Technical report

Chapter 1 of *Intelligent Infrastructures*:

Intelligence in transportation infrastructures via model-based predictive control

R.R. Negenborn, H. Hellendoorn

If you want to cite this report, please use the following reference instead:

R.R. Negenborn and H. Hellendoorn, "Intelligence in transportation infrastructures via model-based predictive control," Chapter 1 in *Intelligent Infrastructures* (R.R. Negenborn, Z. Lukszo, and H. Hellendoorn, eds.), Dordrecht, The Netherlands: Springer, ISBN 978-90-481-3597-4, pp. 3–24, 2010.

Delft University of Technology, Delft, The Netherlands

Chapter 1
Intelligence in Transportation Infrastructures via Model-Based Predictive Control

R.R. Negenborn and H. Hellendoorn

Abstract In this chapter we discuss similarities and differences between transportation infrastructures like power, road traffic, and water infrastructures, and present such infrastructures in a generic framework. We discuss from a generic point of view what type of control structures can be used to control such generic infrastructures, and explain what in particular makes *intelligent* infrastructures intelligent. We hereby especially focus on the conceptual ideas of model predictive control, both in centralized, single-agent control structures, and in distributed, multi-agent control structures. The need for more intelligence in infrastructures is then illustrated for three types of infrastructures: power, road, and water infrastructures.

R.R. Negenborn, H. Hellendoorn
Delft University of Technology, Delft Center for Systems and Control, Delft, The Netherlands,
e-mail: {r.r.negenborn, j.hellendoorn}@tudelft.nl

1.1 Transportation infrastructures

Transportation infrastructures, like power distribution networks [27], traffic and transportation systems [12], water distribution networks [7], logistic operations networks [28], etc., (see Figures 1.1–1.3) are the corner stones of our modern society. A smooth, efficient, reliable, and safe operation of these systems is of huge importance for the economic growth, the environment, and the quality of life, not only when the systems are pressed to the limits of their performance, but also under regular operating conditions. Recent examples illustrate this. E.g., the problems in the USA and Canada [44], Italy [42], Denmark and Sweden [16], The Netherlands, Germany, Belgium, and France [43], and many other countries [36] due to power outages have shown that as power network operation gets closer to its limits, small disturbances in heavily loaded lines can lead to large black-outs causing not only huge economic losses, but also mobility problems as trains and metros may not be able to operate. Also, as road traffic operation gets closer to its limits, unexpected situations in road traffic networks can lead to heavy congestion. Not only the huge traffic congestion after incidents such as bomb alerts are examples of this, also the almost daily road-traffic jams due to accidents illustrate this convincingly.

Expanding the physical infrastructure of these networks could help to relieve the issues in transportation networks, although at extremely high costs. As alternative to spending this money on building new infrastructure, it is worth spending effort on investigating improved use of the existing infrastructure by employing intelligent control techniques that combine state-of-the-art techniques from fields like systems and control engineering [4], optimization [6], and multi-agent systems [47], with domain-specific knowledge.

The examples of networks just mentioned are only some particular types of networks within the much larger class of transportation networks. Common to transportation networks is that at a generic level they can be seen as a set of nodes, representing the components or elements of the network, and interconnections between these nodes. In addition, transportation networks have some sort of commodity, that is brought into the network at source nodes, that flows over links to sink nodes, and that is influenced in its way of flowing over the network by elements inside the network, as illustrated in Figure 1.4. Other characteristics that are common to transportation networks are:

- they typically span a large geographical area;
- they have a modular structure consisting of many subsystems;
- they have many actuators and sensors;
- they have dynamics evolving over different time scales.

In addition to this, transportation networks often contain both continuous (e.g., flow evolution) and discrete dynamics (e.g., on and off switching), and can therefore also be referred to as hybrid systems [45]. This mixture of characteristics makes that transportation networks can show extremely complex dynamics.

Figure 1.1: The water network of The Netherlands.

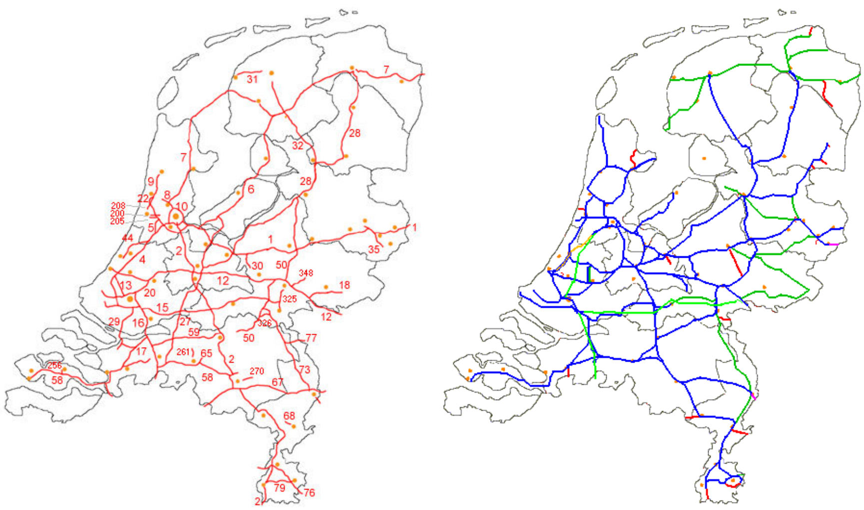

(a) (b)

Figure 1.2: (a) The national road network of The Netherlands; (b) The national train network of The Netherlands.

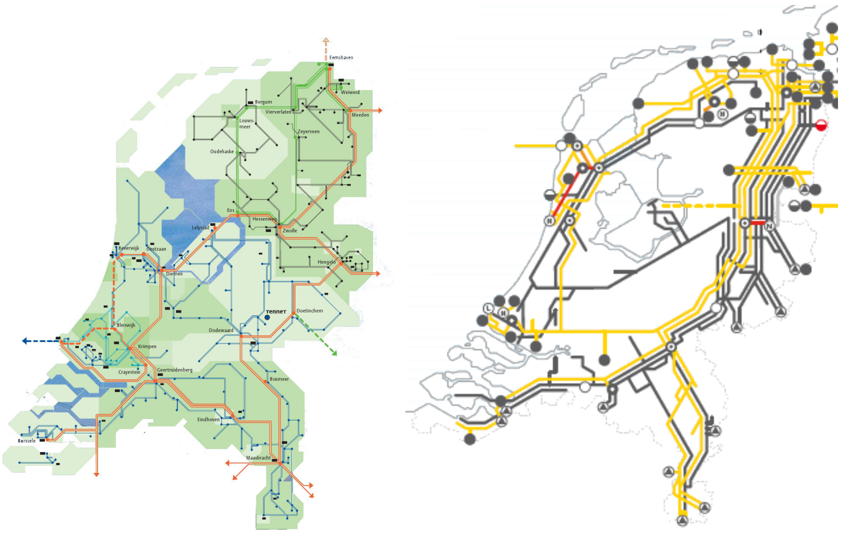

(a) (b)

Figure 1.3: (a) The national electricity grid of The Netherlands (Illustration courtesy of TenneT); (b) The national gas network of The Netherlands (Illustration courtesy of Gasunie).

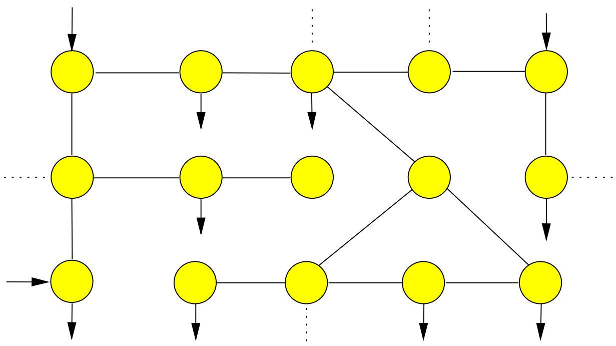

Figure 1.4: Networks of a generic transportation infrastructure. Commodity enters the network at sources (circles with an arrow pointing towards them), flows over links to other elements in the network that alter the flows (at each circle), and leaves the network at sinks (circles with an arrow pointing outward). Dotted lines represent connections with other parts of the network.

All the networks described in this book generically consist of sinks, sources, transition nodes, and interconnections. The commodity that is being transported over the network, however, determines how these components function. In a water network, the water always flows downwards unless a pump is deployed, when two waterways come together in a node the flow will increase, there may be intermediary sinks, ultimately there is one sink, the sea. The source may be distant mountains or rainfall, the amount of water can fairly precisely be predicted. In a traffic network, the sources in the morning are residential areas, the sinks are business districts, in the afternoon sources and sinks change roles. When two roads come together in a node the traffic flow will slow down or stagnate. Traffic lights will give the network a start-and-stop character. Highway and city traffic are very different in character. Electricity networks are determined by Kirchhoff's circuit laws that state that in any node in an electrical circuit the sum of currents flowing into a node is equal to the sum of the currents flowing out of that node. Electricity will always flow and can hardly be stored or stopped, sinks are industries and families, sources are power plants and increasingly often wind turbine parks and solar panels. Gas networks are determined by pressure with usually one source and many sinks. In railway networks it is hard to overtake, which gives these networks a strong sequential character, a chain which is very dependent on the weakest links.

Even though transportation networks differ in the details of commodity, sources, sinks, etc., it is worth to consider them in a generic setting. On the one hand, methods developed for generic transportation networks can be applied to a wide range of specific domains, perhaps using additional fine-tuning and domain-specific enhancements to improve the performance. On the other hand, approaches specifically developed for a particular domain can be applied to other domains after having transferred them to the generic framework.

1.2 Towards intelligent transportation infrastructures

There are many users, controllers, players, actors, and operators involved in the evolution of transportation networks. Each of these concepts refers to entities that directly or indirectly change the way commodity is flowing. Different users may have different objectives, and these objectives may be conflicting. Depending on their objectives, the users choose different actions, resulting in a different operation of the network.

1.2.1 The system versus its control structure

In order to formalize the operation of transportation networks, consider Figure 1.5. The figure illustrates the overall picture of a *system* on the one hand and a *control structure* on the other. The system is the entity that is under control, and the control structure is the entity that controls the system. Hence, the control structure is the concept used to indicate the structure that produces actuator settings. The control

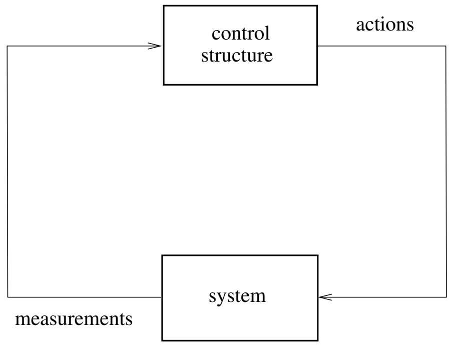

Figure 1.5: The relation between a general system and the control structure that controls the system.

structure monitors the system by making measurements and based on these chooses control actions that are implemented on the system. The system evolves subject to these actions to a new state, which is again measured by the control structure. The control structure consists of one or more components, called *control agents*. These control agents try to determine settings for the actuators inside the system in such a way that their own objectives are met as closely as possible and any constraints are satisfied. In our case, the system consists of the transportation network, and the components of the control structure consist of all the users, controllers, operators, players, etc., from now on only referred to as the control agents.

1.2.2 Intelligent infrastructures

Traditionally, operators play an important role in infrastructural networks. They monitor the system in large local or regional decision rooms, they open valves or pumps, they start up generators or shift transformer taps, and they can remotely close road lanes. More and more operators are replaced by ICT systems. In road traffic and railway control this development is well developed, in waste and drinking water networks this development has hardly begun. ICT solutions are in the beginning focused on local optimization and do not regard the whole network. In the case of, e.g., traffic lights this is not so much a problem, but in railway or electricity networks local optimization is of minor importance.

In the beginning operators observe the ICT solutions in decision rooms, but when confidence in the ICT system increases local control actions will be functioning fully automatically. Simultaneously with the development of locally more optimal control, it becomes apparent that coordination with other parts of the network becomes necessary. Green waves in road traffic networks are simple examples of this that show well how far more complicated this larger system is in a city road network with many junctions and crossing lanes. Communication between local controllers is necessary and intelligent systems have to be built that take into account all the flows and that have the capability of forecasting the near future. As confidence in these larger systems increases, coordination and negotiation between nodes will be a regular practice and transportation infrastructures have become *intelligent transportation infrastructures*.

1.2.3 Control structures

The control structure is a very general concept and can have many different shapes. A first important distinguishing feature between control structures is the number of control agents that constitute the control structure. E.g., the control structure can consist of a single control agent or multiple control agents. Some other properties in which control structures can differ are:

- the access that the control agents have to the sensors and actuators,
- the communication that the control agents have among one another,
- the way in which the control agents process sensor data to obtain actions,
- the authority relations between the control agents,
- the beliefs, desires, and intentions of the control agents.

Defining different types of control structures is difficult due to the large amount of properties that they can have. However, some general types of control structures can be identified, that have increasing complexity, that are commonly encountered in theory and practice, and that will also be of particular interest in the subsequent chapters of this book:

- When it is assumed that there is only one control agent, that has access to all actuators and sensors of the network and thus directly controls the physical network, then this control structure is referred to as an *single-agent* control structure, as illustrated in Figure 1.6(a). The control structure is in a sense ideal structure, since in principle such a control structure can determine actions that give optimal performance, although this may be at significant computational costs.
- When there are multiple control agents, each of them considering only its own part of the network and being able to access only sensors and actuators in that particular part of the network, then the control structure is referred to as a *multi-agent single-layer* control structure, as illustrated in Figure 1.6(b). If in addition the agents in the control structure do not communicate with each other, the control structure is *decentralized*. If the agents do communicate with each other, the control structure is *distributed*.
- When there are multiple control agents, and some of these control agents have authority over other control agents, in the sense that they can force or direct other control agents, then the control structure is a *multi-layer* control structure, as illustrated in Figure 1.6(c) A multi-layer control structure typically is present

when one control agent determines set-points to a group of other control agents, that work in a decentralized or distributed way. Due to the authority relationship between agents or groups of agents, the multi-layer control structure can also be referred to as a supervisory control structure, or a hierarchical control structure.

1.2.4 Control structure design

Suppose that a particular network is given and that any control structure can be implemented on it. The question that then arises is the question of how it can be determined what the best control structure is. Unfortunately, theories for determining general control structures are lacking. However, motivations for preferring one type of control structure over another can be given.

Advantages of single-agent control structures are in general that they can deliver the best performance possible, and that they have been studied extensively in the literature, in particular for small-scale systems. However, there are several issues that complicate the use of single-agent control structures for large-scale transportation networks such as:

- undesirable properties with respect to robustness, reliability, scalability, and responsiveness;
- technical issues related to communication delays and computational requirements;
- commercial, legal, and political issues related to unavailability of information and restricted control access.

These reasons motivate the use of multi-agent control structures [40, 41, 47], which are expected to be able to deal or at least relieve these issues. Multi-agent control structures can in principle:

- improve robustness and reliability, since if one control agents fails, another can take over, and improve responsiveness, since the control agents typically use only local measurements and therefore can react quicker to changing situations;
- reduce communication delays, since the control agents operate locally and therefore solve problems that may be smaller, and since communication typically takes place among nearby control agents;
- deal with unavailability of information and restricted control access, since the control agents only require information of their own part of the network and since they determine actions only for their own part of the network.

However, multi-agent control structures typically have a lower performance than the performance of ideal single-agent control structures; implementing schemes that give desired performance is far from trivial.

An advantage of the decentralized over the distributed multi-agent single-layer control structures is that there is no communication between the controllers, resulting in lower computational requirements and faster control. However, this advantage will typically be at the price of decreased overall performance. The advantage

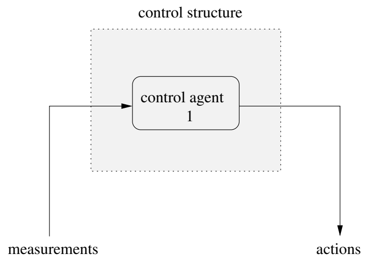

(a) Single-agent control structure. The single control agent makes measurements of the system and provides actions to the network.

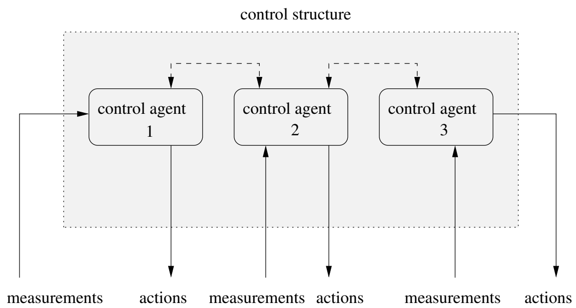

(b) Multi-agent single-layer control structure. Multiple control agents make measurements and provide actions to the network. Communication between the control agents is optionally present (dashed line).

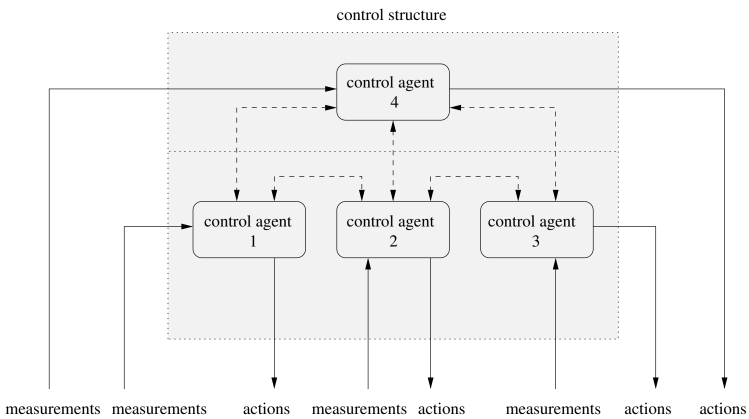

(c) Multi-layer control structure. A higher-layer control agent can make measurements and provide actions to the network and can in addition direct or steer a lower control layer.

Figure 1.6: Some important types of control structures.

of a distributed multi-agent single-layer control structure is therefore that improved performance can be obtained, although at the price of increased computation time due to cooperation, communication, and perhaps negotiation among control agents. However, even though improved performance can be obtained, the performance will still typically be lower than the performance of an ideal single-agent control structure.

The multi-agent multi-layer control structure provides the possibility to obtain a trade-off between system performance and computational complexity. A higher layer considers a larger part of the system and can therefore direct the lower control layer to obtain coordination. Such a multi-layer control structure can thus combine the advantages of the single-agent control structure with the multi-agent single-layer control structure, i.e., overall system performance with tractability. It is noted, however, that communication in a multi-agent multi-layer control structure is typically more complex than in a single-agent control structure and a multi-agent single-layer control structure.

Note that in practice often a particular control structure is already in place, and that the control structure cannot be redesigned from scratch. The question in this case is not so much the question of what control structure is best, but of how the currently existing control structure can be changed, such that the performance is improved. Of course, here it has to be defined what the performance is, and in a control structure with control agents with conflicting objectives it may not be possible to reach consensus on this.

1.2.5 Assumptions for design and analysis

In this book control strategies for several control structures are developed. Due to the complexity of transportation networks, the scope of control problems that is considered is narrowed down. The focus will mostly be on the most fundamental of transportation network control problems: the operational control of transportation networks, in which amounts of commodity to be transported over the network are given, and controllers have to ensure that transport over the network can take place at acceptable service levels, while satisfying any constraints, both under normal and emergency operating conditions.

In order to make the analysis and the design of the control structures more tractable, assumptions have to be made, both on the network and the control structure. Assumptions relating to the network are made on the dynamics of the network, i.e., the way in which the components in the network function. E.g., the dynamics can be assumed to evolve over continuous time or in discrete-time, they can be assumed to involve only continuous dynamics, or both continuous and discrete dynamics, and they can be assumed to be instantaneous or not. In each chapter we explicitly point out which particular assumptions are made on the network. With respect to the control structure, assumptions are made on the following:

- the control agents are already present (however, it is not known yet how they should behave);

- the control agents control fixed parts of the network, and they can access actuators and sensors in these parts of the network;
- the control agents know what qualitative behavior is desired for the parts of the network they control;
- the control agents strive for the best possible overall performance of the network;
- the control agents can measure the state of the parts of the network that they control.

Under such assumptions it remains to be decided how the agents in the control structure derive actuator settings from their measurements, i.e., what protocols, computations, and information exchanges take place inside the control structure. Assumptions on these are made in other chapters of this book. In the following section we discuss a promising approach for use by the control agents in a multi-agent control structure for transportation network control: model predictive control.

1.3 Model predictive control

To find the actions that meet the control objectives as well as possible, the control agents have to make a trade-off between the different available actions. In order to make the best decision and hence find the best actions, all relevant information about the consequences of choosing actions should be taken into account. For power networks, typical information that is available consists of forecasts on power consumption and exchanges [20], capacity limits on transmission lines, dynamics of components like generators, capacitor banks, transformers, and loads [27]. Furthermore, typically area-wide measurements of voltage magnitude and angles across the network can be made to provide an up-to-date status of the situation of the network. A particularly useful form of control for transportation network that in principle can use all information available is model predictive control (MPC) [10, 29].

1.3.1 Single-agent MPC

Over the last decades MPC (also knowns as receding horizon control or moving horizon control) has become an important methodology for finding control policies for complex, dynamic systems. MPC in a single-agent control structure has shown successful application in the process industry [29, 31], and is now gaining increasing attention in fields like, amongst others, power networks [18, 35], road traffic networks [21], railway networks [13], steam networks [30], multi-carrier systems [3], greenhouse control [38], and drug delivery [8].

1.3.1.1 Concept

MPC is a control methodology that is typically used in a discrete-time control context, i.e., control actions are determined in discrete control cycles of a particular

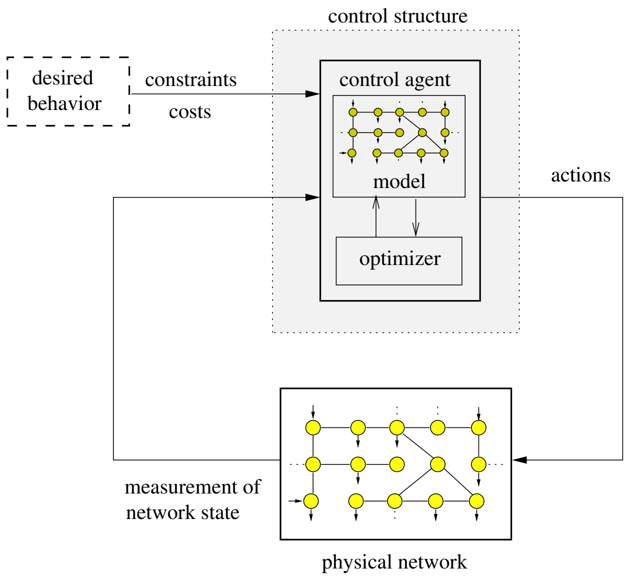

Figure 1.7: Single-agent MPC.

duration which in itself is expressed in continuous time units. From the beginning of one control cycle until the beginning of the next control cycle, the control actions stay fixed, i.e., a zero-order hold strategy is employed.

In each control cycle the MPC control agent uses the following information, as illustrated in Figure 1.7:

- an *objective function* expressing which system behavior and actions are desired;
- a *prediction model* describing the behavior of the system subject to actions;
- possibly *constraints* on the states, the inputs, and the outputs of the system (where the inputs and the outputs of the system correspond to the actions and the measurements of the control agent, respectively);
- possibly known information about future disturbances;
- a *measurement* of the state of the system at beginning of the current control cycle.

The objective of the control agent is to determine those actions that optimize the behavior of the system and minimize costs as specified through the objective function. In order to find the actions that lead to the best performance, the control agent uses the prediction model to predict the behavior of the system under various actions over a certain prediction horizon, starting from the state at the beginning of the control cycle. Once the control agent has determined the actions that optimize the system performance over the prediction horizon, it implements the actions until the beginning of the next control cycle, at which point the control agent determines new actions over the prediction horizon starting at that point, using updated information. Hence, the control agent operates in a receding or rolling horizon fashion to determine its actions.

In general it is preferable to have a longer prediction horizon, since by considering a longer prediction horizon, the control agent can better oversee the consequences of its actions. At some length, however, increasing the length of the prediction horizon may not improve the performance, if transients in the dynamics may have become negligible. For computational reasons, determining the actions over a very long horizon typically is not tractable, and in addition due to potential uncertainty in the prediction model and in predictions of future disturbances, a smaller prediction horizon is usually considered. Hence, in practice, the prediction horizon should be long enough to cover the most important dynamics, i.e., those dynamics dominating the performance, and short enough to give tractable computations. It should hereby also be noted that if a prediction horizon is used that is too short, the system could arrive in states from which it cannot continue due to the presence of constraints, e.g., on the actions. The prediction horizon should thus have such a length that arriving in such states can be avoided.

1.3.1.2 MPC Algorithm

Summarizing, a control agent in a single-agent control structure using MPC to determine its actions performs at each control cycle the following:

1. Measure the current state of the system.
2. Determine which actions optimize the performance over the prediction horizon by solving the following optimization problem:

minimize	the objective function in terms of actions over the prediction horizon
subject to	the dynamics of the whole network over the prediction horizon,
	the constraints on, e.g., ranges of actuator inputs and link capacities,
	the measurement of the initial state of the network at the beginning of the current control cycle.

3. Implement the actions until the next control cycle, and return to step 1.

1.3.1.3 Advantages and issues

Advantages of MPC are that in principle it can take into account all available information and that it can therefore anticipate undesirable situations in the future at an early stage. Additional advantages of MPC are [29]:

- its explicit way of handling constraints on actions, states, and outputs;
- its ability to operate without intervention for long periods;
- its ability to adapt to slow changes in the system parameters;
- its ability to control systems with multiple inputs and multiple outputs;
- its relatively easy tuning procedure;

- its built-in robustness properties.

However, there are also some issues that have to be addressed before a control agent using an MPC methodology can be implemented successfully:

- the control goals have to be specified;
- the prediction model has to be constructed;
- the measurement of the system state has to be available;
- a solution approach (optimization method) has to be available that can solve the MPC optimization problem;
- the solution approach has to be tractable and efficient.

Basic issues, e.g., stability and robustness, have extensively been studied for MPC in single-agent control structures [31], in particular for linear time-invariant systems. For other classes of systems there are still many open issues. E.g., tractability issues of MPC for nonlinear and discrete-event systems, and for systems in which variables take on discrete values, still deserve attention. E.g., in [34] an approach is proposed to make the MPC problem for a system modeled as a Markov decision process more tractable and to deal with changing system dynamics by including experience using reinforcement learning. Another class of systems for which there are still many open questions are hybrid systems, i.e., systems including both continuous and discrete dynamics.

1.3.2 Multi-agent MPC

As mentioned in the previous section, in a multi-agent control structure, there are multiple control agents, each of them controlling only its own subnetwork, i.e., a part of the overall network. Multi-agent MPC issues have been investigated since the 90s, such as in [1, 2, 5, 15, 17, 19, 22, 26, 37, 39], and more recently in [9, 11, 14, 24, 25, 33, 46] and [32].

In multi-agent MPC, multiple control agents in the control structure use MPC, but now they first measure the subnetwork state, then they determine the best actions over the predicted subnetwork evolution, and then they implement actions. Although this may seem like a straightforward extension of single-agent MPC at first sight, when considering the details it is not.

The actions that an agent in a multi-agent control structure takes influence both the evolution of the subnetwork it controls, and the evolution of the subnetworks connected to its subnetwork. Since the agents in a multi-agent control structure usually have no global overview and can only access a relatively small number of sensors and actuators, predicting the evolution of a subnetwork over a horizon involves even more uncertainty than when a single agent is employed. In addition, when a control agent in a multi-layer control structure provides set-points to another agent, this supervisory control changes the way in which the other agent chooses its actions, and thus the higher-layer control agent changes the performance of the system.

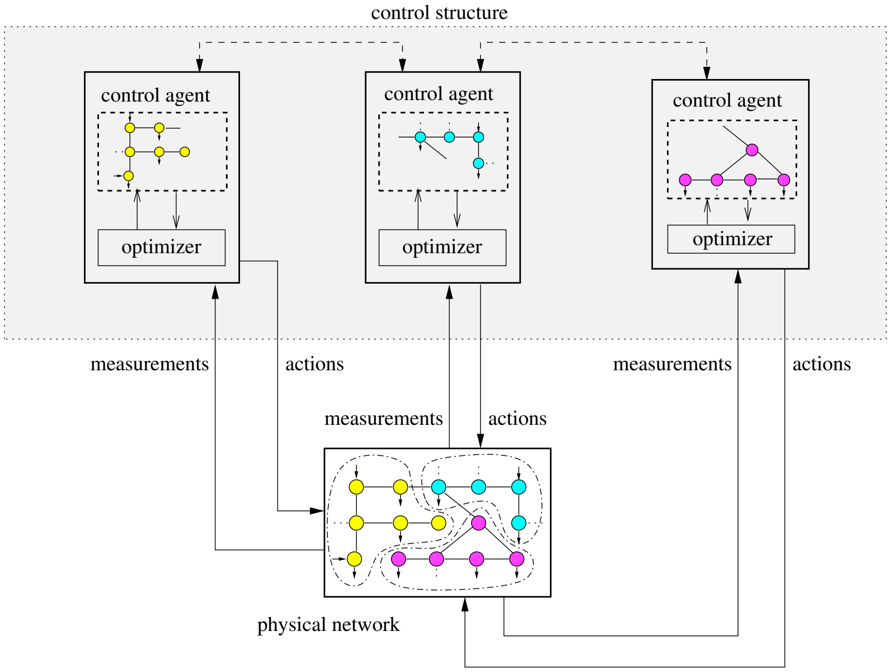

Figure 1.8: Multi-agent single-layer MPC.

Under the assumption that the control agents strive for an optimal overall network performance, the challenge in implementing such a multi-agent MPC methodology comes from ensuring that the actions that the individual agents choose result in a performance that is comparable to the situation when a hypothetical single-agent control structure in which all information is available would be used.

1.3.2.1 Multi-agent single-layer MPC

In the multi-agent single-layer control structure each control agent only has information gathering and action capabilities that are restricted to that part of the network that a particular control agent controls, as illustrated in Figure 1.8. The challenge in implementing multi-agent single-layer MPC comes from predicting the dynamics of the subnetwork, since, as mentioned, its evolution is influenced by the other agents. The underlying problem of MPC for multi-agent control structures can therefore be seen as optimization over a distributed simulation.

Issues

To make accurate predictions of the evolution of the subnetwork, a control agent requires the current state of its subnetwork, a sequence of actions over the prediction horizon, and predictions of the evolution of the interconnections with other subnetworks. The predictions of the evolution of the interconnections with other subnetworks are based on the information communicated with the neighboring control agents. One particular class of methods aims at achieving cooperation among control agents in an iterative way in which in each control cycle control agents perform several iterations consisting of local problem solving and communication. In each iteration agents obtain information about what the plans of neighboring agents are. Ideally at the end of the iterations the agents have found actions that lead to overall optimal performance.

As is the case with MPC for single-agent control structures, having both continuous and discrete dynamics causes computational problems. In transportation networks this combination is commonly encountered, and it is therefore relevant to study models that take this into account. A further complicating issue arises when the subnetworks that the agents control are overlapping. Existing strategies assume that the subnetworks that the control agents control are non-overlapping. However, in some applications the subnetworks considered by the control agents are overlapping. This has to be taken into account explicitly.

1.3.2.2 Multi-agent multi-layer MPC

In the multi-layer multi-agent MPC case there are multiple control layers in the control structure, i.e., there are authority relationships between the agents in the sense that some agents provide set-points or directions to other agents. The agents at higher layers typically consider a larger region of the network and consider slower time scales than agents in lower layers. Figure 1.9 illustrates this.

MPC can also be used by a control agent in a higher layer of the control structure. This higher-layer control agent can then coordinate the lower layer, which may consist of control agents using multi-agent single-layer MPC, or of control agents that use alternative control strategies. The higher-layer control agent then coordinates the lower control layer by enforcing penalty terms, providing additional constraints, or providing set-points. The advantage of the higher-layer control agent is in particular clear when the control agents of the lower layer are working decentralized, i.e., not communicating with one another.

Issues

An important issue to be addressed when designing MPC for multi-agent multi-layer control structures is the choice of the prediction model that the higher-layer control agent uses. A higher-layer control agent has to be able to make relevant predictions of the physical system, but since the physical system is under control of the lower-control layer, the lower-control layer has to be taken into account by the higher-layer control agent as well. In addition, the prediction model that the higher-layer control

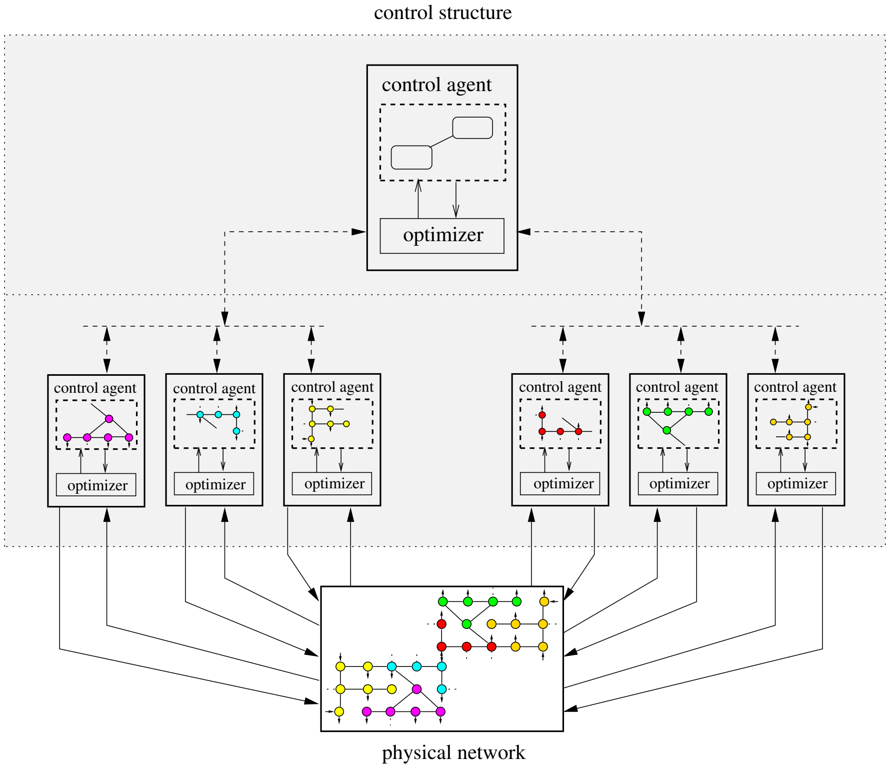

Figure 1.9: Multi-agent multi-layer MPC.

agent uses will typically involve both continuous and discrete elements, since it has to consider a larger part of the network than lower-layer agents. This makes the resulting MPC control problem more complex, and efficient ways have to be found to solve it efficiently.

1.4 MPC for intelligent infrastructures

As stated before, the commodity to be transported in a network determines the meaning of nodes and interconnections. Commodities are electrons in power networks, cars in road networks, water molecules in water networks, trains in railway networks, gas molecules in gas networks, etc. It is important to understand the origin and destination of commodities in networks and the operations that can be performed on them. Recent technological developments, in particular in communication technology, have increased the number of possible operations in various networks considerably. Below, intelligent power, road, and water networks are described as they are the main examples of intelligent infrastructures in this book.

1.4.1 Intelligent power infrastructures

Power networks [27] are large transportation networks consisting of a large number of components. The generators produce power that is injected into the network on the one side, while the loads consume power from the network on the other side. The distribution of the power in the network is dictated by Kirchhoff's laws and influenced by the settings of the generators, loads, transformers, and potentially also by capacitor banks and FACTS devices.

Conventionally, in power networks, power was generated in several large power generators. This power was then transported through the transmission and distribution network to the location where it was consumed, e.g., households and industry. The number of control agents was relatively low. Due to the ongoing deregulation in the power generation and distribution sector in the U.S. and Europe, the number of players involved in the generation and distribution of power has increased significantly. In the near future the number of source nodes of the power distribution network will even further increase as also large-scale industrial suppliers and small-scale individual households will start to feed electricity into the network [23].

As a consequence, the structure of the power distribution network is changing from a hierarchical top-down structure into a much more decentralized system with many generating sources and distributing agencies. This multi-player structure thus results in a system with many interactions and interdependencies. To still guarantee basic requirements and service levels, such as voltage levels, frequency, bounds on deviations, stability, elimination of transients, etc., and to meet the demands and requirements of the users, new control techniques are being developed and implemented resulting in an intelligent power infrastructure. The chapters in Part II of this book focus on this.

1.4.2 Intelligent road infrastructures

Electric traffic lights have now existed for almost a hundred years and are the best well-known control measures in road traffic networks. In the beginning, traffic lights were fixed-time stand-alone systems, later on they evolved into intelligent control systems using sensor information and communication with neighboring junctions. In the last decades a large number of traffic control measures have been added: dynamic speed signs, dynamic route information panels, ramp metering, and parking guidance systems. Furthermore, due to the increased possibilities of traffic measuring traffic information has become available via radio, internet, and route navigation systems. At the same time the amount of traffic has increased enormously, congested inner cities and traffic jams have become a regular phenomenon that cannot be solved by single operating traffic measures anymore. In traffic management the combination of information, prohibition, limitation, and coordination is apparently needed. Model-based predictive control does already play a large role in road networks, but the need for more intelligence and cooperation is huge. Future developments like intelligent (partly) autonomous vehicles in combination with roadside control will open up new possibilities for traffic management and control. The chapters of Part III of this book focus on this.

1.4.3 Intelligent water infrastructures

In the near future the importance of an efficient and reliable flood and water management system will keep on increasing, among others due to the effects of global warming (higher sea levels, more heavy rain during the spring season, but possibly also drier summers). Due to the large scale of water networks, control of such networks in general cannot be done in a centralized way, in which from a single location measurements from the whole system are collected and actions for the whole system are determined. Instead, control is typically decentralized over several local control bodies, each controlling a particular part of the network [41, 47]. Local control actions include activation of pumps or locks, filling or draining of water reservoirs, or controlled flooding of water meadows or of emergency water storage areas. MPC and distributed MPC are suitable techniques for intelligently determining how these local control actions should be implemented. The chapters of Part IV of this book focus on this.

1.5 Conclusions and future research

In this chapter we have discussed various infrastructures and have illustrated that they are similar in structure with regard to their nodes and interconnections, which is in particular visualized in Figures 1.1–1.3, but very different in character due to the very different commodities that are transported over the network. It was made apparent that transporting gas or water molecules is very different from transporting cars and trains, and that this is again different from transporting power. Some commodities are passive, like gas and water, others are active, in particular cars, where drivers play an important role and continuously interact with the control measures. Furthermore, we have discussed several control constellations: single-agent and multi-agent control structures, single-layer and multi-layer control structures, hierarchical and distributed systems, centralized and decentralized systems, and combinations of all of these. We have discussed the pros and cons of each structure and the role of computational power and optimality in each of them. The status quo, the vulnerability of the network, and the technical possibilities determine which control structure is most appropriate to meet the challenges in the networks and the user demands.

We have introduced several kinds of Model Predictive Control (MPC) and explained why this methodology is useful for control of large networks. We have illustrated how MPC and multi-agent techniques can be combined, how this combination can tackle problems in subnetworks and how agents in subnetworks can cooperate and negotiate with each other to find an optimal solution for the whole network.

Finally, we have introduced in somewhat more detail recent technological and communications developments in power networks, road networks, and water networks.

Acknowledgements This research is supported by the BSIK project "Next Generation Infrastructures (NGI)" and the Delft Research Center Next Generation Infrastructures.

References

1. L. Acar. Some examples for the decentralized receding horizon control. In *Proceedings of the 31st IEEE Conference on Decision and Control*, pages 1356–1359, Tucson, Arizona, 1992.
2. M. Aicardi, G. Casalino, R. Minciardi, and R. Zoppoli. On the existence of stationary optimal receding-horizon strategies for dynamic teams with common past information structures. *IEEE Transactions on Automatic Control*, 37:1767–1771, November 1992.
3. M. Arnold, R. R. Negenborn, G. Andersson, and B. De Schutter. Model-based predictive control applied to multi-carrier systems. In *Proceedings of the IEEE PES General Meeting 2009*, Calgary, Canada, June 2009.
4. K. J. Åström and B. Wittenmark. *Computer-Controlled Systems*. Prentice-Hall, Upper Saddle River, New Jersey, 1997.
5. M. Baglietto, T. Parisini, and R. Zoppoli. Neural approximators and team theory for dynamic routing: A receding-horizon approach. In *Proceedings of the 38th IEEE Conference on Decision and Control*, pages 3283–3288, Phoenix, Arizona, 1999.
6. D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, Beltmore, Massachusetts, 2003.
7. P. R. Bhave and R. Gupta. *Analysis of Water Distribution Networks*. Alpha Science International, Oxford, UK, 2006.
8. L. G. Bleris, P. D. Vouzis, J. G. Garcia, M. G. Arnold, and M. V. Kothare. Pathways for optimization-based drug delivery. *Control Engineering Practice*, 15(10):1280–1291, October 2007.
9. M. W. Braun, D. E. Rivera, M. E. Flores, W. M. Carlyle, and K. G. Kempf. A model predictive control framework for robust management of multi-product, multi-echelon demand networks. *Annual Reviews in Control*, 27:229–245, 2003.
10. E. F. Camacho and C. Bordons. *Model Predictive Control in the Process Industry*. Springer-Verlag, Berlin, Germany, 1995.
11. E. Camponogara, D. Jia, B. H. Krogh, and S. Talukdar. Distributed model predictive control. *IEEE Control Systems Magazine*, 1:44–52, February 2002.
12. C. F. Daganzo. *Fundamentals of Transportation and Traffic Operations*. Pergamon Press, New York, New York, 1997.
13. B. De Schutter, T. van den Boom, and A. Hegyi. A model predictive control approach for recovery from delays in railway systems. *Transportation Research Record*, (1793):15–20, 2002.
14. W. B. Dunbar and R. M. Murray. Distributed receding horizon control for multi-vehicle formation stabilization. *Automatica*, 42(4):549–558, April 2006.
15. H. El Fawal, D. Georges, and G. Bornard. Optimal control of complex irrigation systems via decomposition-coordination and the use of augmented Lagrangian. In *Proceedings of the 1998 International Conference on Systems, Man, and Cybernetics*, pages 3874–3879, San Diego, California, 1998.

16. Elkraft Systems. Power failure in Eastern Denmark and Southern Sweden on 23 September 2003 – preliminary report on the course of events. Technical report, Elkraft Systems, Holte, Denmark, 2003.
17. G. Georges. Decentralized adaptive control for a water distribution system. In *Proceedings of the 3rd IEEE Conference on Control Applications*, pages 1411–1416, Glasgow, UK, 1994.
18. T. Geyer, M. Larsson, and M. Morari. Hybrid emergency voltage control in power systems. In *Proceedings of the European Control Conference 2003*, Cambridge, UK, September 2003. Paper 322.
19. M. Gomez, J. Rodellar, F. Vea, J. Mantecon, and J. Cardona. Decentralized predictive control of multi-reach canals. In *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*, pages 3885–3890, San Diego, California, 1998.
20. E. González-Romera, M. Á. Jaramillo-Morán, and D. Carmona-Fernández. Forecasting of the electric energy demand trend and monthly fluctuation with neural networks. *Computers and Industrial Engineering*, 52:336–343, April 2007.
21. A. Hegyi, B. De Schutter, and J. Hellendoorn. Optimal coordination of variable speed limits to suppress shock waves. *IEEE Transactions on Intelligent Transportation Systems*, 6(1):102–112, March 2005.
22. R. Irizarry-Rivera and W. D. Seider. Model-predictive control of the Czochralski crystallization process. Part I. Conduction-dominated melt. *Journal of Crystal Growth*, 178(4):593–611, 1997.
23. N. Jenkins, R. Allan, P. Crossley, D. Kirschen, and G. Strbac. *Embedded Generation*. TJ International, Padstow, UK, 2000.
24. D. Jia and B. Krogh. Min-max feedback model predictive control for distributed control with communication. In *Proceedings of the 2002 American Control Conference*, pages 4507–4512, Anchorage, Alaska, May 2002.
25. D. Jia and B. H. Krogh. Distributed model predictive control. In *Proceedings of the 2001 American Control Conference*, pages 2767–2772, Arlington, Virginia, June 2001.
26. M. R. Katebi and M. A. Johnson. Predictive control design for large-scale systems. *Automatica*, 33(3):421–425, 1997.
27. P. Kundur. *Power System Stability and Control*. McGraw-Hill, New York, New York, 1994.
28. G. Lodewijks. *Dynamics of Belt Systems*. PhD thesis, Delft University of Technology, The Netherlands, 1996.
29. J. M. Maciejowski. *Predictive Control with Constraints*. Prentice-Hall, Harlow, UK, 2002.
30. Y. Majanne. Model predictive pressure control of steam networks. *Control Engineering Practice*, 13(12):1499–1505, December 2005.
31. M. Morari and J. H. Lee. Model predictive control: past, present and future. *Computers and Chemical Engineering*, 23(4):667–682, 1999.
32. R. R. Negenborn. *Multi-Agent Model Predictive Control with Applications to Power Networks*. PhD thesis, Delft University of Technology, Delft, The Netherlands, December 2007.
33. R. R. Negenborn, B. De Schutter, and J. Hellendoorn. Multi-agent model predictive control for transportation networks: Serial versus parallel schemes. *Engineering Applications of Artificial Intelligence*, 21(3):353–366, April 2008.
34. R. R. Negenborn, B. De Schutter, M.A. Wiering, and H. Hellendoorn. Learning-based model predictive control for Markov decision processes. In *Proceedings of the 16th IFAC World Congress*, Prague, Czech Republic, July 2005. Paper 2106 / We-M16-TO/2.
35. R. R. Negenborn, M. Houwing, B. De Schutter, and J. Hel lendoorn. Model predictive control for residential energy resources usin g a mixed-logical dynamic model. In *Proceedings of the 2009 IEEE International Conference on Netwo rking, Sensing and Control*, pages 702–707, Okayama, Japan, March 2009.
36. P. G. Neumann. Widespread network failures. *Communications of the ACM*, 50(2):112, 2007.
37. S. Ochs, S. Engell, and A. Draeger. Decentralized vs. model predictive control of an industrial glass tube manufacturing process. In *Proceedings of the 1998 IEEE Conference on Control Applications*, pages 16–20, Trieste, Italy, 1998.

38. S. Piñón, E. F. Camacho, B. Kuchen, and M. Peña. Constrained predictive control of a greenhouse. *Computers and Electronics in Agriculture*, 49(3):317–329, December 2005.
39. S. Sawadogo, R. M. Faye, P. O. Malaterre, and F. Mora-Camino. Decentralized predictive controller for delivery canals. In *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics*, pages 3380–3884, San Diego, California, 1998.
40. K. P. Sycara. Multiagent systems. *AI Magazine*, 2(19):79–92, 1998.
41. D. D. Šiljak. *Decentralized Control of Complex Systems*. Academic Press, Boston, Massachusetts, 1991.
42. UCTE. Final report of the investigation committee on the 28 September 2003 blackout in Italy. Technical report, Union for the Coordination of Transmission of Electricity (UCTE), Brussels, Belgium, 2003.
43. UCTE. Final report system disturbance on 4 November 2006. Technical report, Union for the Coordination of Transmission of Electricity (UCTE), Brussels, Belgium, 2006.
44. U.S.-Canada Power System Outage Task Force. Final report on the August 14, 2003 blackout in the United States and Canada: causes and recommendations. Technical report, April 2004.
45. A. J. van der Schaft and J. M. Schumacher. *An Introduction to Hybrid Dynamical Systems*, volume 251 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag, London, 2000.
46. A. N. Venkat, I. A. Hiskens, J. B. Rawlings, and S. J. Wright. Distributed output feedback MPC for power system control. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, California, December 2006.
47. G. Weiss. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, Cambridge, Massachusetts, 2000.