REAL-TIME OPTIMISATION OF LARGE RIVER AND RESERVOIR SYSTEM OPERATIONS

BY HENRIK MADSEN & ANNE KATRINE FALK

We present a new approach for real-time optimisation of large river and reservoir systems that combines physically-based hydrological-hydrodynamic models with surrogate control models within a Model Predictive Control framework. The approach is demonstrated in the optimisation of storage operations of the Murrumbidgee River system in New South Wales, Australia. This system consists of multiple reservoirs with pronounced travel times between the reservoirs. The test shows a large potential for improving operational efficiency of large water systems.

Real-time optimisation

There is a large potential for improving the operation of multi-purpose, large-scale river and reservoir systems by real-time optimisation. Reservoir operation rule curves have traditionally been optimised using a simulation-optimisation approach in which a high-fidelity, i.e. a detailed hydrological-hydrodynamic simulation model of the system is coupled with an optimisation algorithm (e.g. Ngo et al., 2007). However, for real-time optimisation of large systems such an approach is infeasible due to the curse of dimensionality and high computational requirements.

To facilitate real-time optimisation of very large water systems, we have developed a Model Predictive Control (MPC) framework that combines high-fidelity hydrological-hydrodynamic simulation models with simpler surrogate control models. The surrogate model is a computationally fast emulator of the high-fidelity model, which is sufficiently accurate for predicting the change in the system state due to changes in system operations and uncontrolled inflows and extractions. We apply linear surrogate models, which allow formulation of a fast-solvable optimisation model with 10,000s of optimisation variables.

Model predictive control framework

A central element of MPC is that it optimises the current operation while taken into account also the operation over a future time horizon. This is accomplished by optimising the operation over the full time horizon but only implementing the optimised operation for the first part. The optimisation is then repeated with new initial conditions and a new forecast of the system boundaries.

In our framework, a surrogate model is used as the internal dynamic model of the MPC. The surrogate model is derived and calibrated from the high-fidelity hydrological-hydrodynamic model of the river and reservoir system. The surrogate model is schematised using different building blocks as illustrated in Figure 1. The reservoir block describes reservoir storage as a function of inflows and outflows. The reach block describes water flow in a river reach using

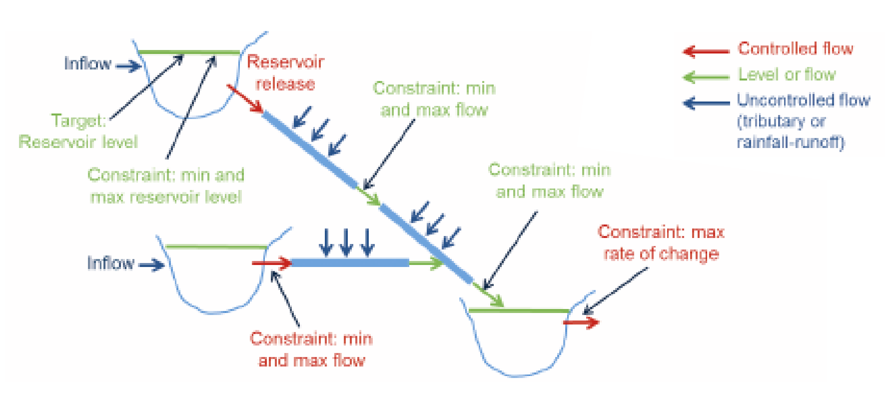

Figure 1. Example of schematisation of surrogate model and definition of constraints for the MPC optimisation model

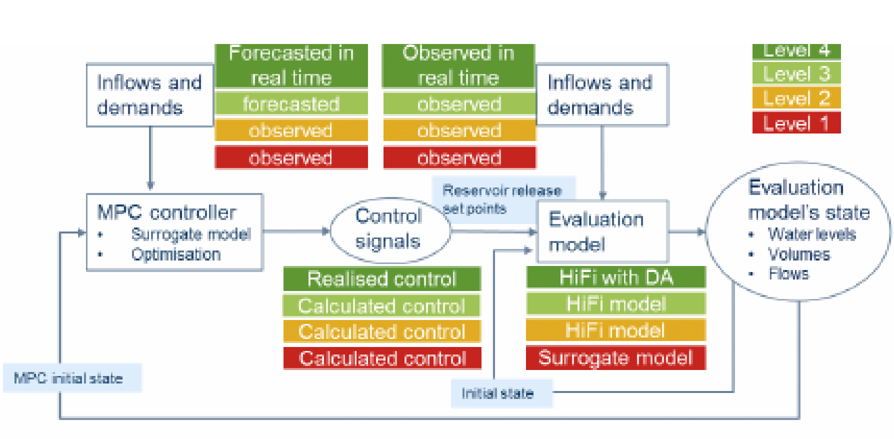

Figure 2. Closed loop test hierarchy for tuning and evaluation of the MPC controller. The colours from red to green illustrate the four test levels with use of different data for inflows and water demands for the MPC controller and evaluation model, control signals, and evaluation model (HiFi: high-fidelity, DA: data assimilation)

a linear routing model and takes into account inflows and extractions along the reach. Figure 1 also illustrates constraints that are defined for the MPC optimisation model, including both physical constraints (e.g. minimum river flow and maximum reservoir level) and operational goals (e.g. reservoir level target, and maximum river flow). Physical constraints are implemented as hard constraints (i.e. are not allowed to be violated), whereas operational goals typically are implemented as soft constraints (i.e. are allowed to be violated, but with an associated penalty).

An important element of our MPC framework is the dynamic interaction between the high-fidelity model and the surrogate models. For testing, the high-fidelity model is used as evaluation model to tune and evaluate the performance of the MPC controller. The high-fidelity model provides initial conditions to the surrogate model when a new MPC optimisation is initiated, which ensures that the surrogate model is not drifting away from the true state of the system. In the real-time implementation, water level and discharge observations are used to update the high-fidelity model using data assimilation.

A test hierarchy for tuning and evaluation of the MPC controller in a closed-loop test environment is illustrated in Figure 2. The test hierarchy includes four levels, which gradually increase the discrepancies between the surrogate model and the evaluation model. At Level 1, the MPC controller is evaluated against the surrogate model itself, whereas at Level 2 it is evaluated against the high-fidelity model. For both Level 1 and 2, observed inflows and water demands are used as input to the MPC controller and the evaluation model. At level 3, the MPC controller is forced with forecasted inflows and water demands to evaluate its robustness to forecast errors. Finally, at Level 4, the MPC controller is tested in a real-time environment where the MPC controller is forced with real-time forecast data. The high-fidelity model uses the realised control, real-time data of inflows and water demands, and water level and discharge measurements for data assimilation.

Optimising operations of the Murrumbidgee River system

We have applied the MPC framework for testing optimisation of the Murrumbidgee River system in New South Wales, Australia. The regulated part of the river is about 1,300 km from Burrinjuck Dam to Balranald (see illustration of the system in Figure 3). Along the river, a number of inline reservoirs regulate the flow and divert water to the 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. The system is operated for provision of water for irrigation and environmental flow requirements. Water is supplied from two upstream reservoirs, Blowering and Burrinjuck, and from natural inflows downstream.

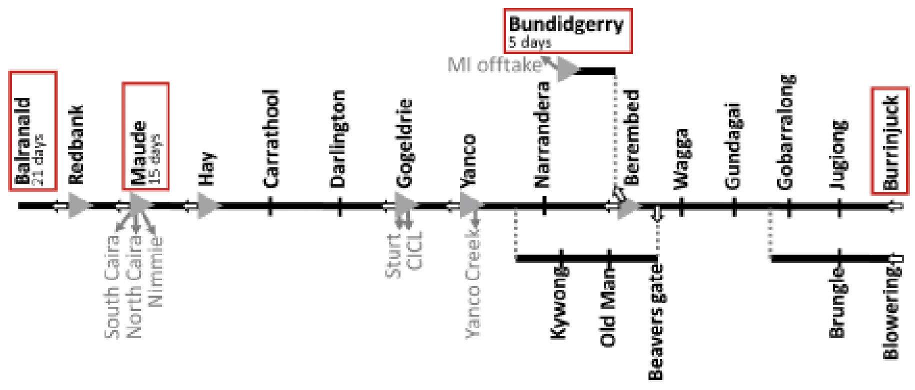

Figure 3. Surrogate model schematisation for the Murrumbidgee River system. The figure shows the river reaches (black lines), inline reservoirs (grey triangles), controllable gates (white arrows), and major offtakes (grey arrows). Tributary inflows and individual water users distributed along the river are not shown. For the selected locations shown in Figure 4 approximate travel times from Burrinjuck Reservoir are shown. Adapted from Falk et al. (2016)

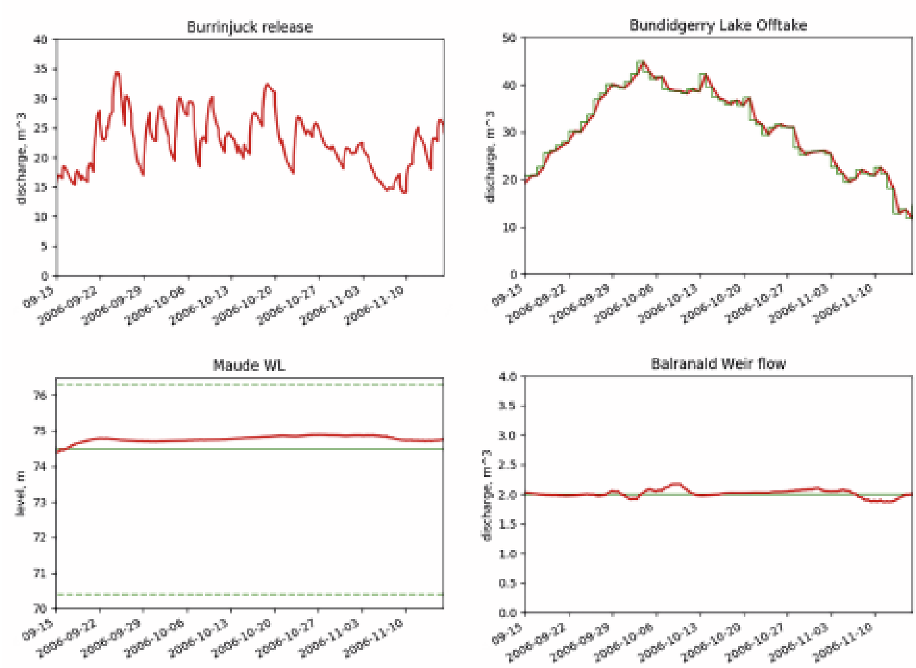

Figure 4. MPC test results: Release discharge [m3/s] at Burrinjuck Reservoir (upper left), Offtake discharge [m3/s] at Bundidgerry Lake (upper right), water level [m] at Maude (lower left), and end-of-system flow [m3/s] at Balranald (lower right). Red line: high-fidelity model results forced with MPC releases, green line: target levels/flows and irrigation water demands, green dashed line: upper and lower bounds

The MPC framework has been set up for part of the Murrumbidgee River system and tested at Level 2 in the test hierarchy. The schematisation of the surrogate model is shown in Figure 3. We use the hydrological-hydrodynamic model that was developed for the Computer Aided River Management (CARM) project (van Kalken et al., 2012) as high-fidelity model.

The operation objectives are: (1) supply ordered water to major irrigation areas and individual users, (2) minimise spills at end-of-system, and (3) keep 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 (lower and upper water levels) in the inline reservoirs. At the end-of-system at Balranald a target flow corresponding to environmental flow requirements is defined.

The MPC optimises release hydrographs at Burrinjuck Reservoir, the six inline reservoirs, and Beavers gate (see positions of controllable gates in Figure 3). In the test, only releases from Burrinjuck Reservoir are optimised, whereas historical releases from Blowering Reservoir are used. 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 about 1,100 control variables to be optimised.

The MPC controller has been tested for a two-month test period where the optimisation has been repeated every 24 hours (Madsen et al., 2017). Selected results from this test are shown in Figure 4. The figure shows the optimised release from Burrinjuck Reservoir (upper left). The MPC is able to meet irrigation water demands at the major offtakes, as shown for Bundidgerry Lake (Figure 4, upper right), as well as individual water demands along the river. At all inline reservoirs, water levels are kept close to the defined target levels and operational zones (see results for Maude in Figure 4, lower left). At the end-of-system at Balranald the flow is kept very close to the target flow of 2 m^3/s (Figure 4, lower right), essentially eliminating spills.

With respect to computational requirements, the MPC is very efficient. It takes about two minutes 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. Computational efficiency is essential for real-time implementation of the MPC, and furthermore it allows a thorough tuning and evaluation using the proposed test hierarchy. ■

References

Falk, A.K.V., Mackay, C., Madsen, H., and 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.

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.

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

Ngo, L.L., Madsen, H. and Rosbjerg, D. (2007) Simulation and optimisation modelling approach for operation of the Hoa Binh reservoir, Vietnam, J. Hydrol., 336, 269-281.

Henrik Madsen **is Research & Development Manager at DHI. He is leading DHI's research and development within water resources with a focus on emerging technologies for real-time forecasting and control of water systems. Main areas of research include data assimilation and uncertainty quantification methods in hydrological and hydrodynamic modelling, and optimisation models for water management. He is also appointed Adjunct Professor at the University of Copenhagen.**

Anne Katrine Falk **is Senior Researcher in Applied Mathematics at DHI. She has more than 15 years of experience in research and implementation of state-of-the-art mathematical methods into DHI's software portfolio. Her main field of research include model-predictive control, optimisation, data assimilation, stochastic modelling and uncertainty analysis.**