

Technical report

Distributed model predictive control for irrigation canals

R.R. Negenborn, P.J. van Overloop, T. Keviczky, B. De Schutter

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

R.R. Negenborn, P.J. van Overloop, T. Keviczky, B. De Schutter. Distributed model predictive control for irrigation canals. *Networks and Heterogeneous Media*, vol. 4, no. 2, pp. 359-380, June 2009.

Delft University of Technology, Delft, The Netherlands

DISTRIBUTED MODEL PREDICTIVE CONTROL OF IRRIGATION CANALS

RUDY R. NEGENBORN

Delft Center for Systems and Control
Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

PETER-JULES VAN OVERLOOP

Department of Water Management
Delft University of Technology, Stevinweg 1, 2628 CN Delft, The Netherlands

TAMÁS KEVICZKY

Delft Center for Systems and Control
Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

BART DE SCHUTTER

Delft Center for Systems and Control & Department of Marine and Transport Technology
Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

ABSTRACT. Irrigation canals are large-scale systems, consisting of many interacting components, and spanning vast geographical areas. For safe and efficient operation of these canals, maintaining the levels of the water flows close to pre-specified reference values is crucial, both under normal operating conditions as well as in extreme situations.

Irrigation canals are equipped with local controllers, to control the flow of water by adjusting the settings of control structures such as gates and pumps. Traditionally, the local controllers operate in a decentralized way in the sense that they use local information only, that they are not explicitly aware of the presence of other controllers or subsystems, and that no communication among them takes place. Hence, an obvious drawback of such a decentralized control scheme is that adequate performance at a system-wide level may be jeopardized, due to the unexpected and unanticipated interactions among the subsystems and the actions of the local controllers.

In this paper we survey the state-of-the-art literature on distributed control of water systems in general, and irrigation canals in particular. We focus on the model predictive control (MPC) strategy, which is a model-based control strategy in which prediction models are used in an optimization to determine optimal control inputs over a given horizon. We discuss how communication among local MPC controllers can be included to improve the performance of the overall system. We present a distributed control scheme in which each controller employs MPC to determine those actions that maintain water levels after disturbances close to pre-specified reference values. Using the presented scheme the local controllers cooperatively strive for obtaining the best system-wide performance. A simulation study on an irrigation canal with seven reaches illustrates the potential of the approach.

2000 *Mathematics Subject Classification*. Primary: 68W15; Secondary: 93C99.

Key words and phrases. Model predictive control, distributed control, irrigation canals.

1. Introduction. Water is one of the most vital elements in human life. Among other things, it is used for drinking, agriculture, transportation, recreation, and energy production. Water systems consist of water bodies, such as lakes and reservoirs, connected by natural water courses (channels) and man-made canals. The water flows in the rivers and canals can be manipulated by structures such as pumps and gates.

Although water is vital for humans, people also have always needed to protect themselves from an excess of water coming from extreme precipitation, high river discharges, or high sea water levels. In the near future the importance of water management will increase due to the effects of global warming, such as higher sea levels, heavier rain during the spring season, and possibly also drier summers [16]. Efficient and reliable strategies for flood protection and prevention on the one hand, and irrigation and fair water distribution on the other hand, will have to be developed. In addition to this, world-wide objectives for water management, such as transportation, recreation, and energy production should also be considered.

In this paper we consider the management organization of water systems and we discuss how automatic distributed control schemes can be employed in the control of water systems. We hereby in particular focus on the role that distributed, model-based, predictive control can play for the control of water systems in general, and irrigation canals in particular.

1.1. Management organization of water systems. The wide variety of important roles that water plays in our lives has created a need for organizations and societies that manage the human interaction with water systems. This has resulted in a complex structure of responsibilities that is not governed by the behavior of the water systems, but by the existing societal or organizational structures. Two examples of societal and organizational divisions, that can be found in the water system domains of large rivers and canal networks, are the following:

- A division at a spatial level is apparent in the management of large rivers. These rivers often run through various countries. The management of the river in each country is an important national issue in which the inflows and outflows to the other counties are considered as a given boundary condition.
- A division by working field is apparent from the separate departments that manage a water system with their own isolated objectives. Water boards usually have one department that is responsible for the management of the water quantity variables and processes, such as water availability and flood protection, and another department responsible for water quality variables and processes, such as salinity control and water treatment. These variables and processes are all part of the same canal network and interact physically.

The spatial and working field division of water management is generally regarded as unwanted, but hard to change. Many studies have been done on trans-boundary water management of rivers and the potential of integrated water management of water quantity and quality for canal systems. These studies have resulted in the formation of international agreements on river inflows and outflows at a national level and agreements on target levels for water quantity and water quality variables that the different departments can use. The agreements are usually updated once every 2–3 years, but it is evident that the dynamic behavior of water systems varies at a much faster timescale and that therefore water management could benefit from coordination at a much higher frequency, e.g., daily or even hourly.

So far two types of controlled water systems have already been described, rivers managed by the national water board and a canal network managed by a regional water board. Other water systems that are characterized by an extensive water management organization with specific objectives and interactions with neighboring water systems are listed in the following:

Sewer systems: These systems transport waste water from houses and excess rain water, through closed conduits, to treatment plants and spillways. The system is managed by a municipal authority. The main objectives are to avoid spilling of waste water into surrounding open water systems and onto the streets. Agreements are usually made with the management of the surrounding canals, normally a regional water board, about the maximum number of spills and the maximum spill volume over a year.

Large multi-purpose reservoirs: Dams are constructed in mountainous areas to backup large amounts of water in reservoirs. The storage capacity in the reservoirs can be used for flood protection of the downstream areas, agricultural irrigation during the whole year and energy production by water power generators. Often, multiple reservoirs are constructed in series. The management objective of one reservoir is already complex due to the multi-purpose usage of water in the reservoir. This often results in neglecting coordination with other reservoirs.

Irrigation canals: Irrigation canals transport water over long distances from a source to water users at remote locations. A canal is managed by an irrigation district. The main objective is to deliver the right amount of water at the right location at the right time. This is achieved by continuously striving for a more or less constant available amount of water in the canals (water level control). In case of drought, the irrigation district needs to negotiate with the neighboring districts and the managers of the water sources to maximize the irrigation benefits in their district as much as possible.

From the above descriptions it is clear that the organizations managing these systems have their own objectives and are less interested in the total availability of water or consequences for downstream water users in terms of water quantity and water quality. Obviously, there are opportunities for automatic control applications utilizing systems theory to coordinate the overall behavior consisting of multiple, more locally oriented, management objectives. At all scales there is a need to develop knowledge on negotiation algorithms and on information content that has to be exchanged between distributed systems, especially as water of good quality is becoming more scarce in certain parts of the world, while other parts face a threat of water excess. In the next section we describe two automatic control philosophies that are aimed at distributing the decision making process involved in controlling large-scale dynamical systems such as water networks.

1.2. Automatic control of water systems. Various automatic control techniques have been used for control of water systems over the years, see [22, 23, 25, 31] for extensive overviews of techniques used, actuators, sensors, and application systems. The first attempts to implement control were based on feed-forward control, i.e., based on choosing actions in a pre-defined way using measured disturbances only. The reason for this was that accurate models were available that could be used for inverse modeling and no knowledge on feedback control (i.e., choosing which action to take based on measurements of the actual output of the system) was available

among civil engineers at that time. The first successful implementations however were based on feedback control [33, 23] either without or with feed-forward control added. These controllers were able to keep water levels close to set-points and in this way ensured availability of water in the canal to be delivered to users.

1.2.1. *Decentralized and distributed control.* Due to the complexity and size of water networks, control of such systems in general is not done in a centralized way. Controlling such systems in a centralized way would have a too large computational burden. Instead of collecting measurements from the whole system and determining actions in a single location, control is typically spread over several local control bodies, each controlling a particular part of the network [27, 35, 41]. Local control actions include activation of pumps, filling or emptying of water reservoirs, or controlled flooding of water meadows or of emergency water storage areas. Due to the continuing developments in information and communications technology, exchange of information between local controllers becomes practically and economically possible, such that the local controllers have the possibility to take one another's actions into account. We make a distinction between decentralized and distributed control architectures based on the available information, control authority, and interactions between the local controllers as follows:

Decentralized control: In decentralized control settings, controllers are developed for local control, without taking into account the effects that local actions have on the overall system performance, and without taking into account potential cooperation, negotiation and communication with other controllers.

Distributed control: Local controllers may be designed in such a way that they take into account the effects of local actions at a system-wide level using information exchange. The local controllers are thus able to perform cooperation and negotiation with other controllers with the aim of achieving the best system-wide performance. This is called distributed control.

1.2.2. *Potential benefits.* When improvements are required in large-scale water systems consisting of multiple controlled subsystems, a new way of looking at the global control problem is required. Instead of simply combining the local control problems into one large control problem, one needs to consider the interactions between the subsystems and investigate distributed decision-making possibilities by relying on coordination between multiple interacting subsystems. The potential for coordination in various water systems is described in the following list:

Rivers: Calamity inundation areas are constructed along rivers. Each country has its own operating rules that are directly applied in the event that these measures have to come into operation. They base their operating rules on measurements in their own country. For downstream lying countries, such as The Netherlands, it can be vital that inundation areas in upstream countries, in this case, e.g., Germany, Belgium, and France, are actuated, even if there is no need to do so in the upstream country. In this way, a significant amount of damage can be avoided. The flood protection systems of multiple countries can be linked and negotiation algorithms that minimize total damage can be implemented.

Canal networks: Various regional water boards in The Netherlands discharge their surplus water to other water boards that transport this water to a river or sea. The control and decision-support systems make calculations using fixed values for maximum discharge flows that are recorded in agreements

between the water boards. Instead, there should be the flexibility to exceed temporarily these values if the embankments in the area of one water board are about to collapse and the other water board still has storage capacity available. The control systems can be linked creating a distributed decision-support environment.

Sewer systems: The total volume of water that a sewer system is allowed to spill into an open water system is based on values of the water quality variables in the sewer network averaged over multiple years. When actual measurements indicate much lower values or if the receiving open water canals are less sensitive to temporal pollution (e.g., in the winter period), the spill volume should be allowed to be higher temporarily. This functionality can be implemented with coordination between the decision support systems of the sewer system and the canal network.

Reservoirs: When more reservoirs are operated in one catchment area, there coordination among them can improve flood prevention. Depending on the spatial distribution of extreme precipitation, certain reservoirs can still have space for storing water, while others are at the verge of spilling towards a downstream village or city. Reservoirs can also be used in the production of energy. In fact, coordination in the production of energy using multiple reservoirs will become more and more important in the future, especially as the potential energy in water is a source that can be used in periods when other energy sources are insufficient. During the day extra energy can be produced, while during the night water is pumped to a higher location and stored there in reservoirs, as is already done in Belgium and Norway. Using distributed control the energy flows demanded between multiple reservoirs can be coordinated.

Irrigation canals: Cooperating control systems for irrigation districts that have inter-dependent water demand schedules can yield a better spreading of the available water towards areas that are under high stress. On a larger scale, less water will be spilled. An illustration of the importance of coordination between sub-systems is the avoidance of disturbance amplification in canals consisting of canal reaches in series. When the water level in separate canal reaches are controlled simultaneously with Proportional Integral (PI) controllers, that are tuned to give a high performance, problems could occur during operation. Disturbances that occur at the downstream side are amplified at each control gate further upstream. Coordination between the canal reaches is required or a global tuning procedure for all PI-controllers needs to be used that minimizes the deviations from the set-points in all reaches [37].

1.3. Model-based control. In water systems, physical constraints play an important role. In order to take physical constraints into account, so-called model predictive controllers (MPC) [20] can be used. Over the last decades MPC (also known as receding horizon control or moving horizon control) has become an important strategy for finding control policies for complex, dynamic systems. MPC for centralized control has shown successful applications in the process industry [20], and is now gaining increasing attention in fields like power networks [12], road traffic networks [14], and steam networks [21].

MPC is a promising form of control for making a trade-off among the various control actions available and for determining the actions that meet the control objectives and satisfy constraints of local control bodies as well as possible. In the

remainder of this paper we discuss how MPC could be used for distributed control of the water systems introduced above. In particular, we focus on the role that distributed MPC can play in the control of irrigation canals.

The remainder of the paper is organized as follows. In Section 2, the MPC methodology is introduced and the approaches of centralized MPC that have been proposed in the literature for control of irrigation canals are discussed. In Section 3, a distributed MPC framework is presented and related literature on distributed MPC in general and with particular application to irrigation canals is surveyed. In Section 4, we propose a particular distributed MPC scheme, well-suited for distributed control of water levels in water systems. A case-study using this control scheme to control an irrigation canal is provided in Section 5. Section 6 concludes the paper and contains directions for future research.

2. Centralized MPC.

2.1. Principle of operation. MPC is an online optimization-based control approach. The motivation for using such an approach arises from the following. In water systems, costs can be associated to actions (e.g., pumping, changing flows through gates) and states (e.g., water levels). This holds both in the case of exceedingly high water levels (which may result in floods) or very low groundwater or surface water levels (which have a negative impact on agriculture, irrigation, and drinking water supplies). Models can be constructed that describe how particular water systems behave. By making predictions over a certain prediction horizon using these models, a controller can determine which actions have to be chosen in order to obtain the best performance.

An MPC controller determines which action to take at discrete *control cycles*. At each control cycle the controller first obtains the current state of the system it controls. It then formulates and solves an optimization problem that determines the actions over the prediction horizon that give the best predicted performance. The controller implements these actions until the beginning of the next control cycle, at which time the MPC controller repeats these steps in a receding horizon fashion, i.e., by obtaining new information about the current state and by reformulating the optimization problem starting from the new control cycle.

2.2. Main characteristics of MPC for water systems. Several approaches of MPC have been proposed for control of water systems. The primary motivation for using MPC for control of water systems is that the dynamics of water systems typically involve long time lags. MPC can take these delays into account. In addition, constraints on inputs and states representing physical limitations can be considered. MPC also has the ability to deal with multiple inputs and multiple outputs, and the ability to include predictions on, e.g., future rainfall. MPC has the potential to improve system-wide performance by coordinating local decentralized controllers through set-point updates, or by replacing these local decentralized controllers altogether. Since the dynamics of the water systems are relatively slow, time is available for performing the computations involved in the optimization within MPC.

The MPC approaches proposed in the literature differ in the control objectives encoded in the performance criterion, the actuators or control measures considered, the type of prediction model used to represent the water system, the physical constraints considered, the approach with which the optimization problem is solved,

and the water systems on which experiments are performed. In such centralized control systems, the optimization problem is solved at a single central location using information from the whole system to determine actions for the entire water system. Below we list distinguishing factors between the various approaches proposed for centralized MPC:

Control objectives: The control objectives of the MPC controller are typically to regulate (downstream) water levels [2,26,13,4,3], to minimize costs on water supply, treatment, and elevation [5], to minimize costs on pressure regulation, flow regulation and water quality [5], to make sure that the right amount of water is at the right place [8,40], to ensure that operational spills are avoided [40], and to make sure that reservoirs are emptied as fast as possible [3]. Such control objectives are usually rephrased as objectives with respect to water levels staying close to a given reference set-point [24, 2, 30, 26, 40, 39, 4], input values being minimized [24], or changes in inputs being minimized [2, 30, 40, 39, 4].

Control measures: To achieve the control objectives, the available control measures have to be manipulated. Usually the control measures consist of the (upstream) discharges through gates [2, 5, 26, 13, 40], although in some cases the gate openings [4, 3], the water flow towards reservoirs [8], or pump flows [5] are controlled.

Prediction models: For making the predictions of the behavior of particular water systems a wide variety of prediction models has been considered in the literature, including an analytically obtained linear time-invariant state-space model derived from Saint Venant's equations discretized through the Preismann implicit scheme [24], a linear time-varying state-space model [3], polynomial models based Diophantine equations [2, 30], a nonlinear discrete-time state-space model [5], a black-box identified linear model transferred into linear state-space form [4], a linear model known as the Muskingum model in combination with a discretized version of a differential equation for representing storage [26, 13], the integrator delay model [40, 39], and artificial neural networks using radial basis functions [8].

Constraints: Various operational constraints may be present as illustrated in [5, 8, 39, 4, 3]. These operational constraints include constraints on lower and upper bounds on water levels, on the capacity of canals, and on the variations of water levels around operating points. In addition, there may be lower and upper bounds on control variables, gates may not be allowed to shut down all the way or to come out of the water, and there may be gate opening and opening rate limits.

Obtaining future exogenous inputs: Future rainfall and/or water offtakes have to be predicted when using MPC. In the literature, future discharges are either predicted [13], or assumed perfectly known [3]. In a test environment the future discharges can be perfectly known. However, in reality predictions will always be uncertain.

Solving the MPC problem: For water systems the MPC optimization problem is usually solved numerically using a nonlinear optimization problem solver, such as CONOPT in [5], quadratic programming in [4], or a combination of quadratic programming with an approach for nonlinear optimization in [3]. In some cases, when no explicit constraints are considered, the MPC optimization problem is solved analytically, such as in [2].

Water system: There are various types of water systems for which MPC has been employed, such as irrigation canals [24, 2, 30, 26, 40, 39, 4], river systems [3], water supply networks [8], and water distribution systems [5].

Test canals: The performance of the MPC approaches has been illustrated on various test canals, including the ASCE test canal 1 (first 2 canal reaches [24], the first 3 canal reaches [30], and the complete canal [39]), ASCE test canal 2 [2, 26, 13], a model consisting of large portions of the Arizona Canal, the Grand Canal, and the South Canal in Arizona [40], a model of the river basin Demer in Belgium [3], and a practical prototype lab setup [4].

Simulation packages: Models of the test canals have been implemented using different simulation packages, such as SOBEK [40, 39, 3], SIC [24, 30], and Simulink [26, 13].

Performance criteria: In order to evaluate the performance of the controllers, different performance criteria have been considered, such as the performance indices defined by the ASCE committee [24, 2] (including the maximum absolute error, the integrated absolute error, the steady-state error, and the change in integrated absolute discharge). Alternatively, controller performance is assessed by visual inspection [3].

2.3. Drawbacks of centralized MPC. The control problems for the water systems described in Section 1.2.2 (for rivers, canal networks, sewer systems, reservoirs and irrigation canals) can, with the present knowledge, be formalized in simplified models, performance criteria, and constraints. These control problems in principle can be solved with centralized MPC. The implementation of multi-objective controllers such as MPC could solve the issues related to the organizational division. At present, the integration of various objectives on different variables by means of MPC is being investigated, showing promising results. However, the limits of complexity have almost been reached. Solving larger control problems is not tractable, due to the computational limitations. At the same time we should notice that solving the control problems of local water systems, and having them negotiate on the interactions with neighboring systems (managed by other organizations), is in line with the present practice in large-scale water systems and probably the only feasible solution for automatic control solutions in the near future. In the next section we further elaborate on distributed, model-based automatic control solutions and we review available methodologies in the literature.

3. Distributed MPC.

3.1. Importance of distributed, model-based control. Although the local water management bodies usually only control or manage the water levels in a relatively small region, the evolution of the water levels is influenced by what happens over a much larger area, often extending far beyond the neighborhood of the given region (e.g., due to water arriving via rivers or via subsurface diffusion flows). This suggests that decisions made using purely local information are necessarily conservative and suboptimal on a system-wide level. Furthermore, the elements in a large-scale water system are all dynamic in nature with often significant time-delays. It is therefore thus appropriate to describe the evolution of these components and their future behavior using mathematical models. In addition, future events such as precipitation should be incorporated into the decision-making process.

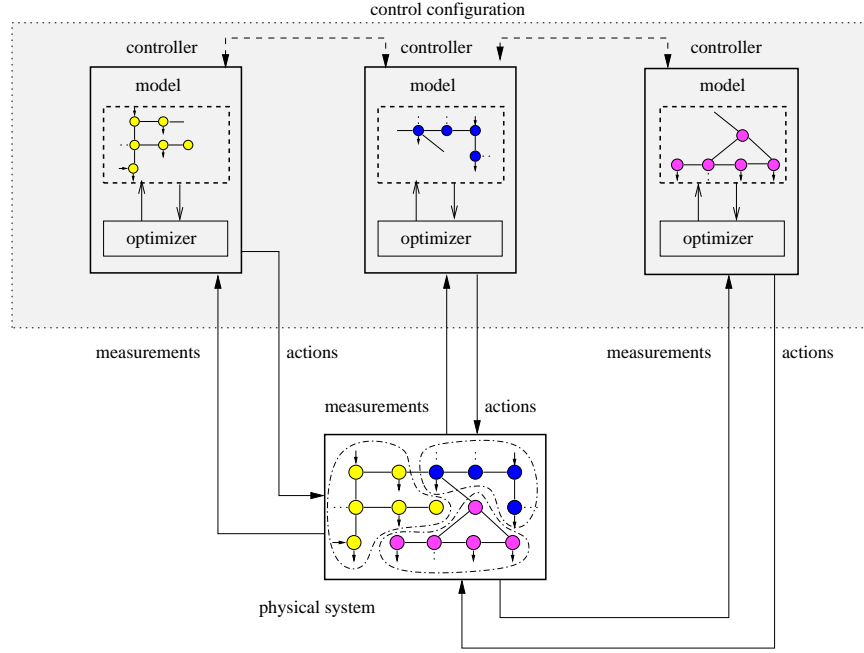


FIGURE 1. Illustration of distributed MPC control of a network.

Improvements in the operation of water systems can be achieved by cooperative distributed automatic controllers that coordinate their actions and take into account predictions or forecasts of future rainfall, future droughts, future arrival of increased water flow via rivers using various weather and hydrological sensors, and prediction models. One possible approach to accomplish this is the so-called distributed MPC philosophy, which will be summarized next.

3.2. Principle of operation. In a distributed MPC configuration, there are multiple controllers, each of them using MPC to control its own subsystem, i.e., its own part of the overall network, as illustrated in Figure 1. Usually, the division of the overall network into subsystems is based on the parts of the network that existing management organizations control. Similarly as in centralized MPC, the controllers in distributed MPC choose their actions at discrete control cycles. At each control cycle a local MPC controller uses the following information:

- a *prediction model* describing the behavior of its subsystem;
- a *performance criterion* expressing which subsystem behavior and actions are desired or should be penalized;
- *constraints* on the local states, inputs, outputs, and the interconnection variables;
- known information about future *disturbances* and exogenous inputs;
- a *measurement* of the state of the subsystem at the beginning of the current control cycle.

The goal of each controller is to determine those actions that optimize the behavior of the overall system and to minimize costs as specified through the performance criterion. In order to find the actions that lead to the best performance, each controller uses its prediction model to predict the behavior of its subsystem under

various actions over a certain prediction horizon, starting from the state at the beginning of the control cycle. Once the controller has determined the actions that optimize the performance of its subsystem over the prediction horizon, it implements these actions until the beginning of the next control cycle. The problem is then reformulated based on the latest measurements and the controller determines new actions over a shifted prediction horizon. Hence, each controller operates in a receding or rolling horizon fashion to determine its actions.

To make accurate predictions of the evolution of a subsystem over the prediction horizon for a given sequence of actions, each controller requires the current state of its subsystem and predictions of the values of variables that interconnect the model of its subsystem with the model of other subsystems. The predictions of the values of these so-called *interconnecting variables* are based on the information communicated with the neighboring controllers. One particular class of methods aims at achieving cooperation among controllers in an iterative way, in which controllers perform several iterations consisting of local problem solving and communication within each control cycle. In each iteration, controllers then obtain information about the plans of neighboring controllers. This iterative process is designed to converge to local control actions that lead to overall optimal performance. The challenge in implementing a distributed MPC strategy comes from ensuring that the actions that the individual controllers choose result in a joint performance, that is comparable to the performance of a hypothetical centralized control configuration in which all information available at a central location would be used.

3.3. Overview of distributed MPC for water systems. Various distributed MPC schemes have been investigated since the 90s, e.g., in [1, 18, 17, 9, 19, 38, 28]. Below we survey those with applications to water systems.

In [11] a decentralized adaptive control approach for a water distribution system is presented. As test system a 40-kilometer long canal connected to three secondary canals and four main reservoirs in the south-east of France is considered. The system is supplied by the Bourne river and by pumping stations that bring water from the Isere river to the canals. A division of the system into four control sections is given. For each section an MPC controller is proposed. The MPC controllers have to ensure that water levels follow a reference trajectory, which is based on predicted water demands obtained from a medium-term planning. A quadratic performance criterion is formulated to encode this. For making predictions a multiple-input multiple-output polynomial representation of each section is considered. Besides the prediction model, no constraints are taken into account. To obtain the control law for each controller, an explicit solution of the minimization problem defined for the full system is first determined. Then, the auxiliary problem principle and a linearization technique are used to solve the overall problem in a parallel distributed by the controllers responsible for the various sections.

In [32] a decentralized predictive controller for water delivery canals is presented. The overall canal is decomposed into individual subsystems, each of them with an upstream control gate to be manipulated in order to maintain the downstream water level as close as possible to a target value, under external perturbations. This control goal is expressed in an performance criterion that minimizes deviations from given set-points and changes in inputs in a quadratic way. As prediction model transfer functions are used. Besides the prediction model, no constraints are considered. To obtain the control laws for each subsystem, an explicit solution of the control law is computed, which is consequently split into expressions for each gate. The scheme

presented involves no iterations of information exchange between the controllers of subsystems.

In [10] a distributed predictive control scheme is presented for control of irrigation systems. The system considered is the same system as considered in [11]. The performance criterion involves costs due to pumping stations and due to losing of waste water, and profits due to power generation. The prediction model is formulated as a discrete-time system of difference-algebraic equations. Upper and lower bound constraints are taken into account. The decomposition approach used to determine the control laws for individual controllers is the same as in [11], i.e., it is based on the use of the auxiliary problem principle, which results in a parallel distributed MPC scheme. The control scheme is given, but no simulation results are presented.

4. A serial distributed MPC scheme. In [29] a distributed MPC scheme for control of general transportation networks is proposed. Water networks are a particular type of transportation networks, and therefore this scheme can also be suitable for distributed control of water networks. The scheme produces a sequence of control actions iteratively, which converge to globally optimal ones, if certain assumptions on the dynamics, the information available to controllers, and the control objectives are made. Below we briefly outline these assumptions and the steps of the scheme.

4.1. Dynamics. Let the water system be divided into n subsystems. Assume that the dynamics of subsystem $i \in \{1, \dots, n\}$ are given by a deterministic linear discrete-time time-invariant model (possibly obtained after symbolic or numerical linearization of a nonlinear model in combination with discretization):

$$\begin{aligned} \mathbf{x}_i(k+1) &= \mathbf{A}_i \mathbf{x}_i(k) + \mathbf{B}_{1,i} \mathbf{u}_i(k) + \mathbf{B}_{2,i} \mathbf{d}_i(k) + \mathbf{B}_{3,i} \mathbf{v}_i(k) \\ \mathbf{y}_i(k) &= \mathbf{C}_i \mathbf{x}_i(k) + \mathbf{D}_{1,i} \mathbf{u}_i(k) + \mathbf{D}_{2,i} \mathbf{d}_i(k) + \mathbf{D}_{3,i} \mathbf{v}_i(k), \end{aligned} \quad (1)$$

where at control cycle k , for subsystem i , $\mathbf{x}_i(k) \in \mathbb{R}^{n_{\mathbf{x}_i}}$ are the local states, $\mathbf{u}_i(k) \in \mathbb{R}^{n_{\mathbf{u}_i}}$ are the local inputs, $\mathbf{d}_i(k) \in \mathbb{R}^{n_{\mathbf{d}_i}}$ are the local known exogenous inputs, $\mathbf{y}_i(k) \in \mathbb{R}^{n_{\mathbf{y}_i}}$ are the local outputs, $\mathbf{v}_i(k) \in \mathbb{R}^{n_{\mathbf{v}_i}}$ are the remaining variables influencing the local dynamical states and outputs, and $\mathbf{A}_i \in \mathbb{R}^{n_{\mathbf{x}_i} \times n_{\mathbf{x}_i}}$, $\mathbf{B}_{1,i} \in \mathbb{R}^{n_{\mathbf{x}_i} \times n_{\mathbf{u}_i}}$, $\mathbf{B}_{2,i} \in \mathbb{R}^{n_{\mathbf{x}_i} \times n_{\mathbf{d}_i}}$, $\mathbf{B}_{3,i} \in \mathbb{R}^{n_{\mathbf{x}_i} \times n_{\mathbf{v}_i}}$, $\mathbf{C}_i \in \mathbb{R}^{n_{\mathbf{y}_i} \times n_{\mathbf{x}_i}}$, $\mathbf{D}_{1,i} \in \mathbb{R}^{n_{\mathbf{y}_i} \times n_{\mathbf{u}_i}}$, $\mathbf{D}_{2,i} \in \mathbb{R}^{n_{\mathbf{y}_i} \times n_{\mathbf{d}_i}}$, $\mathbf{D}_{3,i} \in \mathbb{R}^{n_{\mathbf{y}_i} \times n_{\mathbf{v}_i}}$ determine how the different variables influence the local states and outputs of subsystem i . We are in this paper primarily interested in distributed decision making and therefore we do not include noise terms representing uncertainty in the model. The $\mathbf{v}_i(k)$ variables in the model appear due to the fact that a subsystem is connected to other subsystems. Hence, the $\mathbf{v}_i(k)$ variables represent the influence of other subsystems on subsystem i . If the values of $\mathbf{v}_i(k)$ would be constant, then the dynamics of subsystem i would be decoupled from the other subsystems. However, in practice, the variables $\mathbf{v}_i(k)$ are equal to some of the variables of models representing dynamics of neighboring subsystems. Communication and coordination between the controller of subsystem i and the controllers of neighboring subsystems is then necessary to obtain agreement on the values of $\mathbf{v}_i(k)$.

Below we refer to the interconnecting input variables $\mathbf{w}_{\text{in},j,i}(k) \in \mathbb{R}^{n_{\text{w}_{\text{in},j,i}}}$ as the variables of subsystem i that are influenced by subsystem j , i.e., a selection of $\mathbf{v}_i(k)$. We refer to the interconnecting output variables $\mathbf{w}_{\text{out},j,i}(k) \in \mathbb{R}^{n_{\text{w}_{\text{out},j,i}}}$ as the variables of subsystem i that influence a neighboring subsystem j , i.e., a selection




FIGURE 2. Illustration of the relation between the models and variables of two subsystems i and j .

of $\mathbf{x}_i(k)$, $\mathbf{u}_i(k)$, and $\mathbf{y}_i(k)$. Figure 2 illustrates the relations between the variables of the models of two subsystems.

Let subsystem i be connected to m_i neighboring subsystems. Let the set of indices of the m_i subsystems connected to subsystem i be denoted by the neighbors set $\mathcal{N}_i = \{j_{i,1}, \dots, j_{i,m_i}\}$. Define the interconnecting inputs and outputs for the control problem of agent i at control cycle k as:

$$\mathbf{w}_{\text{in},i}(k) = \mathbf{v}_i(k) \quad (3)$$

$$\mathbf{w}_{\text{out},i}(k) = \mathbf{K}_i [\mathbf{x}_i^T(k) \quad \mathbf{u}_i^T(k) \quad \mathbf{y}_i^T(k)]^T, \quad (4)$$

where \mathbf{K}_i is an interconnecting output selection matrix that contains zeros everywhere, except for a single 1 per row corresponding to a local variable that relates to an interconnecting output variable. The variables $\mathbf{w}_{\text{in},i}(k)$, $\mathbf{w}_{\text{out},i}(k)$ are partitioned such that:

$$\mathbf{w}_{\text{in},i}(k) = [\mathbf{w}_{\text{in},j_{i,1}i}^T(k), \dots, \mathbf{w}_{\text{in},j_{i,m_i}i}^T(k)]^T \quad (5)$$

$$\mathbf{w}_{\text{out},i}(k) = [\mathbf{w}_{\text{out},j_{i,1}i}^T(k), \dots, \mathbf{w}_{\text{out},j_{i,m_i}i}^T(k)]^T. \quad (6)$$

The interconnecting inputs to the control problem of agent i with respect to agent j must be equal to the interconnecting outputs from the control problem of agent j with respect to agent i , since the variables of both control problems model the same quantity. For agent i this thus gives rise to the following *interconnecting constraints*:

$$\mathbf{w}_{\text{in},ji}(k) = \mathbf{w}_{\text{out},ij}(k) \quad (7)$$

$$\mathbf{w}_{\text{out},ji}(k) = \mathbf{w}_{\text{in},ij}(k), \quad (8)$$

for all $j \in \mathcal{N}_i$.

4.2. Available information. We assume that each of the subsystems $i \in \{1, \dots, n\}$ is controlled by a controller i that:

- has a prediction model of the form (1)–(2) of the dynamics of subsystem i ;
- can measure or estimate the state $\mathbf{x}_i(k)$ of its subsystem;
- can determine settings $\mathbf{u}_i(k)$ for the actuators of its subsystem;
- can estimate exogenous inputs $\mathbf{d}_i(k+l)$ of its subsystem over a certain horizon of length N , for $l = 0, 1, \dots, N-1$;
- can communicate with neighboring controllers, i.e., the controllers controlling the subsystems $j \in \mathcal{N}_i$.

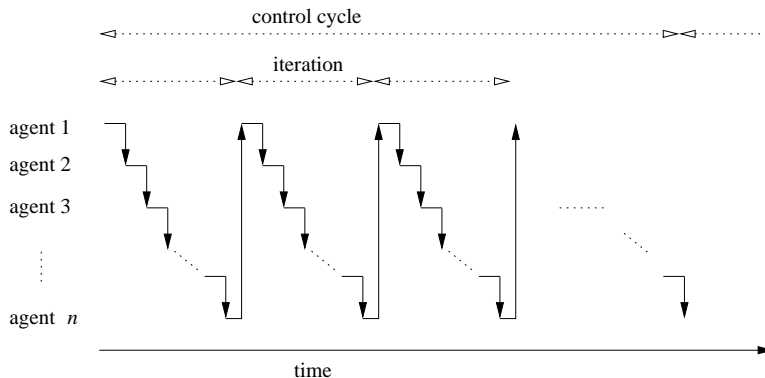


FIGURE 3. Illustration of the order of computations for the serial MPC scheme. One agent after another performs computations. When all agents have performed their computations a single iteration has been finished. After several iterations a control cycle has been finished.

4.3. Control objectives. We design the controllers to be cooperative, meaning that the individual controllers strive for the best overall system performance. In addition, we assume that the objectives of the controllers can be represented by convex functions $J_{\text{local},i}$, for $i \in \{1, \dots, n\}$, which are typically quadratic. Such functions are commonly encountered, in particular for systems that can be represented by (1)–(2), e.g., when the difference between a state variable (such as a water level) and a pre-specified set-point for that state variable (such as a pre-specified water level) are penalized, as is also illustrated in Section 5.5.

4.4. Control algorithm. We will use the tilde notation to represent variables over the prediction horizon e.g., $\tilde{\mathbf{u}}_i(k) = [\mathbf{u}_i^T(k), \dots, \mathbf{u}_i^T(k+N-1)]^T$. The distributed MPC scheme that we employ comprises the following steps at control cycle k :

1. For $i = 1, \dots, n$, controller i makes a measurement of the current state $\mathbf{x}_i(k)$ of the subsystem and estimates the expected exogenous inputs $\mathbf{d}_i(k+l)$, for $l = 0, \dots, N-1$.
2. The controllers cooperatively solve their control problems in the following serial iterative way, as also illustrated in Figure 3:
 - (a) Set the iteration counter s to 1 and initialize the Lagrange multipliers $\tilde{\lambda}_{\text{in},ji}^{(s)}(k)$, $\tilde{\lambda}_{\text{out},ij}^{(s)}(k)$ arbitrarily¹.
 - (b) For $i = 1, \dots, n$, one controller i after another determines $\tilde{\mathbf{x}}_i^{(s)}(k+1)$, $\tilde{\mathbf{u}}_i^{(s)}(k)$, $\tilde{\mathbf{w}}_{\text{in},ji}^{(s)}(k)$, $\tilde{\mathbf{w}}_{\text{out},ji}^{(s)}(k)$ as solutions of the following optimization problem:

$$\min_{\substack{\tilde{\mathbf{x}}_i(k+1), \tilde{\mathbf{u}}_i(k), \tilde{\mathbf{y}}_i(k), \\ \tilde{\mathbf{w}}_{\text{in},i}^{(s)}(k), \tilde{\mathbf{w}}_{\text{out},i}^{(s)}(k)}} J_{\text{local},i}(\tilde{\mathbf{x}}_i(k+1), \tilde{\mathbf{u}}_i(k), \tilde{\mathbf{y}}_i(k)) + \sum_{j \in \mathcal{N}_i} J_{\text{inter},i}^{(s)}(\tilde{\mathbf{w}}_{\text{in},ji}^{(s)}(k), \tilde{\mathbf{w}}_{\text{out},ji}^{(s)}(k)), \quad (9)$$

¹The Lagrange multipliers can in principle be initialized arbitrarily. However, initializing them with values close to the optimal Lagrange multipliers will increase the speed of convergence of the decision making process. Therefore, initializing the Lagrange multipliers with values obtained from the previous control cycle is beneficial, since typically these Lagrange multipliers will be good initial guesses for the new solution. This is referred to as *warm start*.

subject to the local dynamics (1)–(2) and (3)–(4) of subsystem i over the horizon, the current state $\mathbf{x}_i(k)$, and the known exogenous inputs $\tilde{\mathbf{d}}_i(k)$. The additional performance criterion $J_{\text{inter},i}$ in (9) at iteration s is defined as

$$J_{\text{inter},i}^{(s)}(\tilde{\mathbf{w}}_{\text{in},ji}(k), \tilde{\mathbf{w}}_{\text{out},ji}(k)) = \begin{bmatrix} \tilde{\lambda}_{\text{in},ji}^{(s)}(k) \\ -\tilde{\lambda}_{\text{out},ij}^{(s)}(k) \end{bmatrix}^T \begin{bmatrix} \tilde{\mathbf{w}}_{\text{in},ji}(k) \\ \tilde{\mathbf{w}}_{\text{out},ji}(k) \end{bmatrix} + \frac{\gamma_c}{2} \left\| \begin{bmatrix} \tilde{\mathbf{w}}_{\text{in,prev},ij}(k) - \tilde{\mathbf{w}}_{\text{out},ji}(k) \\ \tilde{\mathbf{w}}_{\text{out,prev},ij}(k) - \tilde{\mathbf{w}}_{\text{in},ji}(k) \end{bmatrix} \right\|_2^2,$$

where the 2-norm $\|\mathbf{a}\|_2$ of a vector \mathbf{a} with elements a_1, \dots, a_{n_a} is defined as $\sqrt{a_1^2 + \dots + a_{n_a}^2}$, where $\tilde{\mathbf{w}}_{\text{in,prev},ij}(k) = \tilde{\mathbf{w}}_{\text{in},ij}^{(s)}(k)$ and $\tilde{\mathbf{w}}_{\text{out,prev},ij}(k) = \tilde{\mathbf{w}}_{\text{out},ij}^{(s)}(k)$ is the information computed at the current iteration s for each controller $j \in \mathcal{N}_i$ that has solved its problem *before* controller i in the *current* iteration s , and where $\tilde{\mathbf{w}}_{\text{in,prev},ij}(k) = \tilde{\mathbf{w}}_{\text{in},ij}^{(s-1)}(k)$ and $\tilde{\mathbf{w}}_{\text{out,prev},ij}(k) = \tilde{\mathbf{w}}_{\text{out},ij}^{(s-1)}(k)$ is the information computed at the *last* iteration $s-1$ for the other controllers. Furthermore, γ_c is a positive scalar that penalizes the deviation from the interconnecting variable iterates that were computed by the controllers before controller i in the current iteration and by the other controllers during the last iteration. The results $\tilde{\mathbf{w}}_{\text{out},ji}^{(s)}(k)$ and $\tilde{\lambda}_{\text{in},ji}^{(s)}(k)$ of the optimization are sent to controller j .

- (c) Update the Lagrange multipliers,

$$\tilde{\lambda}_{\text{in},ji}^{(s+1)}(k) = \tilde{\lambda}_{\text{in},ji}^{(s)}(k) + \gamma_c \left(\tilde{\mathbf{w}}_{\text{in},ji}^{(s)}(k) - \tilde{\mathbf{w}}_{\text{out},ij}^{(s)}(k) \right). \quad (10)$$

Send $\tilde{\lambda}_{\text{in},ji}^{(s+1)}(k)$ to controller j and receive the multipliers from controller j to be used as $\tilde{\lambda}_{\text{out},ij}^{(s+1)}(k)$.

- (d) Move on to the next iteration $s+1$ and repeat steps 2b–2c. The iterations stop when the following stopping condition is satisfied:

$$\left\| \begin{bmatrix} \tilde{\lambda}_{\text{in,err},j_{1,1}1}^{(s+1)}(k) \\ \vdots \\ \tilde{\lambda}_{\text{in,err},j_n,m_n n}^{(s+1)}(k) \end{bmatrix} \right\|_{\infty} \leq \gamma_{\epsilon}, \quad (11)$$

where $\tilde{\lambda}_{\text{in,err},ji}^{(s+1)}(k) = \tilde{\lambda}_{\text{in},ji}^{(s+1)}(k) - \tilde{\lambda}_{\text{in},ji}^{(s)}(k)$, γ_{ϵ} is a small positive scalar and $\|\cdot\|_{\infty}$ denotes the infinity norm, defined for a vector \mathbf{a} with elements a_1, \dots, a_{n_a} as: $\max(|a_1|, \dots, |a_{n_a}|)$, where $|\cdot|$ denotes the absolute value operator. Note that since each individual component of the infinity norm depends only on variables of one particular controller, it is easy to devise very simple communication schemes that check satisfaction of this stopping condition after each iteration (e.g., by using a binary flag that is passed around).

3. The controllers implement the actions until the beginning of the next control cycle, when the procedure is repeated from step 1.

5. Case study: Control of an irrigation canal. In this section we formulate the dynamics and control scheme of a particular water system, viz. an irrigation canal. Irrigation canals are used mostly to transport water from source locations, such as

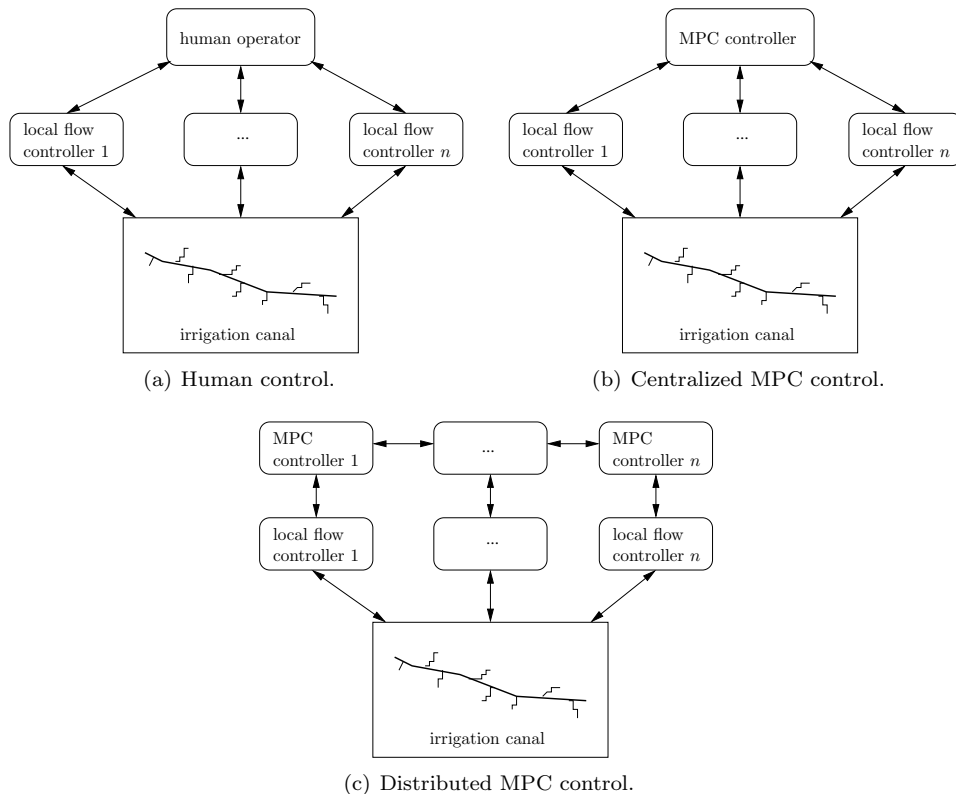



FIGURE 4. Three control configurations for the control of irrigation canals.

lakes and large rivers, to sink locations, such as small rivers and pipes transporting water to agricultural fields of farmers. Irrigation canals consist of several connected canal reaches, the inflow or outflow of which can be controlled using structures such as so-called overshoot or undershot gates, which restrict the flow of water flowing from an upstream canal reach into a downstream canal reach [6]. These structures usually have a local flow controller that regulates the position of the gates in order to obtain a certain flow. We focus on determining the set-points for the local flow controllers at these structures.

5.1. Control configurations. Figure 4 illustrates three possible control configurations for irrigation canals. Currently the configuration of Figure 4(a) is typically used in practice. A human operator manually adjusts the set-points for the local flow controllers at the undershot and overshoot gates. This manual process is expensive, since the human operator has to travel from one control structure to the next, possibly several times per day [36]. A more advanced control configuration is depicted in Figure 4(b). In this case, the determination of the set-points for the local flow controllers has been centralized and automated. Although implementation of such a centralized control configuration may be feasible in practice for relatively small water systems, the increasing computational requirements prevent application to larger systems. A centralized control configuration does not scale well and moreover constitutes a single point of failure. Instead of a centralized

FIGURE 5. Illustration of canal reach i and associated variables.

control configuration, the control configuration in Figure 4(c) may be employed, i.e., a distributed control configuration may be installed, in which set-points are autonomously decided upon by the distributed controllers based on local communication and cooperation.

5.2. Benchmark system. The irrigation canal that we consider is based on the West-M canal, which is an irrigation canal in the south of Phoenix, Arizona. This canal has been used by the ASCE Task Committee on Canal Automation Algorithms to define Test Canal 1 for testing automatic control schemes [7]. The canal is used to provide water to farmers. The length of the canal is almost 10 km and the maximum capacity of the head gate is $2.8 \text{ m}^3/\text{s}$ [36]. The irrigation canal that we consider consists of 7 canal reaches. At each of the reaches, water can be taken out at offtakes for irrigation purposes. Between each of the reaches control structures are present in the form of undershot gates to control the water flow locally. These control structures are equipped with local flow controllers that adjust the height of the undershot gate in order to meet a set-point for the water flow.

For the benchmark system under study, MPC schemes have been proposed based on a single controller determining in a centralized way the set-points for the local flow controllers, cf. Figure 4(b). MPC has been proposed for controlling the first 2 canal reaches of the benchmark system in [24], for controlling the first 3 canal reaches of the system in [30], and for controlling all canal reaches in [39, 36].

In this paper, we propose to use the distributed MPC scheme of Section 4.4 to take over the centralized control task, cf. Figure 4(c). Using a distributed approach there is no need for a central control location in the system, local information is available to local controllers, which share them with their neighboring local controllers.

5.3. Dynamics of irrigation canals. The dynamics of irrigation canals can be modeled in detail, e.g., using the Saint Venant equations [6] resulting in systems of highly-nonlinear partial differential-algebraic equations. However, using such highly-detailed models for MPC results in significant requirements on the computational power, and precise values of parameters are often not available. Therefore, similarly as in [36], we employ the discrete-time linear integral delay model [34] to model the dynamics of a canal reach. This model has shown to adequately capture relevant dynamics [34], and it reduces computations required for simulation of the dynamics significantly.

Let time be discretized into control cycles k and let the length of one control cycle be T_c s. Each canal reach is considered to have an inflow from an upstream canal

reach, as illustrated in Figure 5. Let this inflow into reach i be given by $q_{\text{in},i}(k) \in \mathbb{R}_0^+$ (m^3/s), with \mathbb{R}_0^+ the set of the positive real numbers. A canal reach has an outflow to a downstream canal reach. Let $q_{\text{out},i}(k) \in \mathbb{R}_0^+$ (m^3/s) denote this outflow. In addition to this inflow and outflow due to upstream and downstream canal reaches, there can be additional local inflow (e.g., due to rainfall) and outflow (e.g., due to outflow caused by farmers). Let such inflow be represented by $q_{\text{ext},\text{in},i}(k) \in \mathbb{R}_0^+$ (m^3/s) and such outflow by $q_{\text{ext},\text{out},i}(k) \in \mathbb{R}_0^+$ (m^3/s). The inflow $q_{\text{ext},\text{in},i}(k)$ and outflow $q_{\text{ext},\text{out},i}(k)$ are assumed to be known or predicted accurately in advance.

Depending on how the inflows and outflows change over time, the levels of the water in the reaches will change. Instead of considering the levels of the water at each location in the reaches, we only consider the levels of the water at the downstream end of each reach, since that is where the water offtakes take place and where the water level have to be kept close to the given set-points. In addition to the amount of inflow and outflow, also the surface of the reach influences how much the level of the water will change. Let $h_i(k) \in \mathbb{R}_0^+$ (m) denote the level of the water in canal reach i , and let the surface of reach i be $c_i \in \mathbb{R}_0^+$ (m^2). It takes some time for a change in the inflow of reach i to result in a change of the water level at the downstream end of the reach. Let this delay be $k_{\text{d},i}$ control cycles for reach i .

Using the variables defined above, the model describing how the level of the water in the canal reach changes from one control cycle k to the next control cycle $k + 1$ is given by:

$$\begin{aligned} h_i(k+1) = h_i(k) &+ \frac{T_c}{c_i} q_{\text{in},i}(k - k_{\text{d},i}) - \frac{T_c}{c_i} q_{\text{out},i}(k) \\ &+ \frac{T_c}{c_i} q_{\text{ext},\text{in},i}(k) - \frac{T_c}{c_i} q_{\text{ext},\text{out},i}(k). \end{aligned} \quad (12)$$

Canal reaches are connected to one another. When two canal reaches are connected to each other, the inflow of one canal reach is equal to the outflow of the other. Hence, for neighboring reaches i and j this interconnection is expressed by

$$q_{\text{out},i}(k) = q_{\text{in},j}(k). \quad (13)$$

The dynamics of canal reach i are conveniently written down in the state-space form (1)–(2) by defining

$$\begin{aligned} \mathbf{x}_i(k) &= \begin{bmatrix} h_i(k) \\ q_{\text{in},i}(k - k_{\text{d},i}) \\ \vdots \\ q_{\text{in},i}(k - 1) \end{bmatrix} & \mathbf{d}_i(k) &= \begin{bmatrix} q_{\text{ext},\text{in},i}(k) \\ q_{\text{ext},\text{out},i}(k) \end{bmatrix} \\ \mathbf{u}_i(k) &= q_{\text{in},i}(k) & \mathbf{v}_i(k) &= q_{\text{out},i}(k) & \mathbf{y}_i(k) &= \mathbf{x}_i(k) \end{aligned}$$

and

$$\mathbf{A}_i = \begin{bmatrix} 1 & \frac{T_c}{c_i} & 0 & \dots & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & 1 \\ 0 & \dots & \dots & \dots & \dots & 0 \end{bmatrix}$$

$$\begin{aligned} \mathbf{B}_{1,i} &= \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix} & \mathbf{B}_{2,i} &= \begin{bmatrix} \frac{T_c}{c_i} & -\frac{T_c}{c_i} \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \end{bmatrix} & \mathbf{B}_{3,i} &= \begin{bmatrix} -\frac{T_c}{c_i} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \\ \mathbf{C}_i &= [1 \ 0 \ \cdots \ 0] \\ \mathbf{D}_{1,i} &= 0 & \mathbf{D}_{2,i} &= [0 \ 0] & \mathbf{D}_{3,i} &= 0, \end{aligned}$$

and

$$\begin{aligned} \mathbf{w}_{\text{in},j_{i,\text{down}}i}(k) &= q_{\text{out},i}(k) \\ \mathbf{w}_{\text{out},j_{i,\text{up}}i}(k) &= q_{\text{in},i}(k), \end{aligned}$$

where $j_{i,\text{up}}$ and $j_{i,\text{down}}$ are the index of the upstream and downstream canal reach of reach i , respectively.

5.4. Available information. There are n controllers, and controller i is responsible for canal reach i . Controller i can measure the water level in its canal reach, can adjust the set-point for the flow controller at its upstream gate, and can communicate with the controllers of the canal reaches immediately upstream and downstream of the canal reach. In addition, controller i can obtain the expected water offtakes and rainfall with respect to its canal reach.

The actions that are optimal for each of the controllers depend on one another, since if one controller decides to increase its inflow, the water level in the upstream reach will decrease and therefore influences the decision making process of the upstream controller.

5.5. Control objectives. The set-points determined by the controllers and provided to the local flow controllers of the undershot gates should be chosen in such a way that

1. the deviations of water levels h_i from provided set-points $h_{\text{ref},i} \in \mathbb{R}_0^+$, for all $i \in \{1, \dots, n\}$ are minimized;
2. the changes in the set-points u_i provided to the local flow controllers are minimized to reduce equipment wear.

The performance criterion $J_{\text{local},i}$ can be written as

$$J_{\text{local},i}(\cdot) = \sum_{l=0}^{N-1} p_{h,i} (h_i(k+1+l) - h_{\text{ref},i})^2 + \sum_{l=0}^{N-1} p_{u,i} (u_i(k+l) - u_i(k-1+l))^2,$$

where for controller i , $p_{h,i} \in \mathbb{R}_0^+$ is the quadratic cost on the water level deviation, $p_{\Delta h,i} \in \mathbb{R}_0^+$ is the quadratic cost on a change in the water level deviation, and $p_{u,i} \in \mathbb{R}_0^+$ is the quadratic cost on a change in the set-point provided to the local flow controller.

5.6. Simulation. In this section, we describe a simulation result to illustrate the performance of the controllers discussed in this paper. We have implemented the model of the benchmark irrigation canal consisting of 7 canal reaches in Matlab v7.3. For solving the optimization problems at each control cycle step, we use the ILOG CPLEX v10.0 quadratic programming solver through the Tomlab v5.7 interface [15] to Matlab.

TABLE 1. Values of the parameters of the model, taken from [36].

reach i	1	2	3	4	5	6	7
$k_{d,i}$ (steps)	1	3	1	1	9	3	5
c_i (m ²)	397	653	503	1530	1614	2000	1241

We compare the performance of the distributed MPC scheme with a centralized scheme, i.e., we compare the performance of a control configuration of Figure 4(c) with a corresponding control configuration of Figure 4(b).

5.6.1. *Scenario.* The parameters used for the model of the irrigation canal are shown in Table 1. The control cycle length T_c is 240 s. The controllers use a perfect prediction model, in the sense that the combination of the prediction models of the individual controllers is the same as the model used for the simulation of the complete irrigation canal. The controllers use as parameters $N = 31$, $\gamma_c = 1$, $\gamma_\epsilon = 1.10^{-4}$. A prediction horizon length of 31 is chosen to take into account the total delay present in the irrigation canal [36]. The cost coefficients that the controllers use are chosen as $p_{h,i} = 100$ and $p_{u,i} = 10$, for $i \in \{1, \dots, n\}$.

As the simulation scenario we consider a sudden increase of $0.1 \text{ m}^3/\text{s}$ at control cycle $k = 30$ in the water offtake of canal reach 3 and a sudden decrease of $0.1 \text{ m}^3/\text{s}$ at control cycle $k = 70$ in the same canal reach.

5.6.2. *Results.* Figure 6(a) shows the changes in the set-points decided upon by the various controllers. Figure 6(b) shows the closed-loop evolution of the deviations of the water levels from the reference values. It can be seen that the inflow of canal reach 1 is increased right before the additional offtake increase (at $k = 30$) takes place to prevent having a too low water level after the additional offtake. It can also be observed that the deviations of the water levels after the offtake increase are minimal due to the changes in the set-points. We observe that after about 25 control cycles, the set-points settle, while the deviations of the water levels from the references are minimal, and that thus the controllers have performed their tasks adequately. Around control cycle $k = 70$, i.e., the control cycle at which the offtake decrease takes place, we clearly observe the inverse reaction.

The costs computed over the full simulation using the distributed MPC scheme are 0.1095 units. A centralized MPC controller based on the same objectives obtains costs over the full simulation of 0.1093 units. This difference in performance is negligible (and can in fact be made even smaller by decreasing the value of γ_ϵ), and hence, in this case, the distributed controllers have achieved a performance comparable to the performance obtained by a centralized controller.

6. Conclusions and future research. In this paper, we have surveyed the present literature on control of water systems. Most of the control problems of individual water systems can be solved. However, when global objectives are set for multiple water systems in a larger area, the traditional local solution process cannot be applied. A solution to this type of problem requires coordination between the water systems in the various areas. A promising solution investigated in this paper is to create a distributed framework for interacting water systems. The potential of distributed control for rivers, canal networks, sewer systems, irrigation canals, and reservoirs has been treated.

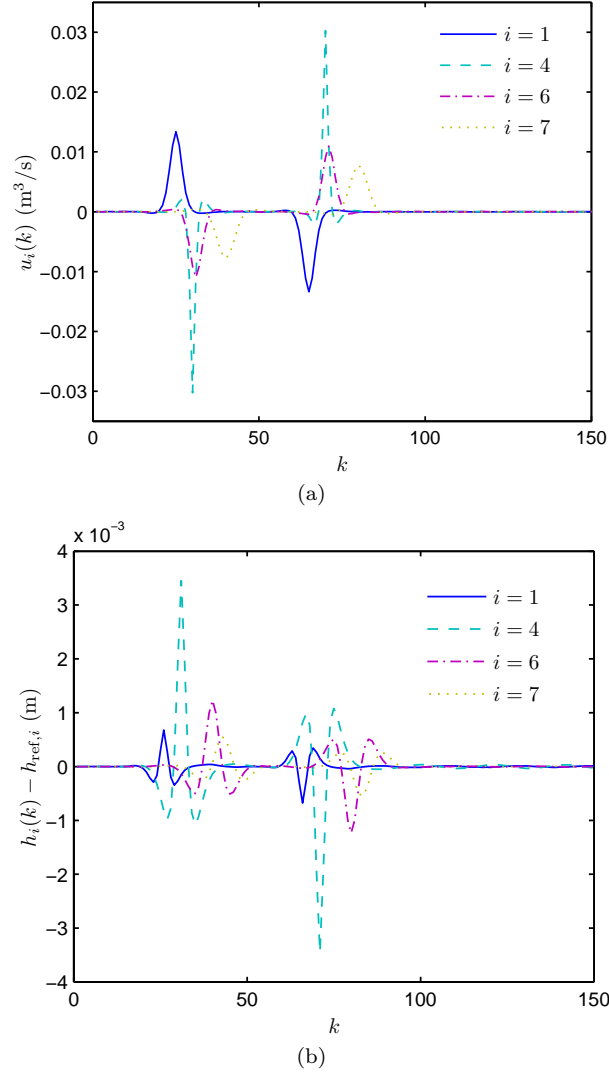


FIGURE 6. Evolution for four representative canal reaches of (a) set-points and (b) deviation of the water levels from reference values.

We have considered model predictive control (MPC) for distributed control of water systems as a particularly promising control approach. We have presented the use of a serial, iteration-based, distributed MPC scheme for the control of irrigation canals. We have illustrated the potential of the approach in a simulation study on a 7-reach irrigation canal.

Future work consists of further assessing the performance of the proposed scheme, extending the system model to include constraints on the minimal and maximal flow possible through undershot and overshot gates, assessing the performance of the distributed MPC scheme using linear prediction models on a nonlinear simulation model of the canal, and, based on this assessment, further improving the system

model if necessary. In addition, in future work, more complex case studies, including branching irrigation canals and canals in which the direction of water flow can change will be considered.

Acknowledgements. This research was funded by the BSIK project “Next Generation Infrastructures (NGI)”, the Delft Research Center Next Generation Infrastructures, the European STREP project “Hierarchical and distributed model predictive control (HD-MPC)”, the European 6th Framework Network of Excellence “HYbrid CONtrol: Taming Heterogeneity and Complexity of Networked Embedded Systems (HYCON)”, and the VIDI project “Multi-Agent Control of Large-Scale Hybrid Systems” (DWV.6188) of the Dutch Technology Foundation STW.

REFERENCES

- [1] L. Acar, *Some examples for the decentralized receding horizon control*, Proceedings of the 31st IEEE Conference on Decision and Control (Tucson, Arizona), 1992, 1356–1359.
- [2] K. Akouz, A. Benhammou, P. O. Malaterre, B. Dahbou, and G. Roux, *Predictive control applied to ASCE canal 2*, Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics (San Diego, California), October 1998, 3920–3924.
- [3] T. Barjas Blanco, P. Willems, B. De Moor, and J. Berlamont, *Flood prevention of the Demer using model predictive control*, Proceedings of the 17th IFAC World Congress (Seoul, Korea), July 2008, 3629–3634.
- [4] O. Begovich, V. M. Ruiz, G. Besançon, C. I. Aldana, and D. Georges, *Predictive control with constraints of a multi-pool irrigation canal prototype*, Latin American Applied Research, **37** (2007), 177–185.
- [5] G. Cembrano, G. Wells, J. Quevedo, R. Pérez, and R. Argelaguet, *Optimal control of a water distribution network in a supervisory control system*, Control Engineering Practice, **8** (2000), 1177–1188.
- [6] V. T. Chow, “Open-Channel Hydraulics”, Mc-Graw-Hill Book, New York, New York, 1959.
- [7] A. J. Clemmens, T. F. Kacerek, B. Grawitz, and W. Schuurmans, *Test cases for canal control algorithms*, Journal of Irrigation and Drainage Engineering, **124** (1998), 23–30.
- [8] M. Damas, M. Salmerón, and J. Ortega, *ANNs and GAs for predictive controlling of water supply networks*, Proceedings of the International Joint Conference on Neural Networks (Como, Italy), July 2000, 365–370.
- [9] W. B. Dunbar, *Distributed receding horizon control of dynamically coupled nonlinear systems*, IEEE Transactions on Automatic Control, **52** (2007), 1249–1263.
- [10] H. El Fawal, D. Georges, and G. Bornard, *Optimal control of complex irrigation systems via decomposition-coordination and the use of augmented Lagrangian*, Proceedings of the 1998 International Conference on Systems, Man, and Cybernetics (San Diego, California), 1998, 3874–3879.
- [11] G. Georges, *Decentralized adaptive control for a water distribution system*, Proceedings of the 3rd IEEE Conference on Control Applications (Glasgow, UK), 1999, 1411–1416.
- [12] T. Geyer, M. Larsson, and M. Morari, *Hybrid emergency voltage control in power systems*, Proceedings of the European Control Conference 2003 (Cambridge, UK), September 2003, Paper 322.
- [13] M. Gómez, J. Rodellar, and J. A. Mantecón, *Predictive control method for decentralized operation of irrigation canals*, Applied Mathematical Modelling, **26** (2002), 1039–1056.
- [14] 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** (2005), 102–112.
- [15] K. Holmström, A. O. Göran, and M. M. Edvall, *User’s guide for Tomlab /CPLEX*, June 2007.
- [16] Intergovernmental Panel on Climate Change, “Climate Change 2001: Impact, Adaptation and Vulnerability”, Cambridge University Press, New York, New York, 2003.
- [17] D. Jia and B. H. Krogh, *Distributed model predictive control*, Proceedings of the 2001 American Control Conference (Arlington, Virginia), June 2001, 2767–2772.
- [18] M. R. Katebi and M. A. Johnson, *Predictive control design for large-scale systems*, Automatica, **33** (1997), 421–425.

- [19] T. Keviczky, F. Borrelli, and G. Balas, *Decentralized receding horizon control for large scale dynamically decoupled systems*, *Automatica*, **42** (2006), 2105–2115.
- [20] J. M. Maciejowski, “Predictive Control with Constraints”, Prentice-Hall, Harlow, UK, 2002.
- [21] Y. Majanne, *Model predictive pressure control of steam networks*, *Control Engineering Practice*, **13** (2005), 1499–1505.
- [22] P. O. Malaterre, *Regulation of irrigation canals*, *Irrigation and Drainage Systems*, **9** (1995), 297–327.
- [23] P. O. Malaterre and J.-P. Baume, *Modeling and regulation of irrigation canals: Existing applications and ongoing researches*, *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics (San Diego, California)*, October 1998, 3850–3855.
- [24] P. O. Malaterre and J. Rodellar, *Multivariable predictive control of irrigation canals. design and evaluation on a 2-pool model*, *Proceedings of the International Workshop on Regulation of Irrigation Canals: State of the Art of Research and Applications (Marrakech, Morocco)*, April 1997, 239–248.
- [25] P. O. Malaterre, D. C. Rogers, and J. Schuurmans, *Classification of canal control algorithms*, *Journal of Irrigation and Drainage Engineering*, **124** (1998), 3–10.
- [26] J. A. Mantecón, *A Simulink-based scheme for simulation of irrigation canal control systems*, *SIMULATION*, **78** (2002), 485–493.
- [27] I. Mareels, E. Weyer, S. K. Ooi, M. Cantoni, Y. Li, and G. Nair, *Systems engineering for irrigation systems: Successes and challenges*, *Annual Reviews in Control*, **29** (2005), 191–204.
- [28] R. R. Negenborn, “Multi-agent model predictive control with applications to power networks”, Ph.D. thesis, Delft University of Technology, Delft, The Netherlands, December 2007, Second printing September 2008.
- [29] 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** (2008), 353–366.
- [30] V. M. Ruiz and J. Ramirez, *Predictive control in irrigation canal operation*, *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics (San Diego, California)*, October 1998, 3897–3901.
- [31] V. M. Ruiz-Carmona, A. J. Clemmens, and J. Schuurmans, *Canal control algorithm formulations*, *Journal of Irrigation and Drainage Engineering*, **124** (1998), 31–39.
- [32] S. Sawadogo, R. M. Faye, P. O. Malaterre, and F. Mora-Camino, *Decentralized predictive controller for delivery canals*, *Proceedings of the 1998 IEEE International Conference on Systems, Man, and Cybernetics (San Diego, California)*, 1998, 3380–3384.
- [33] J. Schuurmans, “Control of water levels in open channels”, Ph.D. thesis, Delft University of Technology, Delft, The Netherlands, October 1997.
- [34] J. Schuurmans and M. Ellerbeck, *Linear approximation model of the WM canal for controller design*, *Proceedings of the First International Conference on Water Resources Engineering (San Antonio, Texas)*, June 1995, 353–357.
- [35] D. D. Šiljak, “Decentralized Control of Complex Systems”, Academic Press, Boston, Massachusetts, 1991.
- [36] P. J. van Overloop, “Model Predictive Control on Open Water Systems”, Ph.D. thesis, Delft University of Technology, Delft, The Netherlands, June 2006.
- [37] P. J. van Overloop, J. Schuurmans, R. Brouwer, and C. M. Burt, *Multiple model optimization of PI-controllers on canals*, *Journal of Irrigation and Drainage Engineering*, **131** (2005), 190–196.
- [38] A. Venkat, J. Rawlings, and S. Wright, *Stability and optimality of distributed model predictive control*, *Proceedings of the 44th IEEE Conference on Decision and Control and 2005 European Control Conference (Sevilla, Spain)*, December 2005, 6680–6685.
- [39] B. T. Wahlin, *Performance of model predictive control on ASCE test canal 1*, *Journal of Irrigation and Drainage Engineering*, **130** (2004), 227–238.
- [40] B. T. Wahlin and A. J. Clemmens, *Automatic downstream water-level feedback control of branching canal networks: Simulation results*, *Journal of Irrigation and Drainage Engineering*, **132** (2006), 208–219.
- [41] E. Weyer, *Control of irrigation channels*, *IEEE Transactions on Control Systems Technology*, **16** (2008), 664–675.

E-mail address: r.r.negenborn@tudelft.nl