




EPiC Series in Engineering

Volume 3, 2018, Pages 1281–1288

HIC 2018. 13th International
Conference on Hydroinformatics



A Model Predictive Control Framework for Real-Time Optimisation of Water System Operations

Henrik Madsen¹, Anne Katrine Falk¹ and Rasmus Halvgaard¹

¹DHI, Hørsholm, Denmark

hem@dhigroup.com, akf@dhigroup.com, rha@dhigroup.com

Abstract

We have developed a versatile Model Predictive Control (MPC) framework, which can handle real-time control of a large variety of water systems. The framework combines a fast-solvable optimisation model (a quadratic program) with evaluation and realignment by a detailed hydrological-hydrodynamic model. The flexibility of the MPC framework is highlighted by two case studies: (1) a large-scale river system with several weeks of travel time, and (2) an urban storm and wastewater system with a concentration time of about half an hour to one hour. Both case studies demonstrate a large potential for improving operations by system-wide real-time optimisation.

1 Introduction

Traditionally, real-time operation of water systems is carried out using rule-based regulation of individual or small groups of coordinated control structures. There is a large potential for improving operational efficiency by dynamic optimisation of the operation of all control structures within the water system. We have developed a Model Predictive Control (MPC) framework that uses a dynamic surrogate model of the system to optimise operations. The model predicts the change in the state of the system caused by changes in control actions and uncontrolled system forcing (e.g. meteorological forcing), and these predictions are used to optimise the control actions.

There are two major challenges for application of MPC for real-time optimisation: (1) computational requirements of the dynamic model, and (2) large dimension of the optimisation problem for water systems with a large number of control structures. We make use of a computationally efficient surrogate model in the MPC. The surrogate model is formulated as a linear model because this allows optimisation of very large problems with thousands of optimisation variables within seconds or few minutes.

2 Model Predictive Control Framework

A core element of MPC is the receding horizon principle. The control actions of the water system are optimised over a fixed control horizon based on a forecast of the system forcing. We use flow set points as control actions. The optimised set points are time varying, with a time resolution that depends on the system in question. Only the first part of the optimised operation is applied, then a new MPC optimisation is carried out with updated initial conditions and a new forecast of the uncontrolled system forcing.

Our MPC framework combines a high-fidelity hydrological-hydrodynamic simulation model of the water system with a fast surrogate model. The surrogate model is a computational efficient emulator of the high-fidelity model and is used as the internal dynamic model in the MPC. In between two MPC optimisations, the high-fidelity model is advanced and updated with available sensor measurements from the water system using data assimilation. Prior to a new MPC optimisation, the surrogate model is updated with initial conditions that are inferred from the high-fidelity model. The dynamic interaction between the high-fidelity and surrogate models is important to ensure that the surrogate model is continuously kept close to the true state of the system.

We have formulated a linear surrogate model, because in a linear model the future state of the system is an explicit (and linear) function of three factors: the initial conditions, the uncontrolled forcing, and the control (the regulated flows). The linearity is a key feature when defining bounds on the future states (e.g. future reservoir volumes). This forms the basis for defining an optimisation model that has purely linear constraints, and such optimisation models are very fast to solve. Specifically, we have chosen to formulate the optimisation as a quadratic program.

The constraints include both physical constraints and operational goals. Physical constraints are implemented as hard constraints that cannot be violated, whereas operational goals typically are formulated as soft constraints that can be violated with an associated penalty. The objective function includes terms that measure: (i) deviations from the desired system state and output, (ii) violations of soft constraints, and (iii) control activity.

We have developed a flexible modular framework for connecting network elements together to a surrogate model. Basic building blocks include:

- A reach/pipe block that describes the flow of water in river reaches and drainage pipes. It is based on a linear routing model and accounts for inflows and outflows distributed along the reach/pipe.
- A reservoir/basin block that describes the change of storage in reservoirs and detention basins as a function of inflows and outflows.
- A linear reservoir block that transforms inflow to outflow according to linear reservoir theory.
- A reservoir with overflow block that describes storage and overflow of a detention basin (Halvgaard and Falk, 2017).

The different building blocks are calibrated to resemble the high-fidelity model. For each block a set of features are defined, including inflows and outflows (both controlled and uncontrolled), hard and soft constraints, and contribution to the objective function. When the building blocks have been configured and linked together, the MPC framework automatically generates the quadratic program.

3 Case studies

We demonstrate the MPC framework of two real-time optimisation problems: (i) a large-scale river-reservoir system, and (ii) an urban storm and wastewater system. Spatial and temporal characteristics of the two systems are shown in Table 1. The table lists both physical characteristics (storage volume and total travel time) and MPC parameters (optimisation horizon, time step and how often the optimisation is repeated).

The choice of MPC parameters ultimately determines the size of the optimisation model in terms of the total number of optimisation variables as listed in Table 2 and detailed in the following subsections. The total number of optimisation variables is the sum of the number of control actions and the number of slack variables. The control actions are related to the trajectories of optimised flows at the regulators, whereas the slack variables are related to penalties for violating constraint thresholds (soft constraints).

Both case-studies are tested with perfect forecasts, which implies that there is no difference in the uncontrolled forcing of the dynamics of the MPC internal model and the high-fidelity model.

	Storage volumes	Total travel time	Optimisation horizon	Time step	Re-optimise every
River-reservoir	0.7-12.3 mill. m ³	3-4 weeks	2 weeks	3 hours	24 hours
Urban drainage	2,000-16,000 m ³	0.5-1 hour	2 hours	5 min	5 min

Table 1: Key figures that highlight the different spatial and temporal scales of the two case studies.

	Regulated flows	Control actions	Slack variables	Total number of optimisation variables
River-reservoir	9	1,008	3,024	4,032
Urban drainage	6	144	144	288

Table 2: Key figures that highlight the size of the optimisation models.

3.1 River-reservoir system

The MPC framework was applied for testing optimisation of storage operations of the Murrumbidgee River system in New South Wales, Australia (Falk et al., 2016). It is based on the Computer Aided River Management (CARM) project that was initiated by WaterNSW in 2011 with the aim to improve operational efficiency and reduce end-of-system surplus flows in the Murrumbidgee River system (van Kalken et al., 2012). The surrogate model is derived from the high-fidelity hydrologic-hydrodynamic model developed for the CARM project, which is based on the MIKE 11 river modelling system. Schematisation of the surrogate model is shown in Figure 1. The regulated part of the river is about 1,300 km from Burrinjuck Dam to Balranald with a total travel time of three to four weeks. Water is supplied from two upstream dams and from natural inflows along the river. A number of inline reservoirs regulate the flow and divert water to major irrigation areas. Three major irrigation areas account for approximately 70% of the total irrigation demand, and the remaining part is extracted by about 700 individual users.

Operation objectives include: (1) supply of ordered water to the major irrigation areas and individual users, (2) minimisation of spills at the end-of-system, and (3) keeping the river in a lean state to minimise evapotranspiration losses and make storage available for natural inflows. In the MPC optimisation model, the lean state is defined as either a target water level or an operational zone in the inline reservoirs. At the end-of-system at Balranald a target flow corresponding to environmental flow requirements is defined.

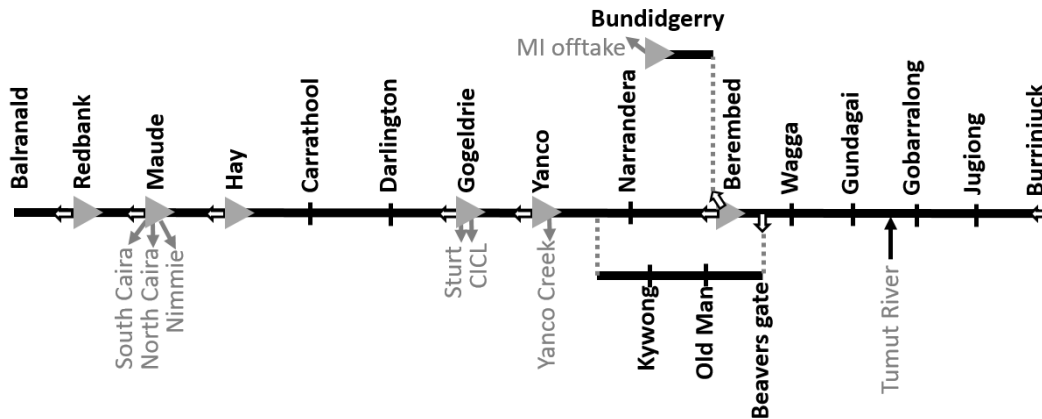


Figure 1: Surrogate model schematisation of the Murrumbidgee River system. The figure shows the river reaches (black lines), inline reservoirs (grey triangles), controllable gates (white arrows), and major offtakes (grey arrows). Tributary inflows and individual water users distributed along the river are not shown. Adapted from Falk et al. (2017).

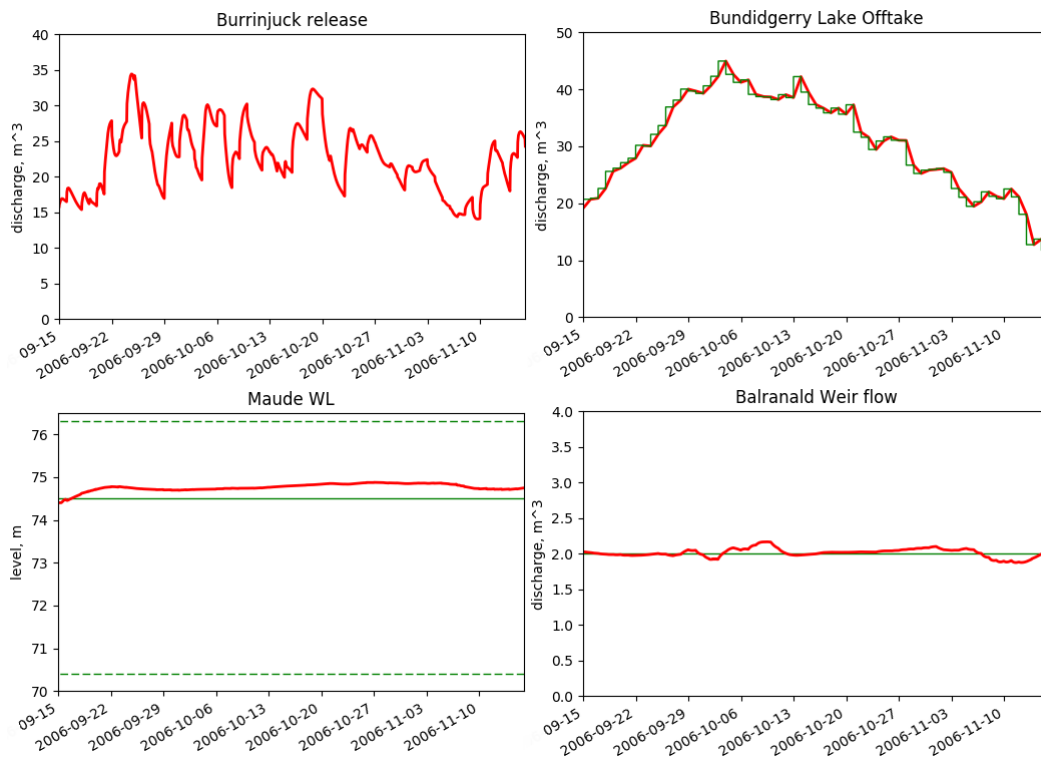


Figure 2: Test results for optimisation of the Murrumbidgee River system using the MPC framework: (i) release discharge [m^3/s] at Burrinjuck Dam (upper left), (ii) offtake discharge [m^3/s] at Bundidgerry Lake (upper right), (iii) water level [m] at Maude (lower left), and (iv) end-of-system flow [m^3/s] at Balranald (lower right). The figure includes the high-fidelity model results forced with optimised MPC releases (red line), target levels/flows and irrigation water demands (green line), and upper and lower bounds (green dashed line). Adapted from Madsen et al. (2017).

The MPC optimises release hydrographs at Burrinjuck Dam, the six inline reservoirs, and Beavers gate (see positions of controllable gates in Figure 1). The control horizon of the MPC is set to 14 days with a 3-hour temporal resolution for each release hydrograph, which gives in total 1,008 control actions to be optimised. Furthermore, 3,024 slack variables enter the optimisation model as additional optimisation variables. The slack variables are related to soft constraints on minimum flow in river reaches and soft constraints on upper and lower bound on reservoir levels. With respect to computational requirements, the MPC is very efficient. It takes about 60 seconds for one MPC optimisation on a standard laptop, and additionally 1.5 minutes to advance the high-fidelity model for initialising the next MPC optimisation.

The MPC was evaluated for a two-month test period (Madsen et al., 2017). Selected results from this test are shown in Figure 2. The figure shows the optimised release from Burrinjuck Dam (upper left). The MPC is able to meet all irrigation water demands at the major offtakes as well as individual water demands along the river for the entire test period. Water demand and offtake discharge for the Murumbidgee Irrigation (MI) area at Bundiderry Lake are shown in Figure 2 (upper right). At all inline reservoirs, water levels are kept close to the defined target levels and operational zones (see results for Maude in Figure 2, lower left), thus keeping the river in the lean state. At the end-of-system at Balranald the flow is kept very close to the target flow of 2 m³/s (Figure 2, lower right), and thereby essentially eliminating spills.

3.2 Urban storm and wastewater system

The MPC framework was set up and tested for control of an urban storm and wastewater system in Aarhus, Denmark with a concentration time of about half an hour to one hour (Halvgaard et al., 2017). The surrogate model has been derived and calibrated from a detailed hydrological-hydrodynamic model of the system, which is based on the MIKE URBAN modelling system. It models the major detention basins and the associated pipe network to flow set point locations as reservoir with overflow blocks. This means that each block represents the volume in a physical basin plus additional volume in surrounding pipes. Overflow occurs when this compound volume reaches a threshold. The schematisation of the surrogate model is shown in Figure 3. The main operation objective is to optimise use of the available storage in the detention basins under heavy rainfall in order to reduce combined sewer overflows. When the system is saturated, the control should be able to prioritise overflow between the different locations.

The MPC optimises the flow set points for water that leaves each basin. Each set point is transferred to a PID controller in the existing coordinating control layer. The control horizon of the MPC is 2 hours and the optimisation is repeated every 5 min. The set-point time series have a time resolution of 5 min, which means that for each of the six set-point locations there are 24 control actions within the two-hour control horizon, in total 144 control actions. Additionally, 144 slack variables enter the optimisation model. The slack variables are related to soft upper bounds on the reservoir volumes. Overflow occurs when this threshold is exceeded.

The MPC was evaluated for historical rain events that caused combined sewer overflows. Results for the largest of these rain events are shown in Figure 4. The MPC is tuned to prioritise overflow at the TB basin while minimising overflow at the MB basin. We achieved a total overflow volume reduction of 24.6% compared to the baseline MIKE URBAN model, which uses a rule-based control strategy. Overflow volumes are reduced at all detention basins: 10.7% at TB, 9.0% at HB, 35.6% at KB and 75.3% at JP. No overflows occur at the MB and CB basins for both the MPC solution and the baseline control.

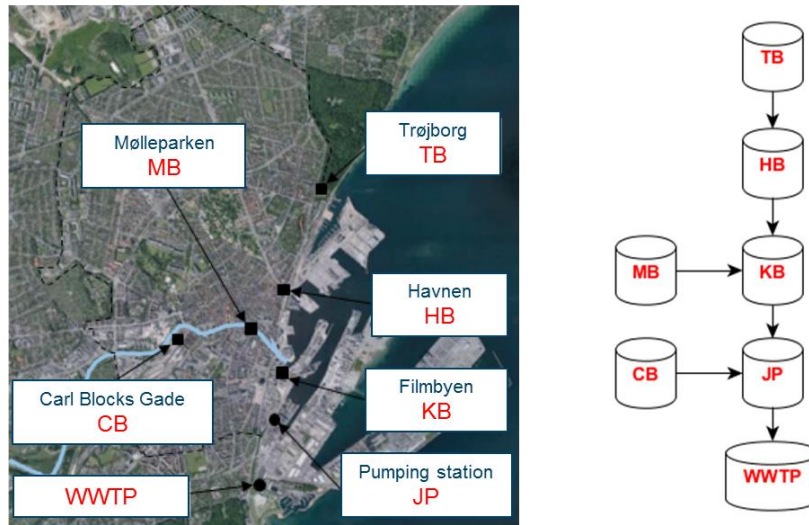


Figure 3: Surrogate model schematisation for the Marselisborg Nord part of the Aarhus sewer system. The map to the left shows the location of the physical detention basins, the diagram to the right shows the reservoir with overflow blocks in the MPC model. Each arrow represents a location where the flow set point is optimised.

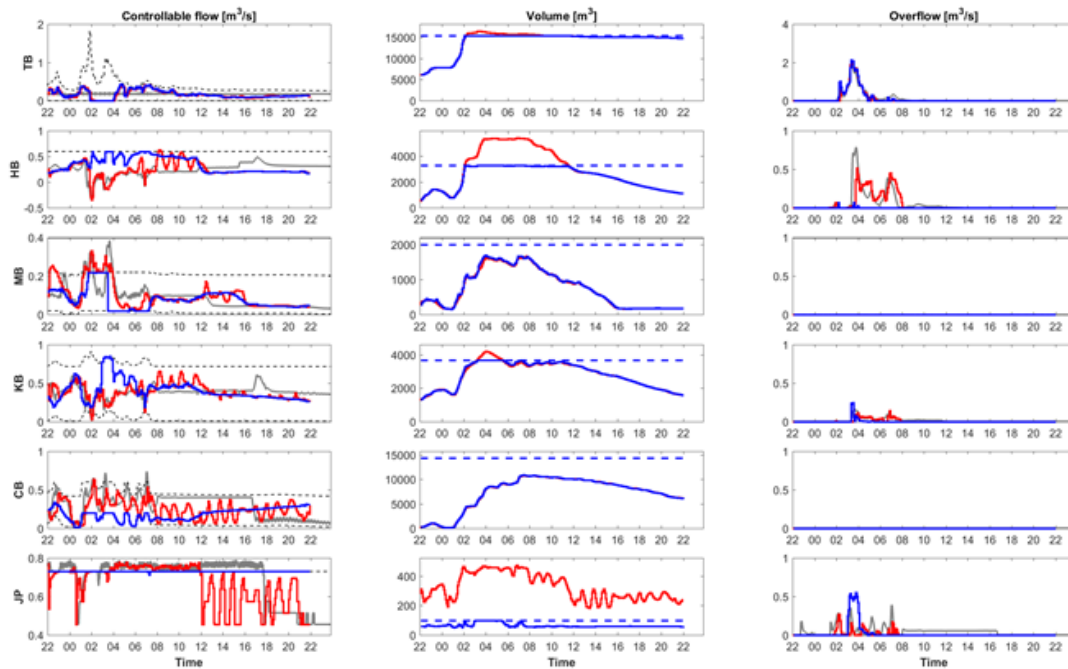


Figure 4: MPC results for a rain event in July 2014. Each row shows results for a detention basin in terms of optimised set point (1st column), stored volume (2nd column), and overflow discharge (3rd column). The figure includes high-fidelity model results (red line), MPC results (blue line), and simulated results using a baseline control strategy (grey line). The dashed lines are constraints. Adapted from Halvgaard et al. (2017).

The results in Figure 4 show some discrepancies between the high-fidelity and MPC surrogate model results. There are different reasons for these. First, the MPC surrogate model does not account for backwater effects. In some situations, the MPC may suggest unrealistic high flow set points while downstream conditions constrain the controllable flow. Secondly, the MPC performance depends heavily on the design of the underlying coordinating control layer. In some parts of the system, there is a travel time between the actuator and the downstream sensor of 15-30 minutes. This time delay significantly reduces the performance of the PID controllers.

4 Conclusions

We have developed an MPC framework, which can be configured for different types of water systems with different spatial scales and temporal dynamics. Two case studies are presented that demonstrate the flexibility and scalability of the framework. The river-reservoir system shares some basic features with the urban storm and wastewater system, namely storage of water and water routing characterised by system travel time. Thus, no matter whether the storage is in millions of cubic meters or only in thousands, and no matter whether the travel time is measured in weeks or hours, these basic features have been exploited to create a common MPC framework for water networks.

Both case studies demonstrate a large potential for improving real-time operation by system-wide optimisation using MPC. In the river-reservoir study all water demands are met, while at the same time the river is kept in a lean state and spills are essentially eliminated. The urban storm and wastewater case shows that even though there are some discrepancies between the surrogate model and the MIKE URBAN model, the MPC reduces the overflow volume significantly compared to the existing rule-based control.

The MPC framework is computationally very efficient, thus facilitating optimisation of large-scale water systems that include thousands of optimisation variables. This is crucial for real-time implementation. In addition, it allows a thorough testing and tuning of the MPC before being deployed in an operational system.

Acknowledgements

We would like to thank WaterNSW for providing data and make the MIKE 11 model of the Murrumbidgee River available for our study. Realdania supported the study through the “Klimaspring” programme as part of the project “Smart Cities Water Solutions”.

References

- Falk, A.K.V., Mackay, C., Madsen, H., Godiksen, P.N., 2016, Model predictive control of a large-scale river network, 12th International Conference on Hydroinformatics (HIC 2016) - Smart Water for the Future (Eds. Joong Hoon Kim, Hung Soo Kim, Do Guen Yoo, Donghwi Jung and Chang Geun Song), Procedia Engineering, 154, 80-87. <https://doi.org/10.1016/j.proeng.2016.07.422>
- Halvgaard, R., Falk, A.K.V., 2017, An efficient formulation of overflow structures for Model Predictive Control of urban drainage systems, Proceedings of 12th IWA Specialized Conference on Instrumentation, Control and Automation (ICA).
- Halvgaard, R., Falk, A.K.V., Lund, N.S.V., Borup, M., Mikkelsen, P.S., Madsen, H., 2017, Model predictive control for urban drainage: testing with a non-linear hydrodynamic model, Proceedings of 14th IWA/IAHR International Conference on Urban Drainage (ICUD), 2471-2474.

van Kalken, T., Nachiappan, N., Berry, D. and Skinner, J., 2012, An Optimized, Real Time Water Delivery and Management System for the Murrumbidgee River, In Proceedings of the 34th Hydrology and Water Resources Symposium, HWRS 2012, Sydney, Australia.

Madsen, H., Falk, A.K.V., Suter, S.O., 2017, Improving efficiency in river management operations using Model Predictive Control, 13th Hydraulics in Water Engineering Conference, 13-16 November 2017, Dockside, Sydney, Australia.