{th} agent shares an interconnecting variable. For example, in Negenborn et al. (2008) $J_i^{\text{inter}}(k)$ is used to allow agents to reach consensus on interconnecting variables over the prediction horizon in an iterative fashion, and $\tilde{\boldsymbol{\theta}}_i$ are the values of the interconnecting variables that the i^{th} agent would like to receive. Typically these algorithms achieve, at best, a Nash optimal response, which provides performance somewhere between that of a selfish MPC algorithm, where agents only optimise for

$J_i^{\text{local}}(k)$, without communicating with other agents, and an algorithm which achieves Pareto optimal performance. Often these algorithms are referred to as non-cooperative DMPC algorithms.

The preceding paragraphs were presenting a common basis with the companion paper (McNamara et al., 2018). They do not offer an exhaustive account of the range of distributed MPC algorithms that have been developed, and merely serve to give a general flavour of the way in which distributed MPC can be solved. The vast DMPC literature contains an array of techniques that have been developed based on varying mathematical approaches, and system and communication architectures. For more technical descriptions the reader is referred to Maestre and Negenborn (2014); Negenborn and Maestre (2014). Having briefly introduced the DMPC mechanisms, the following sections will discuss insights into the organisation of social systems which can be gained from analyses based on results from the DMPC literature and the selection of the sources of information for the decision making.

3. INSIGHTS INTO THE ARRANGEMENT OF GROUP DECISION MAKING

In any large group of interacting agents it is of interest to observe the phenomena related to the way in which agents both choose their decision makers and organise decision making processes. The following techniques from the DMPC literature provide insights into the types of SHPs that others are likely to wish to deal with in decision making processes, and insights are given regarding the ways groups may optimally choose between decision making structures, which in turn could be useful in analysing the dynamics of systems in found social structures, such as the electoral system.

3.1 Reliability of SHPs and stability?

When a SHP acts in a shared environment with others, they must gain knowledge upon the other SHPs' behavior, in order to try predict those individuals' future actions or states, and possibly how their actions/states affect their own state. Through the development of general models for how another behaves over time (possibly identified through interaction), a SHP can then predict to a certain extent how another SHP will behave. Naturally, these predictions are affected by uncertainty. To guarantee safe interactions with other individuals, uncertainty upon their behaviour must be suitably counteracted.

The DMPC method presented in Farina and Scattolini (2012) is based on the following idea: each subsystem assumes that the future state/input trajectories of the interacting subsystems lie in a pre-defined bound of some nominal ones (which are known in advance) and seeks to minimize a local cost without having to make any further assumption about the behaviour of the neighbors. A similar philosophy is followed in Grancharova and Olaru (2015) for modelling the uncertainty by means of (polytopic) bounds on the parameters of the linear prediction model of each subsystem (1). In a few words, in these DMPC strategies each agent seeks to maximize its own "worst case" utility, similarly to max-min solutions of non-cooperative games.

When designing non-cooperative DMPC, the sizes of the uncertainty sets allowed to all subsystems are not freely chosen but must be identified (or, in a way, negotiated) in a collective, although aggregated, fashion. Due to the interdependencies between the uncertainties of each agent, the uncertainty sets of each agent need to be chosen carefully in order to ensure safe and reliable operation of the overall control system.

The uncertainty related to the information (i.e., the predicted state/input trajectories) transmitted by each subsystem to the neighboring ones, in its turn affects the uncertainty related to the predictions that can be computed by the neighbors and, in a few words, propagates throughout the subsystem network. Therefore, only with a careful (and somehow centralized) choice of the uncertainty allowed to each subsystem one can guarantee a safe and reliable operation of the overall (distributed) control system.

Two facts are observed. First, systems which are internally robust (in the sense that the gain between the input provided by the neighbors states/inputs and the internal state variables is low) are able to reject the uncertainty of neighbors efficiently, and therefore avoid propagation (i.e., amplification) of the uncertainty throughout the overall system. Second, the uncertainty related to subsystems influencing many other subsystems (through the model equations) normally heavily propagates through the network of subsystems.

It was possible to infer the following rules of thumb for the choice of the uncertainty levels related to the behavior of each subsystem, also verified through empirical and simulation studies:

- (1) The uncertainty set allocated to more internally robustly stable subsystems should be small, in order to exploit their natural capability to reject/absorb disturbances. On the other hand, the uncertainty set related to less internally robustly stable subsystems should/can be allowed to be greater.
- (2) The uncertainty sets allocated to subsystems which are neighbors of many other ones must be small. To put it another way, in order to guarantee the stability of the overall system, one must design the system in order to make subsystems coupled with many others as robust as possible.

This defines somehow a hierarchy between interacting subsystems: both the most resilient subsystems and the ones affecting many other ones must also be the most reliable and trustful. On the other hand, the less resilient subsystems are allowed to transmit information with more uncertainty but, in order to guarantee stable behaviour in the overall network, their influence upon other agents must be limited.

This in turn can be related to multi-SHP decision making processes. As mentioned above when dealing with other SHPs in decision making processes, individual SHPs will have a perceived model of the other SHPs in the group, and will have an idea of the uncertainties associated with these other SHPs (or an idea of how consistent these SHPs are). Indeed, one tries to be as robust as possible with respect to the unpredicted behavior of neighboring agents assumed to be unreliable (this is done by accounting for the fact that

the uncertainty in their behavior may be great), while the behavior of trustworthy neighbors is assumed to be more close to the presumed one.

A simple example where a distributed robustness-based predictive control algorithm idea is naturally applied is when we drive a car: we assume that the neighboring cars follow some defined "nominal" trajectories (in this case we identify their nominal trajectories based on their actual position, their velocity, the type of road, etc.) but we also assume that our guess can be affected by some uncertainty, i.e., related to the level of confidence. This is basically how we (try to...) avoid making accidents all the time. In case, for example, a neighboring car reacts in an unreliable way to the external environment (sudden changes in direction, velocity, unexpected breaks), the uncertainty level related to its expected trajectory will be high, and we try to increase the distance from it. On the other hand, we normally drive relatively close to cars that run smoothly and reliably.

Related to this, the observation could be made that in general, in decision making processes, SHPs tend to gravitate towards other SHPs who appear to be highly stable and reliable, and on the other hand, the more connections an individual SHP has, the more reliable that SHP is considered to be.

3.2 Choosing representative bodies

In social systems there are many examples of systems that need regulation of a number of outputs, and for these outputs, decisions are made at regular intervals as to the representative structures that should be used to regulate these systems. The illustrative example that will be used here is the process of Democratic Elections (DEs).

DEs could be considered to operate as follows. There are various systems such as our economic, legal, educational, etc., systems that need to be managed. These systems additionally provide various indicators as to how well they are being managed. For example, it is desirable that inflation would be kept at reasonable levels from the perspective of our economic systems, and it is desirable that high percentages of the population are receiving a good education. Political parties could then be considered to provide various models for the management of these systems, which look at the outputs, and based on their perceived model of the system, regulate the inputs to the system, e.g., funding for education, interest rates, etc. Based on the performance of different political parties, who use various combinations of outputs and models to control these systems, then the public make decisions at fixed intervals as to which parties should be used to control the system next. Thus this defines a model of the democratic system of elections. While it may not be possible to comment on the optimality of the decision to elect a party, an optimal system of choosing amongst these various decision agents could provide useful insights into the consequences of choosing representatives in this fashion.

In Stoican and Oлару (2013) a predictive method is developed for choosing between arrays of healthy and faulty sensors, which in turn are used for informing the control

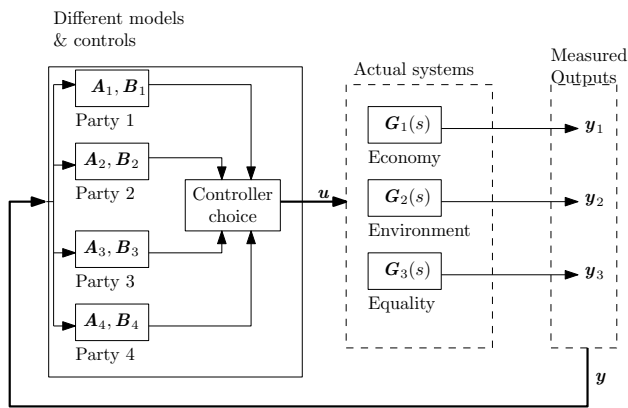


Fig. 1. The method of choosing healthy and faulty sensors, with democratic electoral system equivalents.

of the system. An example system for which this may be applied is given in Fig. 1 and it is now shown how analogies can be drawn between this system and the formation of representative systems, such as the DE process. The real systems are given by G_1, G_2, G_3 , and for the DE analogy these represent the different systems to be controlled by the government. Each of these systems produce the outputs y_1, y_2, y_3 which are measured, and fed back into a number of different controllers, which have models A_i related to the certain estimated states of the system, and B_i related to certain inputs of the system. The aim of the work in Stoican and Olaru (2013) was to be able to optimally determine which sensors were healthy and which were faulty, in order to improve system performance through the removal of the faulty sensors.

Parties can be considered as combining A_i, B_i models with sensor outputs as proposed in Stoican et al. (2014) for switched control configurations. Thus the same analysis that is used for choosing which sensors to use for control can be extended for deciding which models/parties to optimally choose for controlling the system. Therefore several of the issues related to the choice of sensors could be equally related to the system of deciding which parties or combinations of parties should be used to control the system.

For instance, one of the issues outlined in Stoican and Olaru (2013) is that it is possible for faulty sensors to remain in a pool of healthy sensors. The effect of these faulty sensors may not be differentiable from an unknown disturbance, and could potentially bias the entire closed loop. The lesson here from a political perspective, is that it might not be possible to detect that errors in the regulation of a particular system are due to the approach being taken by the party to regulate it. This is because the healthy parts of the system may compensate for the faulty parts, and it may not be possible to determine if errors are due to the faulty model or unknown disturbances.

Another issue in Stoican and Olaru (2013) is that it is only possible to determine when a sensor is faulty once the system begins to deviate from its optimal path, i.e., the only time when it can be detected that a sensor is faulty is when the system behaves in a way not predicted by the system model. Likewise, in terms of political systems it could be recognised that when systems appear to be

running as desired the ability of a party to manage the system is rarely called into question. It is only when our financial, social aid, etc systems deviate away from their desired performances that people begin to recognise that there might indeed be flaws with the systems in place. The lesson that can be learned from Stoican and Olaru (2013) is that a potential method for ensuring that the right models are being used for control of a system, would be to purposefully make the system deviate from it's desired behaviour even though it is currently achieving it's desired outcome. However, it's unlikely that any political party is likely to willingly adopt such a strategy in reality!

4. CAVEATS ON THE APPLICATION OF THESE INSIGHTS

It is necessary to be careful when analysing social systems using abstracted mathematical formulations. Often engineering systems can be modelled as linear systems in which there is some bounded uncertainty, for the purposes of designing a controller for the system. Social systems may not be so well behaved, and may exhibit highly nonlinear behaviours. Social systems may also exhibit non-equilibrium behaviour that is not accounted for in models which assume the system is in equilibrium. Also, a SHP may not necessarily behave as either a function minimiser or a satisficer Farmer (2012); Borrill and Tesfatsion (2011). If ever there was a strong example of where the use of abstract models for controlling systems can go wrong, it is in economics, where a new era of simulation based analysis has been developed in recent years, following the failure of abstract models such as the Dynamic Stochastic General Equilibrium models to predict economic crises Farmer (2012).

However, lessons can be learned from the previous uses of Game Theory for the prediction of phenomena in group behaviour. There have been several cases in which, once an experiment is carefully designed, the theoretical predictions of the theory were observed in reality, and the use of game theory for auction design in the UK 3G network was highly successful Wooldridge (2012). A caveat however, is that often Game Theory is used for analysis in places where the underlying assumptions are wholly inappropriate and so it is important to be careful of where such analyses are applied. Additionally, where a real life experiment may not be capable of confirming a theoretical result, simulation using learning agents could be used. For example in Krause et al. (2006), the Nash equilibria in a power system market are derived and are found to parallel those that are found by decentralised reinforcement learning based agents.

5. CONCLUSIONS AND OUTLOOK

In this article and the companion paper (McNamara et al., 2018) a number of observations from the Distributed Model Predictive Control (DMPC) literature are used to illustrate the potential of this body of work to provide insights into the operation of social systems. Furthermore once these insights had been discussed a number of caveats were provided as regards applying such analysis to social systems, as opposed to the application of these techniques in their traditional application domains.

While the authors here have provided insights based on their own experience of developing and using DMPC algorithms, there is likely to be more insights gained from a careful analysis of the DMPC literature. Additionally, society is facing a host of environmental, social, economic, and political challenges for which the current suite of DMPC could hold promising solutions. Through attempts to solve these problems there is no doubt that the short comings of DMPC techniques will also be highlighted and thus the development of DMPC could equally be driven further through interactions between the DMPC and social science domains. Recently Barreiro-Gomez (2018) explore this avenue by analyzing the role of population games in the design of optimization-based controllers.

6. ACKNOWLEDGEMENTS

Financial support from Spanish MEC is acknowledged (project ref. DPI2017-86918-R).

REFERENCES

- Barreiro-Gomez, J. (2018). *The Role of Population Games in the Design of Optimization-Based Controllers*. Springer.
- Borrill, P. and Tesfatsion, L. (2011). Agent-based modeling: the right mathematics for the social sciences? *The Elgar Companion to Recent Economic Methodology*, 228.
- Farina, M. and Scattolini, R. (2012). Distributed predictive control: a non-cooperative algorithm with neighbor-to-neighbor communication for linear systems. *Automatica*, 48(6), 1088–1096.
- Farmer, J. (2012). Economics needs to treat the economy as a complex system. *Institute for New Economic Thinking*.
- Grancharova, A. and Oлару, S. (2015). Distributed robust model predictive control of interconnected polytopic systems. In *Developments in Model-Based Optimization and Control*, 73–91. Springer.
- Krause, T., Beck, E., Cherkaoui, R., Germond, A., Andersson, G., and Ernst, D. (2006). A comparison of Nash equilibria analysis and agent-based modelling for power markets. *International Journal of Electrical Power & Energy Systems*, 28(9), 599–607.
- Maestre, J.M. and Negenborn, R.R. (2014). *Distributed Model Predictive Control Made Easy*. Springer, Dordrecht, The Netherlands.
- McNamara, P., Negenborn, R.R., Cañizares, J.C., Farina, M., Maestre, J.M., Oлару, S., and Trodden, P. (2018). Life lessons from and for distributed MPC – Part 1: Dynamics of cooperation. In *IFAC Conference on Technology, Culture and International Stability*. Baku, Azerbaijan.
- Negenborn, R.R., De Schutter, B., and Hellendoorn, J. (2008). Multi-agent model predictive control for transportation networks: Serial versus parallel schemes. *Engineering Applications of Artificial Intelligence*, 21(3), 353–366.
- Negenborn, R.R. and Maestre, J.M. (2014). Distributed model predictive control: An overview and roadmap of future research opportunities. *IEEE Control Systems Magazine*, 34(4), 87–97.
- Stoican, F. and Oлару, S. (2013). Fault Detection and Isolation in Multisensor Systems. *Set-theoretic Fault-tolerant Control in Multisensor Systems*, 27–53.
- Stoican, F., Oлару, S., Seron, M.M., and De Doná, J.A. (2014). A fault tolerant control scheme based on sensor-actuation channel switching and dwell time. *International Journal of Robust and Nonlinear Control*, 24(4), 775–792.
- Venkat, A. (2006). *Distributed Model Predictive Control: Theory and Applications*. Ph.D. thesis, University of Wisconsin-Madison, Wisconsin.
- Wooldridge, M. (2012). Does game theory work? *IEEE Intelligent Systems*, 27(6), 76–80.