

Life lessons from and for distributed MPC – Part 1

Dynamics of cooperation

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

DOI

[10.1016/j.ifacol.2018.11.256](https://doi.org/10.1016/j.ifacol.2018.11.256)

Publication date

2018

Document Version

Final published version

Published in

IFAC-PapersOnLine

Citation (APA)

McNamara, P., Negenborn, R. R., Cañizares, J. C., Farina, M., Maestre, J. M., Trodden, P., & Oлару, S. (2018). Life lessons from and for distributed MPC – Part 1: Dynamics of cooperation. *IFAC-PapersOnLine*, 51(30), 101-106. <https://doi.org/10.1016/j.ifacol.2018.11.256>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Life lessons from and for distributed MPC

– Part 1: Dynamics of cooperation

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

Abstract: This paper and a second accompanying paper (Olaru et al., 2018) explore the potential of Distributed Predictive Control (DMPC) literature to provide valuable insights into social behaviour. In particular this first paper focuses on 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 decisions. DMPC can make optimal decisions but these are only optimal with respect to a given objective and model, both of which must be explicit. SHPs operate, by and large, with only vague, implicit objectives and models – which can be surprisingly accurate – but often make sub-optimal decisions both individually (because of irrationality or poor anticipation and due to a short horizon, bad model or misjudgement of objectives) and in a group sense (for the previous reasons plus selfishness). Thus while SHPs' decisions would typically be suboptimal, with respect to their desired goals, for the aforementioned reasons, it can be expected that SHPs' decision making would evolve towards an optimal solution as groups of SHPs develop more experience within the system they're operating in.

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

Keywords: Distributed Model Predictive Control, Social Systems.

1. INTRODUCTION

Society is composed of groups or individuals, which henceforth we will collectively refer to as Social Human Participants (SHPs), who share common resources. 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. To complicate things further SHPs are limited in terms of their abilities to manipulate their environments, and have to devise ways of counteracting the effects of phenomena which affect the goals they are trying to achieve. Thus SHPs must constantly monitor the relevant indicators in the environments they are trying to manipulate in order to try and achieve their goals, in the face of varying degrees of uncertainty.

For example, in a car the driver is the SHP, the car can be considered as the part of the environment over which the SHP can have an influence via the actions of manipulating the speed and direction of the vehicle, and the goal can be considered as safely arriving at a particular destination. Uncertainties arise in the form of weather conditions, pedestrians, unpredicted road conditions, etc. And of course, other SHPs drive other vehicles on the road, and so SHPs interact and must coordinate their actions so that each driver may reach their destination safely.

Typically each SHP will have developed a model both of how its own actions affect itself, and how these actions affect other SHPs with whom it shares these resources. These models might be built up over time by the SHP based on their experience of the system and may be very simplistic. For example, using the vehicle analogy, the SHP driver will not have an explicit knowledge of the

internal mechanisms of the engine. Instead they will have developed a tacit knowledge over time that, by turning the wheel a certain amount, the vehicle can be expected to go a certain direction; or that, by pressing the accelerator, the car can be expected to speed up by a certain amount.

However, these SHPs will in turn be aware that there is some uncertainty related to these models, and that this uncertainty is in some way known too. For example, when groups drive on a road, an individual driver will be aware of the fact that drivers in adjacent lanes might change lanes very fast occasionally. Thus that driver will in turn account for this potential behaviour when driving, assuming that this can happen for, maybe, 5% of the time. However, drivers will usually know with great certainty that if they stop at a traffic light the person behind them is likely to stop too.

SHPs also know that there are limits associated with the various inputs and outputs of these systems, and that some of these limits may arise through the interaction with other SHPs, e.g., a car can travel at 150 km/h on the motorway but in a city centre, with heavy traffic, a car can only travel safely at 20 km/h.

Then SHPs will have a range of goals that they wish to fulfil and will seek to achieve these goals using the so called 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 a knowledge of what is acceptable

social behaviour in various situations and will have an idea of the likely consequences of their actions in various circumstances.

In some cases it may not be necessary to explicitly communicate with other SHPs in order to satisfy these objectives. For example, when walking down the street individuals can navigate satisfactorily based on their knowledge of how others are likely to act, and so explicit negotiation and communication is not necessary between the relevant individuals. However, for complex problems such as the coordination of workers in a large factory, SHPs need to use greater levels of communication in order to coordinate their responses. Furthermore, the development of suitable hierarchical structures for the efficient processing and relaying of this communication is of vital importance to ensure the smooth operation of a particular system.

Finally, SHPs may use some form of prediction based on their models of a system in order to best determine what course of action to take over a given period. However, SHPs are typically aware of the fact that there will be errors in these predictions and that unpredictable events can change the trajectory of their predictions, and so they will update these predictions over time based on an updated “measurement” of their position with regards to fulfilling their goals, and in turn update their predictions so as to decide how best to act again. By iterating this process, SHPs counteract the inherent uncertainty associated with achieving their goals, rather than simply assuming that their predictions are totally accurate, effectively performing a continuous replanning process. For example, stock market investors will use predictions as to how certain stocks are expected to grow over time in order to inform their decisions as to how they should invest funds. However, the decision as to whether to maintain this investment or not must be updated at regular intervals as it’s highly unlikely that these predictions will be totally accurate (note here that these predictions could be based on mathematical models or objective experience of the investor. Either way, an individual investor will form some belief as to how the market will behave in the future and invest accordingly).

In turn, if a given group of SHPs interact and communicate with each other in order to coordinate their actions over time, all of them will stand a better chance of satisfying their goals (of course provided that SHPs’ models of the system are accurate). In social systems this communication, often, may not take the form of a direct coordinating signal between SHPs. In groups of people often it will be possible to coordinate actions by observing the actions of others in the group and acting appropriately. Often it can be the case that SHPs respond to a coordinating signal. For example, in the stock market scenario, based on aggregated information from the market a government may change the trading value of a currency in order to attempt to affect the actions of investors. In cases where it is perceived that systems for coordinating SHPs are not well designed, analysis of optimal forms of coordination between SHPs could potentially provide valuable insights into methods for coordinating groups of interacting SHPs, allowing for improvements in the construction of these systems.

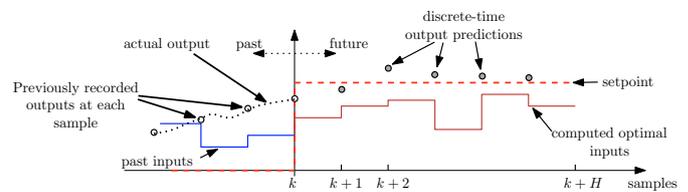


Fig. 1. State-space predictions are used in MPC to determine optimal control inputs over a prediction horizon H .

The idea of planning based decision making has been formalised by engineers in the form of Model Predictive Control (MPC) (Maciejowski, 2002). Here a discrete-time system model is used to form predictions as to how the system behaves over a prediction horizon of H sample steps, as illustrated in Fig. 1 (it should be noted that continuous-time MPC algorithms have been formulated but the results referred to in this article are all related to the discrete-time domain). The system monitors a number of states which indicate how the system is evolving with time, has a number of inputs that can affect the trajectory of the states, is typically subject to measurable or unmeasurable disturbances, and has some outputs which are typically either one of the states of the system or are derived from a combination of these states.

Using predictions based on the state-space model of the system, an optimisation problem is formed which embodies the goals of the system. Solving this optimisation problem, while considering the constraints on the system’s parameters, yields the optimal inputs that should be applied to the system over the full prediction horizon to achieve the desired goal. The inputs for the first sample step of the horizon are applied to the system. Then the control system waits for another sample step and runs this algorithm again, looking one sample step further into the future than in the previous sample.

In recent years there has been extensive research into the area of Distributed MPC (DMPC) (Maestre and Negenborn, 2014). In DMPC a number of autonomous controllers of interconnected systems, called agents, are assigned the control of different interconnected subsystems, and using different types of communication these agents coordinate their control actions so as to fulfil their local objectives. The way in which agents in DMPC make their decisions and the way in which SHPs make their decision in society have many commonalities, as has been described in the previous paragraphs.

Tools such as Game Theory have been used previously to provide new perspectives on the behaviour of groups of interacting agents Dawkins (1976); Axelrod (2006); Binmore (2005). A wide variety of DMPC methods have been developed which provide an array of options for distributing control between agents in a stable fashion, and in turn these algorithms provide new perspectives on group coordination. Thus far the application of DMPC techniques has focused on typical engineering applications such as electrical power systems Hermans et al. (2012); Moradzadeh et al. (2011); Negenborn et al. (2008), process engineering systems Venkat (2006); Liu et al. (2009), water networks Leirens et al. (2010); Zafra-Cabeza et al. (2011), etc. However, there is the potential for the results

from DMPC to provide valuable insights into multi-agent behaviour in general, which in turn could be used for the analysis of non-engineering systems such as social, biological, or economic systems. Indeed, it has been demonstrated previously how the integration of MPC with such systems has the potential to significantly enhance the decision making of individual SHPs, as in the case of Human-in-the-Loop Model Predictive Control van Overloop et al. (2015). However, one must be cautious when making analyses of these systems based on mathematical concepts as one could be drawn to inaccurate conclusions if the assumptions made in the mathematical models are not accurate for the system being studied.

In this paper we explore the potential of various results from the DMPC literature to provide valuable insights into 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 is important to note the nature of these observations. Almost inevitably the group decisions made by SHPs will be suboptimal, as interacting SHPs, in general, do not use computation for decision making in the sense that DMPC does, and the “models” used by SHPs will never be as exact as those used by DMPC for typical engineering applications. However, it can be expected that SHPs’ decision making would evolve towards an optimal solution as groups of SHPs develop more experience within the system they’re operating in. Thus observations of how optimal decision makers behave can be used to provide insight into how certain properties of SHP decision making came to evolve or might evolve into the future. Thus observations made in this paper effectively provide qualitative insights into certain aspects of decision making where groups of SHPs are involved, as opposed to providing explicit formulations of how a particular SHP might make a decision. Given the potential for naively analysing systems of SHPs using DMPC, some caveats need to be considered as regards the straightforward application of these results for social systems analysis.

2. TYPICAL DMPC FORMULATION

Typically the following model is used for the dynamics of a subsystem i ,

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

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

The matrices \mathbf{A}_i , \mathbf{B}_i , \mathbf{V}_i , and \mathbf{C}_i are the relevant state-space matrices that embody the dynamics of the i^{th} subsystem and the influence of control inputs ($\mathbf{u}_i(k)$). The external inputs from other subsystems is represented by $\mathbf{v}_i(k)$. The effects of noise or uncertainties may also be included in these models depending on the particular approach that is used. Predictions can be made as to the i^{th} subsystem’s trajectory over H sample steps, where H is called the prediction horizon. Using a centralised MPC approach, the following optimisation is performed 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 $J_i^{\text{local}}(k)$ embodies 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. Agents then apply only $\mathbf{u}_i(k)$ to the system and repeat this process each sample step. 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.

Some distributed controllers are capable of solving (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 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 most extreme form of these algorithms agents will solve:

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

subject to some constraints. If agents are allowed some 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 similar to the following:

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

subject to constraints, where the $J_i^{\text{inter}}(k)$ cost is designed to allow agent i deal with interconnecting constraints, and $\tilde{\boldsymbol{\theta}}_i$ is a vector 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.

The preceding paragraphs are by no means 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. A vast array of techniques 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).

3. INSIGHTS INTO THE DYNAMICS OF COOPERATION

Results from the DMPC literature provide insights into the dynamics involved in the coordination of groups of agents in control of interacting subsystems. In particular the literature on this area has provided a number of insights into the degree of communication needed between agents

to effectively cooperate their responses and the dynamics of the cooperation itself.

3.1 Transparency, Optimality, and Stability

Right across modern society, in its economic, social, environmental, and political systems, for example, groups of individual decision makers are responsible for making decisions that not only affect the subsystem under their own control, but affect other connected subsystems not under their direct control. With the onset of globalisation, the ties between these individual areas are becoming stronger and the effects of individual actions may have unpredictable consequences. It is perceived that it is these interconnections that are responsible for many of the instabilities that have plagued these systems in the modern era. For example, while the 2008 economic crisis started with institutions based in the United States, the decisions taken with regard to these institutions had knock on effects on economies across the world. In turn the question arises as to what sort of control structures we should expect given such an increasingly interconnected system, and what is the main driving force behind the necessary evolution of these control structures.

From the game theoretic literature, it is widely known that the more information that is shared between agents, the better the overall performance of the system will be. The concepts of the Nash and Pareto equilibria are well known, describing situations where agents reach agreement on those variables connecting their objectives, and the situation where agents have access to the same level of information to make their decisions as would be afforded to a central coordination agent, respectively Venkat (2006). Equally concepts such as the Price of Anarchy (PoA) have been developed in the game theoretic literature in order to quantify the loss in performance when all agents in a system act selfishly as opposed to cooperatively.

The DMPC literature contains a wide variety of techniques that use increasing levels of communication and, as would be expected, the general trend with these control systems is for performance to improve as the inter-controller communication increases Venkat (2006); Hermans et al. (2012); Alvarado et al. (2011); Maestre et al. (2015). Of note, in Venkat (2006), is the fact that as the strength of interconnections between areas increases, it is necessary to increase the level of information that participating agents have access to in order to maintain stable control of the overall system. As the strength of interconnection is increased, first a communication free decentralised controller is driven unstable, and then a Nash equilibrium seeking controller is driven unstable. Only a Pareto equilibrium seeking distributed algorithm is capable of stabilising the system at this stage, which implies all agents have access to the same level of information as a centralised coordinator, that is, all inputs, states, matrices and cost functions are available to the all agents in the system.

Here it can be seen that without a certain level of communication and cooperation in tightly interconnected systems, agents will not be capable of achieving their goals, and this in turn will result in instabilities in the system. This emphasises that instability as opposed to improvements in optimality could be the main driving force in

the development of adopting open, transparent control systems that can maintain stability in heavily interconnected systems. This then provides further insight into the form that the control systems in tightly coupled societies could be expected to take. When societies were not highly interconnected, it was not necessary for governments to coordinate their responses as the decentralised control of the system would not have resulted in instabilities. Through the process of globalisation as countries increasingly interact and become interdependent, it has been necessary for countries to increasingly coordinate their actions in order to maintain stability. A prime example of such interconnection causing instability would be the case of interconnection in the banking system where deleveraging resulted in systematic instabilities due to the high degree of interconnection between financial institutions Caccioli et al. (2014). Given that the strength of these interconnections is increasing continuously it would be expected that it will be necessary for countries to at least consider the objectives of other countries in their responses or for decision making to be centralised such that the objectives of all countries can be considered simultaneously. Indeed, this process would reflect what has been seen throughout history, where the trend for increasingly interconnected societies has been to increase cooperation and empathy between those interconnected agents in society, and to centralise the structures responsible for coordination between the agents in this system Diamond (1998).

Thus, the lesson that is presented by the DMPC literature here is that increasing interconnection between social systems can act as a means by which these social systems will eventually seek to increase their cooperation due to the increased threat of instability in the overall system.

3.2 Decisions based on models of others

A notable aspect related to the problem of group coordination in social systems arises from the fact that their objectives or constraints are often conflicting: actions driven purely by self-interest can lead to compromised outcomes for all. Thus, an interesting question is to what degree should a SHP anticipate the actions or intentions of others in deciding upon his own strategy?

Cooperative DMPC aims to endow control agents with a sense of anticipation about the actions of others. For example, in Trodden and Richards (2013), each agent devises not just its own plan for the future, but also hypothetical plans for other agents, in which their objectives and constraints are taken into account. The idea is that an agent considers, as part of its internal decision making, what other agents might be able to achieve in the future in response to its own actions. Meanwhile, special constraints in the local optimization problem ensure that the coupled constraints, constraints that result in interactions between the two agents optimization problems, are guaranteed to be satisfied regardless of the outcome, that is, whether the other agents take advantage of the situation (by adopting something like the hypothetical plan) or not.

Social group decision problems closely resemble the DMPC problem for dynamically decoupled subsystems sharing constraints and possibly an objective. The question of the degree of anticipation by SHPs in a social system is ana-

logous to the questions in cooperative DMPC of how and when to use cooperation. Social group decision problems closely resemble DMPC problems in which dynamically decoupled subsystems share constraints and possibly an objective. The question of the degree of anticipation by individuals in a social system is analogous to the questions in cooperative DMPC of how and when to use cooperation.

Tuning the cooperative DMPC objective between self-interest and interest only for other agents, maps to different plans between the extremes of acting selfishly and altruistically. For the latter, the local agent is willing to sacrifice its own performance and a less selfish plan is seen to result. No negotiation or iteration is required; an agent unilaterally chooses a cooperative plan. In terms of society, this indicates that mutually satisfactory outcomes can result from SHPs anticipating and accommodating the actions of others in furthering their own interests - the philosophy of enlightened self-interest. Research has shown that there is most benefit to system-wide performance when actively-coupled neighbours cooperate Trodden and Richards (2009). The potential benefit can be estimated from the predicted cost of hypothetical plans. In a societal setting, this suggests that group decision making can be made more mutually beneficial if SHPs consider the objectives, and anticipate the actions, of those with whom their interests are most conflicting.

3.3 Long term coordination is easiest

A general observation that might be made about the ability of large groups of SHPs to coordinate their responses is that it is always easiest for conflicting parties to agree on things in the long term, but it is always the short term plans that are the most difficult to agree on. Everyone will always agree that in long term they want world peace, a clean, pollution-free environment, and resources to be shared fairly amongst the nations of the world. However, usually the agreements on what countries will do next week or next year are far more difficult to settle on.

Interestingly results from the DMPC literature reflect this trend. In Negenborn (2007) distributed controllers are designed that act in an iterative fashion to converge on a Nash equilibrium solution each sample step. In Chap. 2 of Negenborn (2007) a number of figures illustrate how two connected areas typically converge on the agreed values for the final trajectory of an interconnecting variable. It is seen in these diagrams that the values for the variable at the end of the prediction horizon are the first to reach agreement and the values then reach agreement working their way from the final stage of the prediction horizon to the first. A rationale for this could be that the final stage of the horizon offers the most degrees of freedom for reaching cooperation and so requires the least effort in terms of reaching an agreed value. Also, the rules of dynamic programming come to mind here where the optimal trajectory is found by working backwards from the end of the horizon to the start.

The lesson that is found here is that if optimal coordination algorithms find it easiest to coordinate the longest term responses but more difficult to coordinate their short term responses, then it should not be surprising that this would be seen in real world negotiations too.

4. INSIGHTS BASED ON THE WEIGHTING OF DECISION AGENTS

In decision making processes, usually certain decision making agents will have a greater influence than other agents. In the optimisation and control literature this preference is reflected in the different weights allocated to decision makers in the cost function. It is of interest to see how the choice of these weights affects negotiation processes. Interesting insights into these processes can be found in the DMPC literature on weight tuning. In Mc Namara et al. (2013) the weights of a non-cooperative DMPC system were optimised, in order to minimise a setpoint tracking criterion for a highly interconnected power system. It was observed here that each of the agents in the optimal weight case experienced improvements in performance despite some agents' optimal weights being significantly larger than others.

Thus these results imply that it may not be in an SHPs best interests to seek to maximise its own weight in negotiations. Thus this would seem to contradict the intuitive assumption that an SHP should always seek to maximise their weight in negotiations to improve their returns. However, in real life systems the determination of what exactly is an optimal weighting for each SHP is highly non-trivial, and the idea of a SHP purposefully minimising their weight in a negotiation is, in general, undesirable from their perspective and unlikely to occur voluntarily.

However, the consideration of weights as fixed constants is not always the ideal way to model group negotiation. Equally, the modelling of SHPs as maximisers may not always be accurate, particularly in social systems where SHPs may not desire more of a resource once they have satisfied their need for it. A preferable method for modelling SHPs may be to assume that a SHP i is satisfied once their cost function is within a certain level ν_i . In the DMPC literature an agent which coordinates its actions with other agents in this way is called a satisficing agent.

A distributed satisficing MPC method is described in de Lima et al. (2015). It is illustrated in this paper how a distributed algorithm can be designed to minimise the following equivalent centralised satisficing problem:

$$\min_{\mathbf{u}_1(k) \dots \mathbf{u}_N(k)} \sum_{i=1}^N -\log(\nu_i - J_i(k)), \quad (6)$$

In this paper it is in turn shown that (6) is equivalent to the centralised MPC problem given in (3) where the weight associated with area i are given by:

$$w_i = \frac{1}{\nu_i - J_i(k)}. \quad (7)$$

Thus this shows that the satisficing problem is equivalent to a centralised MPC problem in which the weights adaptively change, such that agents who are more satisfied are given less preference, and those who are less satisfied receive more. This view on satisficing as a form of centralised MPC with adaptive weights highlights the advantages groups of satisficers have over groups of maximisers, i.e. satisficers naturally adapt to the situation faced by the group, while maintaining the optimality of a centralised maximisation problem, while groups of maximisers conti-

nue to give preference to the same individual regardless of the situation. This could in turn provide an evolutionary perspective as to why groups would evolve as satisficers instead of maximisers. The recent work Barreiro-Gomez (2018) discuss the role of population games in the design of optimization-based controllers.

5. CONCLUSION

In this paper, 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. It is pointed out that the quantity of information shared among the subsystems influences the global performances and gives a primal role to the interconnection in the stability analysis of both DMPC and SHP. Next, the prediction-based strategy being related to an anticipative capability, it is shown that cooperative MPC relates to altruistic behavior in SHP. Ultimately, the length of the prediction horizon and the weightings in cost function play an important role in the coordination of DMPC as well as in the negotiations of SHP. In a companion paper, additional insights related to arrangement of group decision making are presented and a number of caveats are provided as regards applying such analysis to social systems, as opposed to the application of these techniques in their traditional application domains.

6. ACKNOWLEDGEMENTS

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

REFERENCES

- Alvarado, I., Limon, D., de la Pena, D., Maestre, J., Ridao, M., Scheu, H., Marquardt, W., Negenborn, R., De Schutter, B., Valencia, F., et al. (2011). A comparative analysis of distributed MPC techniques applied to the HD-MPC four-tank benchmark. *Journal of Process Control*, 21(5), 800–815.
- Axelrod, R. (2006). *The evolution of cooperation*. Basic books.
- Barreiro-Gomez, J. (2018). *The Role of Population Games in the Design of Optimization-Based Controllers*. Springer.
- Binmore, K. (2005). *Natural justice*. Oxford Univ. Press.
- Caccioli, F., Shrestha, M., Moore, C., and Farmer, J. (2014). Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46, 233–245.
- Dawkins, R. (1976). *The selfish gene*. Oxford Univ. Press.
- de Lima, M.L., Camponogara, E., Marruedo, D., and de la Pena, D.M. (2015). Distributed Satisficing MPC. *IEEE Transactions on Control Systems Technology*.
- Diamond, J. (1998). *Guns, germs and steel: a short history of everybody for the last 13,000 years*. Random House.
- Hermans, R., Jokić, A., Lazar, M., Alessio, A., Van den Bosch, P., Hiskens, I., and Bemporad, A. (2012). Assessment of non-centralised model predictive control techniques for electrical power networks. *International Journal of Control*, 85(8), 1162–1177.
- Leirens, S., Zamora, C., Negenborn, R.R., and De Schutter, B. (2010). Coordination in urban water supply networks using distributed model predictive control. In *Proceedings of the American Control Conference*, 3957–3962. Baltimore, Maryland.
- Liu, J., Muñoz de la Peña, D., and Christofides, P. (2009). Distributed model predictive control of nonlinear process systems. *AIChE journal*, 55(5), 1171–1184.
- Maciejowski, J. (2002). *Predictive Control with Constraints*. Prentice Hall, Harlow, England.
- Maestre, J.M. and Negenborn, R.R. (2014). *Distributed Model Predictive Control Made Easy*. Springer.
- Maestre, J., Ridao, M., Kozma, A., Savorgnan, C., Diehl, M., Doan, M., Sadowska, A., Keviczky, T., De Schutter, B., Scheu, H., et al. (2015). A comparison of distributed MPC schemes on a hydro-power plant benchmark. *Optimal Control Applications and Methods*, 36(3), 306–332.
- Mc Namara, P., Negenborn, R.R., De Schutter, B., and Lightbody, G. (2013). Weight optimisation for iterative distributed model predictive control applied to power networks. *Engineering Applications of Artificial Intelligence*, 26(1), 532–543.
- Moradzadeh, M., Bhojwani, L., and Boel, R. (2011). Coordinated voltage control via distributed model predictive control. In *Chinese Control and Decision Conference (CCDC)*, 1612–1618. Mianyang, China.
- Negenborn, R.R. (2007). *Multi-Agent Model Predictive Control with Application to Power Networks*. Ph.D. thesis, Delft Univ. of Technology, Delft, The Netherlands.
- 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.
- Olaru, S., McNamara, P., Negenborn, R.R., Cañizares, J.C., Farina, M., Maestre, J.M., and Trodden, P. (2018). Life lessons from and for distributed MPC – Part 2: Choice of decision makers. In *IFAC Conference on Technology, Culture and International Stability*. Baku, Azerbaijan.
- Trodden, P. and Richards, A. (2009). Adaptive cooperation in robust distributed model predictive control. In *Proceedings of the IEEE Control Applications & Intelligent Control Symposium*, 896–901.
- Trodden, P. and Richards, A. (2013). Cooperative distributed MPC of linear systems with coupled constraints. *Automatica*, 49(2), 479–487.
- van Overloop, P.J., Maestre, J.M., Sadowska, A.D., Camacho, E.F., and Schutter, B.D. (2015). Human-in-the-Loop Model Predictive Control of an Irrigation Canal. *IEEE Control Systems*, 35(4), 19–29.
- Venkat, A. (2006). *Distributed Model Predictive Control: Theory and Applications*. Ph.D. thesis, University of Wisconsin-Madison, Wisconsin.
- Zafra-Cabeza, A., Maestre, J., Ridao, M., Camacho, E., and Sánchez, L. (2011). A hierarchical distributed model predictive control approach to irrigation canals: A risk mitigation perspective. *Journal of Process Control*, 21(5), 787–799.