

SYSTEMS ENGINEERING FOR IRRIGATION SYSTEMS: SUCCESSES AND CHALLENGES

Iven Mareels, Erik Weyer, Su Ki Ooi
Michael Cantoni, Yuping Li, Girish Nair

*Department of Electrical and Electronic Engineering,
The University of Melbourne,
Vic 3010, Australia,
i.mareels@unimelb.edu.au*

Abstract: In Australia gravity fed irrigation systems are critical infrastructure essential to agricultural production and export. By supplementing these large scale civil engineering systems with an appropriate information infrastructure, sensors, actuators and a communication network it is feasible to use systems engineering ideas to improve the exploitation of the irrigation system. This paper reports how classical ideas from system identification and control can be used to automate irrigation systems to deliver a near on-demand water supply with vastly improved overall distribution efficiency. *Copyright ©2005 IFAC*

Keywords: large scale systems, irrigation systems, system identification, control, SCADA

1. INTRODUCTION

Over the last seven years researchers at the University of Melbourne have collaborated with Rubicon Systems Australia to achieve near on-demand water supply in open channel irrigation systems. The research and development process started by considering what appeared to be a case of hunting of local controllers in an irrigation network. The journey went through modeling and control design, including the development of new sensor and actuator hardware and pilot projects

culminating in the present industrial scale ongoing project where based on the developed ideas and technology a large irrigation system in Victoria is being fully automated. The ongoing project is reported as the largest supervisory control and data acquisition (SCADA) system of its kind in the world. Upon completion it will be the first entire irrigation system under closed loop control where water orders are met in real time on demand in as much as this is physically realisable.

This paper reports on some aspects of this journey. A corresponding paper trail can be found in the bibliography. The patents (Mareels *et al.*, 2000; Mareels *et al.*, 2001) describe the overall systems engineering philosophy in conjunction with the newly developed hardware for both measuring and controlling water in open channels.

So far the research and development effort has focused on controlling water quantity in open channels, as captured in water flow, volume and

¹ Over the last seven years, different aspects of the work reported here have been variously supported by the Cooperative Research Centre for Sensor, Signal and Information Processing; the Australian Research Council under an industry-linkage grant, and Rubicon Systems Australia, Pty. Ltd. and AusIndustry. Iven Mareels is now also working with the National ICT Australia, Ltd. that will support ongoing research in this general area under its Water Information Networks initiative.

water levels. This paper is solely devoted to this issue. In future developments it is envisaged that water quality will become a major focus for control. Sensors can assess the biological and chemical content of water to determine its quality in real time. The information infrastructure to be discussed can then be used to advantage to manage water quality as well as water quantity.

In Australia irrigation accounts for 70% of all fresh water usage (*The United Nations World Water Development Report, Executive Summary, 2003; Water and the Australian Economy, 1999; Securing our Water Future Together, 2004; Water Savings in Irrigation Distribution Systems, 2000*). This statistic varies much between different countries, and regions, but globally 70% of all fresh water is used for irrigation. The majority of irrigation is achieved through an extensive civil infrastructure of reservoirs and open canals that supplies fresh water to farms. The large scale distribution of water is powered purely by gravity. Also on most farms, a gravity fed system is in use, such that the amount of farm land that can be irrigated through the on-farm water outlet is directly limited by the available water level (potential energy) at the canal outlet, the on-farm supply point.

In the reports (*The United Nations World Water Development Report, Executive Summary, 2003; Water and the Australian Economy, 1999; Water Savings in Irrigation Distribution Systems, 2000*) overall water efficiency in irrigation, that is the ratio of volume of water used by crops and life stock to volume of water extracted from the available fresh water resource is estimated to be less than 50%. In Australia, the losses are reportedly evenly split (*Water and the Australian Economy, 1999; Water Savings in Irrigation Distribution Systems, 2000*) between the large scale distribution losses and on-farm losses. The latter are to a large extent due to poor timing of irrigation, a consequence of manual water scheduling on the supply canals. Furthermore, the on-farm losses if coupled with poor drainage lead to soil degradation (irrigation-drainage imbalance induced salination of soils). Most of the large scale distribution losses are due to the natural tendency to oversupply water, as a lack of water has obviously averse effects on the yield. Some losses are due to evaporation and canal seepage. In Australia, over supplied water simply means that water spills out at the downstream end of the irrigation system, and is no longer available to supply irrigation. The water remains of course part of the natural hydro cycle.

Policy makers in Australia have recognized that is important to re-consider present water practices in view of long term environmental and sustainability issues. Indeed, Australia is a very dry con-

tinental that presumably cannot sustain the present exploitation levels of its natural water resources. Climate change, population and industrial growth pressures compound the problem.

Here it is argued that irrigation efficiency is an ideal objective for closed loop control. Irrigation efficiency can be vastly improved by installing an information technology infrastructure: sensors and actuators linked through a supervisory control and data acquisition (SCADA) communication network. The data derived from this can be exploited using systems identification and control techniques to manage the civil infrastructure to achieve superior efficiency and support new policies in support of long term sustainable water use. Besides the improved distribution efficiency, the near on-demand water delivery that is achieved leads to significant on-farm efficiency improvements as well.

The system identification and control ideas to be discussed are quite classical, and in this way the paper illustrates the strength of the present state of the art in control and systems identification theory. Nevertheless new questions or perhaps new emphases appear and a number of interesting open problems, particular relevant to large scale systems engineering are identified.

The remainder of the paper is organized as follows. First a particular implementation of an automated, SCADA equipped irrigation system, installed in Victoria, Australia is described. The physical constraints in the irrigation infrastructure, which ultimately limit the performance that can be achieved, are touched upon. The control and management objectives are briefly outlined. System modeling is discussed next. Models derived from the St-Venant equations and grey box models are compared as to their utility/effort in the context of modeling for control and fault detection. Next model based control ideas are pursued. A variety of distributed control and centralised control strategies are briefly introduced. More importantly the impact of automation and the diminishing returns of increased control effort and control sophistication are identified. Results from the field describing the quality of the proposed models and controls are discussed.

2. IRRIGATION SYSTEMS

2.1 Canal Infrastructure

Refer to Figure 1. Through a series of open canals, water is distributed from reservoirs to farms under the driving force of the available potential energy at the reservoir. The routing and scheduling of water is realized through the adjustment of regulator structures (*gates*) placed

in the canals. These gates restrict the water flow in the canals. A stretch of canal between up- and down-stream gates is called a *pool*.

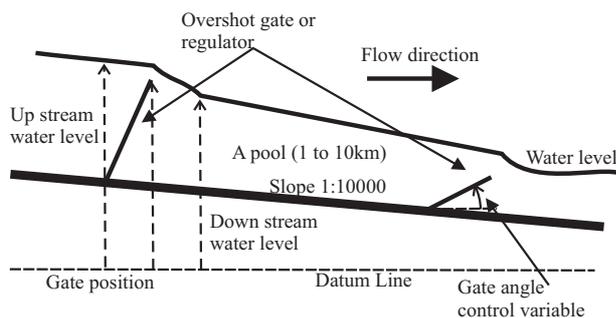


Fig. 1. Schematic representation of open canal infrastructure; vertical cross section through canal centre

A section of a canal network, is represented in Figure 2. The canal topology can be represented like a directed graph (nodes are regulators, arcs are pools, the direction is the flow). In most gravity fed systems the graph is a tree (but this is not always the case).

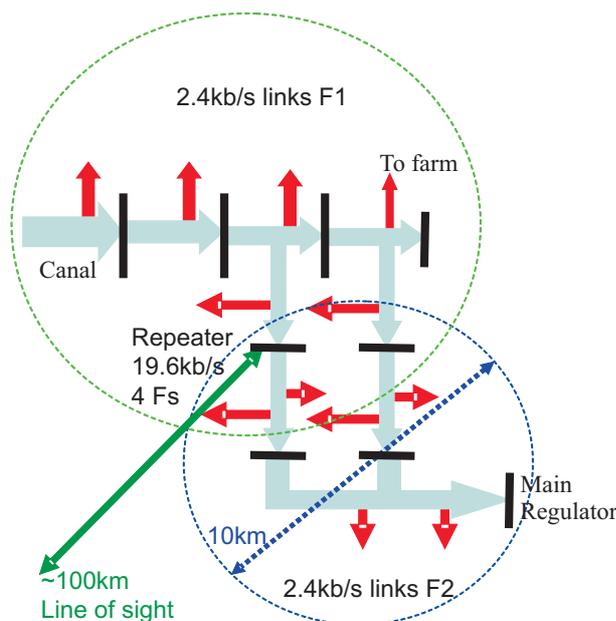


Fig. 2. Schematic representation of open canal topology (actual schematic)

The discussion here focuses on the so-called overshoot gates (FlumeGate^{TM 2}) that were developed in the course of the project. In these gates water flows over the top of the gate leaf, the angle of which can be adjusted to change the amount of water that flows over it. The gates can operate in the overflow regime, where the down stream water level is below the gate level (as on the left in Figure 1), or drowned condition, where the downstream

level is above the gate level (as is illustrated on the right in Figure 1). At the gate water levels immediately downstream and upstream of the gate are measured.

The main farm outlets are found near the downstream end of a pool. In Australia, the canals are designed for flow, and are rather shallow and narrow, with little storage in pools. A typical slope of the canal floor is 1/10,000. The canals have little free board, in most instances water will spill if the water level would reach 0.5m above the design level.

The canal and reservoir infrastructure is substantial. The replacement value (for Australia) is in the order of several tens of billion dollars in investment. To understand the scale, consider the following data for a medium size irrigation district in Australia that was used in the pilot projects described here: 500km of main irrigation canal; about 500 regulator structures in the canals; over 600 on-farm outlets; 1.1Gm² of farm land irrigated; an annual water allocation of about 0.6Gm³. The Goulburn-Murray irrigation district, the largest in Australia, is more than 10 times this size (7000km of canal; 17,000 regulators and 21,000 on-farm outlets, annual average irrigation water allocation in excess of 2Gm³) and irrigation districts in excess of 100 times this size exist in the USA, Pakistan, India and China.

The quality of the irrigation service, from a farm perspective, is determined by the timing of the irrigation water and in case the on-farm irrigation is also gravity fed, also by the water level at the on-farm outlet.

Clearly using open canals with banks that are not water tight to effect water distribution implies losses through evaporation and seepage. The combined effect of these is estimated at around 10% to 15% of supply (*Water and the Australian Economy, 1999*). These losses are insufficient to warrant consideration of replacing the canal infrastructure by a piped water network (which after all also has its own problems with leaks).

2.2 Automation Infrastructure

Figure 3 illustrates in situ the new hardware replacing the manually controlled on-farm outlets. During the course of the project, new sensors and actuators were developed which form the subject of two patent applications (Mareels *et al.*, 2000; Mareels *et al.*, 2001). The actual locations for the gates in the canals are as determined by the existing infrastructure under manual control, and no optimization of gate placement, in view of the new and automated management regime, has been considered thus far. This is an interesting aspect

² These gates are manufactured by Rubicon Systems Australia, Pty. Ltd.

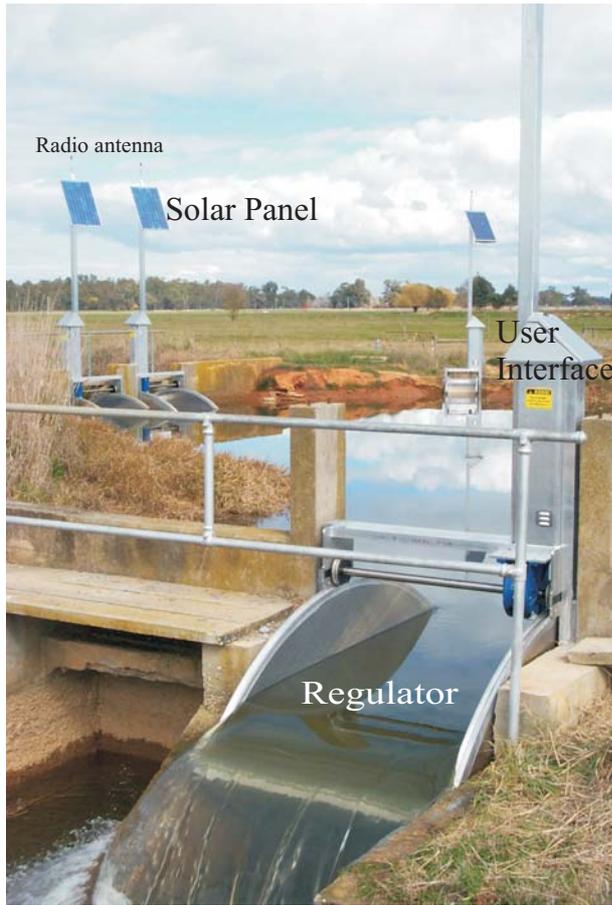


Fig. 3. Automated gates and communication infrastructure

of the overall system design that is under further investigation.

For the purpose of distributing water from reservoir to farm, all the gates both in the main canals as well as the on-farm outlets, are actuated and equipped with sensors for water levels and actual gate position (the gate angle is measured). This enables measurement of actual demand, water supplied onto farms, on the irrigation system. Actuation and monitoring functions are typically collocated, and very few sole purpose actuator or monitoring sites are in use.

Each gate/site can communicate via a cellular packet based radio network to any other site in the network. The communication network enables both broadcasting (e.g. for software upgrades) as well as peer to peer communications on all sites. A typical direct peer group contains up to 10 sites (see also Figure 2). For the above described average irrigation district there are just over a 1000 sites, the maximum number of router hops for site to site communications is 6, whereas a direct communication to the *central* node requires at most 3 router hops. For reasons of security and reliability, the network management overhead in data communication is large, and the actual net

data throughput is very low (measured in bits per second).

The required data rate for control purposes is kept to a minimum by using event driven communication and by using a particular decentralized control structure. Event based communication means that measurement data are only communicated when a water level change (a gate movement or a set point change) exceeds a predetermined threshold. Even then, the sensor data must compete with other event driven communications such as demanded by various hardware and software alarms. As a consequence the time required to successfully complete a communication cannot be guaranteed with certainty. Although and this is confirmed by experience over a few irrigation seasons, the design methods for the communication network are such that with very high confidence the maximum communication delay for a data packet can be specified. The other technique used to minimize the data rate requirements is the particular decentralized control structure that is implemented. It is such that water demand information only has to be passed on to the nearest upstream regulator, which only requires a peer to peer session, without supervisory intervention. Moreover, the central supervisory control node does not need regular sampling of sensors and/or actuator states and can use the event driven data to implement global network policies.

Locally, at the gate level, all data about sensors, actuators and hardware conditions are monitored on a regularly sampled basis, and a long history of all variables is maintained in memory.

At the central node, a global data base of all reported and requested events is maintained. In this data base, the topology of both the information network and the irrigation network is also stored. This topology is further enhanced by descriptors for all models for pools, regulators and sensors, and the corresponding control algorithms.

2.3 Objectives

From a real-time control point of view the main objectives are

- optimization of distribution efficiency, i.e. minimize the total amount of water supplied by matching supply to demand. Maximal distribution efficiency is attained when there is a zero outflow at the bottom end of the canal tree.
- real-time on-demand water supply of the requested water volumes to all irrigation outlets and;

- water level regulation at the on-farm outlet to a given set point (within an agreed dead-band limit);

The supervisory control has to ensure that the physical flow capacities are not exceeded, that the water set points are as agreed by the stake holders and that the distributed water remains within the allocations set by the regulating bodies (as modified by water trade). Clearly, the supervisory control task is in itself a dynamic object. The supervisory control system receives demand requests by the farmers through an automated ordering system. These orders are verified, if the water order is legitimate and if there is the physical capacity to supply before the water order is processed. If rejected, the water order can be rescheduled to an agreed time. The supervisory control system can suggest alternative time slots for the water order. These aspects are not discussed further. That some form of rescheduling is in general unavoidable is clear from the histogram in Figure 4. It shows the distribution of regulator sizes (in flow m^3/s) for one section of the above irrigation system. The total on-farm installed outflow capacity in this section is close to $25\text{m}^3/\text{s}$, which is a little more than 4 times the inflow capacity (the inflow capacity at the top regulator is $6.1\text{m}^3/\text{s}$). This is typical across the entire irrigation system.

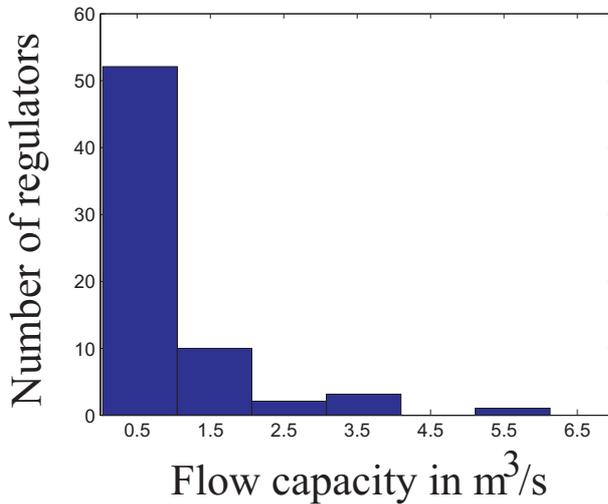


Fig. 4. Histogram of regulator sizes in m^3/s at the downstream end in a typical irrigation district

Equally important from a management perspective is to keep track of water flow and water volumes, as well as the general condition of the irrigation canals (e.g. leakage detection, evaporation losses, hardware failures).

3. MODELS FOR IRRIGATION SYSTEMS

Understanding the behaviour of the irrigation system is key to arriving at an acceptable closed loop control solution. The natural subsystems in irrigation systems are a single pool and the regulator structures. Given the topology of the irrigation system, an entire model is then but the concatenation of pool and regulator models as dictated by the system topology (see Figure 2).

3.1 St-Venant Equation Based Models

Traditionally, open channel dynamics are described using the shallow water equations, or the so called St-Venant equations (Brutsaert, 1971; Chaudhry, 1993; Cunge *et al.*, 1980). One particular representation is given by

$$\frac{\partial A_i}{\partial t} + \frac{\partial Q_i}{\partial x} = -q_{i,off}(x,t), \quad (1)$$

which models the mass balance along the length of the pool, and

$$\begin{aligned} \frac{\partial Q_i}{\partial x} + \left(\frac{gA_i}{B_i} - \frac{Q_i^2}{A_i^2} \right) \frac{\partial A_i}{\partial x} + \frac{2Q_i}{A_i} \frac{\partial Q_i}{\partial x} \\ + gA_i(S_f(Q_i, A_i, x, t) - \bar{S}_i(x)) = 0, \end{aligned} \quad (2)$$

which models the momentum balance along the length of the pool. Here $t \geq 0$ is time, and $0 \leq x \leq L_i$ is a position along the centre line of the pool, where L_i is the length of the i th pool.

In (1) and (2) $A_i(x,t)$ is the water cross section at position x along the center line of pool i at time t (it equals the local depth of water times the width of the canal). $B_i(x)$ is the width of the pool at position x , it is assumed that the cross section of the canal is rectangular. The width information can be obtained either from design drawings, or perhaps more interestingly from satellite images. $Q_i(x,t)$ is the total flow in m^3s^{-1} through the cross section $A_i(x,t)$. $q_{i,off}(x,t)$ represents the flow off-take per meter canal (in m^2s^{-1}) at position x . This term models not only a farm outlet, or an off-take into another pool/canal but also accounts for seepage and evaporation (hence the potential dependence on time, as seepage is subject to seasonal variation and evaporation experiences seasonal and diurnal variations). The $S_f(Q_i, A_i, x, t)$ represents the momentum loss per meter canal, due to friction but also water off-takes, seepage and evaporation. $\bar{S}_i(x)$ represents the slope of the bottom of pool i . The latter represents essentially the available potential energy per meter canal. The constant $g = 9.81\text{ms}^{-2}$ represents gravitational acceleration. The friction

term S_f is a nonlinear function of flow and area, for which various empirical formulas exist. Here modeled as time-invariant, although slow time variations can be expected due to effects such as bank erosion and plant growth.

The boundary conditions are the inflow into the top-end of the pool $Q_i(0, t)$ and the outflow $Q_i(L, t)$ at the bottom-end of the pool. These are to be derived from the hydraulic characteristics of the regulating structures at the up- and downstream end of the pools. For $Q_1(0, t)$ this will typically be a discharge from a reservoir, or major feed canal. $Q_i(L, t)$ for each of the *end* pools is to be minimised, ideally zero outflow is to be achieved. For the other boundary conditions $Q_i(L, t) = Q_{i+1}(0, t) = f_i(A_i(L, t), A_{i+1}(0, t), u_j(t))$, where f_i is the hydraulic characteristic for that regulator, and u_j is the control variable, which enables the variation of the flow through this regulator. Hydraulic characteristics can be determined from experimental data. Hydraulic characteristics essentially relate geometry of the regulator, as captured by the down and up-stream water conditions and the gate position to flow over the regulator.

To complete the model for the purpose of prediction an initial condition has to be provided, a suitable starting point is a steady state flow condition (zero flow for example).

The models (1) and (2) are nonlinear partial differential equations, which can be solved using a numerical integration scheme. Typically, a Preismann scheme, which involves using a discretisation in both space and time domain (Chaudhry, 1993; Cunge *et al.*, 1980) is advocated.

Despite the obvious physical interpretation that is associated with the variables, and the clear validity of the principles on which the St-Venant equations are constructed, these equations represent a considerable modeling cost when they are to be applied to a particular canal. In order for the model to become relevant to a particular canal, or pool it is necessary to fit functions like B_i , S_f , $q_{i,off}$ and \bar{S}_i . Typically this involves providing a parametric representation for these functions, and then estimating the parameters, using say non-linear least squares techniques, from a history of measurements. The mapping from the discrete values of flow Q_i and cross sections A_i at certain positions along the canal, and at certain time instances to the parameters is not trivial. One such approach is discussed in some detail in (Ooi *et al.*, 2003b). It is shown there that the St-Venant equations are indeed adequate for the purpose of modeling of open channel behaviour. This study uses data from Australian irrigation canals. An earlier study (Brutsaert, 1971) also shows the adequacy of the St-Venant equation using a similarity study under laboratory conditions.

Despite this, it is not clear that the cost of establishing an accurate quantitative model using this route is actually justified in the context of irrigation systems. Moreover, it is also not obvious that the St-Venant equations represent the most appropriate starting point for management and/or control of irrigation systems.

In a nutshell the advantage of the St-Venant equations, with the corresponding boundary and initial conditions, are:

- familiar and physically relevant variables are modeled;
- parameters have an obvious physical interpretation;
- interpolation across all time and space (simulation information not limited to discrete sample points);
- an expectation to be reliable across a broad range of operating conditions.

These have to be weighted against the obvious disadvantages

- significant cost in achieving quantitatively reliable models from discrete data;
- significant cost, perhaps prohibitive cost to simulate an entire canal or irrigation system;
- relevance beyond the operational regimes remains questionable as it cannot be verified against real data.

There is a more philosophical and more fundamental objection rooted in the way the St-Venant equations are actually used. As the St-Venant equations do not permit analytic solutions they are by necessity numerically integrated using some discretisation scheme. One could therefore argue that for all practical purposes the St-Venant equations are merely an abstraction (and perhaps a poor one). Clearly only the finite dimensional numerical approximation, a computer algorithm, is the real model. It is this computer algorithm that is tuned against the data and that forms the actual model and that predicts the actual behaviour. This begs a number of natural questions

- Can this finite dimensional and discrete time/space model not be considered as the basis for control design?
- Which discretisation scheme is the most appropriate for control design? After all only a small portion of the open-loop behaviour is targeted in closed-loop operation and there may be better schemes than the Preismann scheme.
- Perhaps the model may be further simplified using model order reduction techniques to represent only adequately the smaller closed loop behaviour (Litrico *et al.*, 2003; Litrico and Fromion, 2004).

Rather than pursuing these questions from a St-Venant equation perspective, which leads to several interesting open problems in the study of infinite dimensional systems, they are pursued directly from a data-to-model system identification perspective.

3.2 Grey box models

The above motivates the consideration of a grey box, lumped model approach to construct a model for a pool adequate for control design in an irrigation system (Ooi, 2003; Ooi *et al.*, 2005; Ooi *et al.*, 2003a; Ooi *et al.*, 2003b; Weyer, 2001; Eurén and Weyer, 2005). It should be observed that significant simplifications as compared to the full complexity of the St-Venant equations may be expected as the management regime in an irrigation system requires tight control over the water levels, despite significant variations in flow. Hence the model does not have to explain the behaviour over all possible water levels.

What is minimally required is a model that links the control input u_i which is the gate position at regulator i (a regulator may have several gates in parallel, but this issue is ignored to simplify the presentation) to water levels and flows at the critical points along the canals, typically the up- ($y_{i,u}$) and down-stream ($y_{i,d}$) water levels at the regulators. With reference to Figure 5 consider a canal where all pools are simply in line. To further simplify the exposition, assume that the regulators are overshoot gates in free-flow conditions. In this situation, the downstream water levels at the regulators are largely irrelevant. The ideas can be applied mutatis mutandis to other situations as well, including the more complicated situations with under shot regulators and drowned overshoot gates (Eurén and Weyer, 2005).

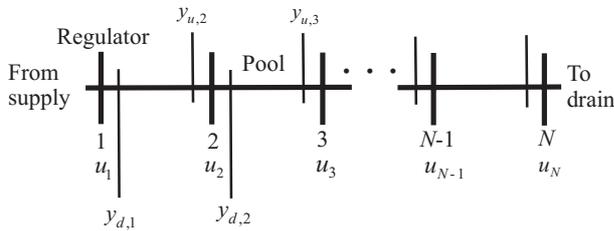


Fig. 5. In-line canal system, a series of regulators and pools

The flow Q_i over the regulator i is a static function determined by the geometry of the regulator, which is varied by the gate position u_i and the supply water level $y_{u,i}$ at the regulator site. This relationship is the hydraulic characteristic H_i for the regulator,

$$Q_i = H_i(u_i, y_{u,i}). \quad (3)$$

Besides the hydraulic characteristic, a model for the dynamics of the pool is required. This model is in its simplest form a mass balance, the volume of water in the pool is increased by the flow over the up-stream regulator and decreased by the flow over the down-stream regulator. In (Weyer, 2001; Mareels *et al.*, 2000; Mareels *et al.*, 2001) a model is proposed that captures conservation of mass, and the first mode of standing waves in a pool:

$$\begin{aligned} p_i \left(\frac{d}{dt} \right) y_{u,i+1}(t) = & - \sum_j q_{i,j} \nu_{i,j}(t) \\ & + c_i H_i(u_i(t - \tau_i), y_{u,i}(t - \tau_i)) \\ & - c_{i+1} H_{i+1}(u_{i+1}(t), y_{u,i+1}(t)). \end{aligned} \quad (4)$$

A continuous time model is used, as in general regular sampling is not available. This is therefore the preferred embodiment of the model. Although for identification purposes it is clear that the above model only requires local information at the upstream and downstream gate of the pool. In this situation it is possible to obtain regularly sampled and synchronized data series of the water levels and gate positions for any particular pool. In (4) $y_{i,u}$ is the water level upstream of regulator i ; τ_i is a time-delay. It is associated with the fastest wave phenomena in the pool between gates i and $i + 1$ (which can be linked to the critical velocity in the pool (Chaudhry, 1993)). An action on the upstream end of the pool is observed at the downstream end of the pool after τ_i time units.

The model dynamics are captured via the polynomial $p_i(\xi)$. For short pools, up to a few km in length, a polynomial of degree 3 suffices to achieve a high fidelity model. For such pools, the roots of $p_i(\xi)$ consist of an integrator and a weakly damped oscillatory pair. The integrator captures the idea of water storage in the pool. The oscillatory roots capture the dynamics associated with the dominant standing wave(s) in the pool. The terms $\nu_{i,j}$ represent water off-takes between regulator i and $i + 1$ and the $q_{i,j}$ are static gains. (Most off-takes are placed near the down-stream end of the pool, which simplifies the model, and ensures better quality of service for the farmer, as the down stream water level can be regulated more easily than the up stream water level.)

Typical order of magnitudes for the various parameters in the Australian irrigation districts where field trials have been conducted thus far are:

- time delay is approximately 3 min per km
- dominant wave period is approximately 9 min per km
- achievable (closed loop) dominant time constant is approximately 10 to 15 min per km

The parameters appearing in (4) are readily obtained via standard linear system identification techniques. This can be done both from open loop experiments as well as from closed loop data (Ooi *et al.*, 2005; Ooi and Weyer, 2001; Ooi *et al.*, 2003a; Ooi *et al.*, 2003b; Weyer, 2001). The latter is critically important when considering a large scale irrigation system.

The grey box model (4) compares favourably with the St-Venant model (1) and (2) in that it has much lower complexity and yet similar or better prediction accuracy (see below and Figure 6).

Validation tests across both low and high flow regimes have been completed. In Figure 6, adapted from (Ooi *et al.*, 2003b), a comparison is made between the St-Venant based models and the grey box models. The Figure 6 represents a set of open-loop simulated model responses compared against the actual field measurements. The inputs to the simulation are the regulator positions over time (using the measured positions), and the initial water level condition. The data set for the modeling and validation are from different irrigation seasons.

The least accurate model, which is still very useful, is the St-Venant model where the model parameters are not tuned to the actual pool but guesstimated based on a rule of thumb. The response is the dashed line in Figure 6. This model is by no means disgraced. The other St-Venant model has its parameters (a model for the friction was adjusted) optimally selected based on field observations. Its response is dash-dot line in Figure 6. These models are compared against a grey box model that uses a third order model, (leaky) integrator and standing waves and a first order grey box model that only captures the (leaky) integrator trends. The grey box models have been identified from open-loop experiments. All models are reasonably accurate. From a simulation perspective, the best model, based on quality for effort, is the third order grey box model. For control purposes the first order grey box model will suffice, when the wave phenomena are appropriately considered in the closed loop design (Weyer, 2003a).

3.3 Automated Model Building

Given the size of a typical irrigation system automated model building and model maintenance is essential. Model maintenance must be performed within the constraints imposed by normal management conditions. As demonstrated in (Ooi and Weyer, 2001) the grey box models (4) allow for simple calibration of the parameters under closed loop conditions. A few hours of data, using permitted changes in the water level reference set

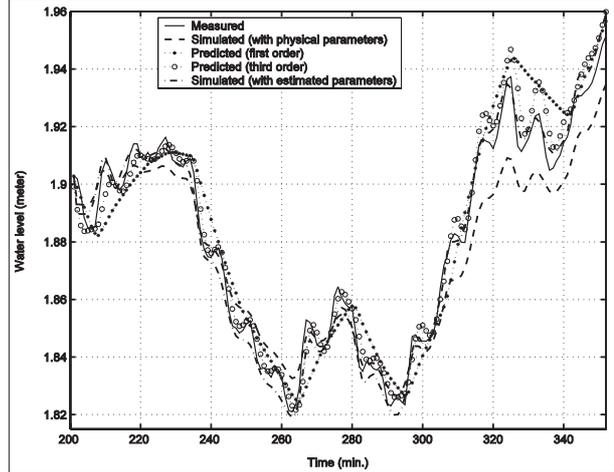


Fig. 6. Experimental validation of grey box and St-Venant based models

points, exploiting the available dead band, suffices to tune the model parameters. Such variation is allowed under normal operating conditions as the typical dead band for the set point regulation is of the order of 10cm (under manual operation, water levels were deemed acceptably regulated when less than 30cm away from set points).

Figure 7 taken from (Ooi and Weyer, 2001) compares a validation data set (black solid line) with a simulation (dotted line) based on a model derived from closed loop data. Only 90 minutes of closed loop data with minimal set point variations were used to derive the relevant parameters of a first order grey box model (4) that sufficed for control design purposes.

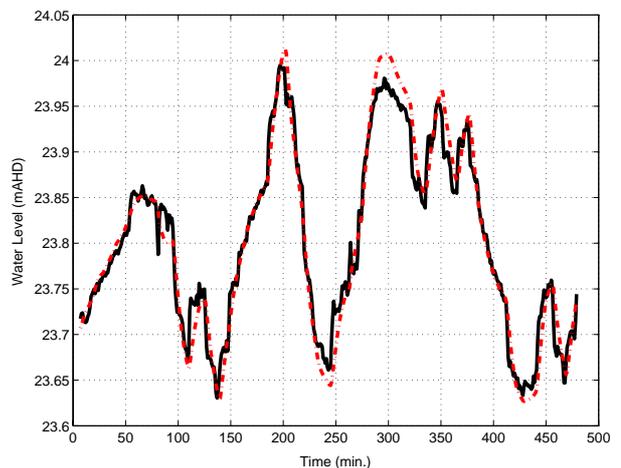


Fig. 7. Validating closed loop system identification, using grey box models

It is important to notice that the tuning of the model for a particular pool is inherently a decentralized task. It involves only data gathered from the up-stream and down-stream regulators. Given the available compute engines at the various monitoring sites a model update can be performed locally using peer-to-peer communications. Once

the modeling or model maintenance is completed, the new model parameters can be communicated to the central node, which requires very little communication resources, as it only involves a few coefficients to be updated.

In case of hardware failures (Choy and Weyer, 2005), like a regulator break-down, or sensor failure, an alarm is raised that also triggers a model update (and if necessary a control law update) to reflect the new dynamics.

A model for an entire irrigation system is constructed from the logical interconnection of the individual pool models.

3.4 An Open Modeling Issue

The apparent data-based validation of both a St-Venant partial differential equation model and a low order grey box model within the present context poses an interesting and open question in model order reduction (Litrico *et al.*, 2003; Litrico and Fromion, 2004). *Under what conditions and in what sense can the grey box model (4) be considered as an approximation for the coupled partial differential equation model (1) and (2).*

This question is not only of academic interest. The basic parameters in the grey box model have a physical meaning, the dominant time constant, the wave period and time-delay are all related to the behaviour of the St-Venant equations. If the relationship between the St-Venant equations and the grey box model is better understood it may be possible to derive a simple map, a rule of thumb, linking physical parameters like pool length, critical velocity, manning coefficient (an important parameter describing the friction losses in open channel flow) to a reasonably first guess for the coefficients in the grey box model. At present such a map is essentially being constructed from experience. An answer to the above open problem would provide a sound theoretical basis for such a rule of thumb.

The question is not trivial in that even taking into account that the water levels are reasonably tightly controlled the St-Venant equations (2) remain highly nonlinear due to their nonlinear dependence on the flow. Hence classical model order reduction techniques do not apply.

The situation is a nonlinear version of the classical lumped parameter Pi and T-section approximations used in electrical power engineering to represent a transmission line. Transmission line equations are linear St-Venant equations.

3.5 On-line Model Uses

The need for a low complexity model structure like the grey box model cannot be argued solely on the basis of simulation or control design requirements. Indeed simulation and most control design techniques only use the model in an off-line capacity in which complexity plays a lesser role. (The use of model predictive control methods has not as yet been pursued, but a low complexity model is definitely an advantage in this case.)

The model is used in the supervisory control to determine the limits of physical capability of the irrigation system to meet demand. Off-line an approximate table can be constructed that indicates which demand can be met in real time, but it is simpler and more accurate to actually calculate from the present state of the model forward to verify that the physical constraints can be satisfied (the present scheduling policy works on a first come first served basis).

The model has however far more important on-line uses, which justify a low order complexity model. The main on-line use of the model is in fault detection and fault isolation. The single most valuable use of the model in this respect is the identification in real-time of leaks, which are identified and localized, as a (significant) difference between water levels predicted by the model and actually achieved water levels as measured by the sensors.

The development of fault detection methods that exploit the full capabilities of the model and sensory data is ongoing.

4. CONTROL FOR IRRIGATION SYSTEMS

In the following discussion, consider a series of in-line pools (see Figure 5), and assume overshoot gates at the regulators. . Other approaches, not discussed here, for control of canals can be found in (Halleux *et al.*, 2003; Schuurmans *et al.*, 1999; Malaterre and Baume, 1998).

From a control perspective, and to simplify the discussion, the predictive models discussed above can be further simplified. Indeed, as the wave dynamics of the pools are not to be excited, because the free space in the canals above the water level is not designed to cope with standing waves, this part of the model may be ignored for control design as long as the design ensures that the waves are not going to be excited. The main wave frequency from the model is thus an important input into the control design process, determining the cut-off frequency of a low pass filter applied to the control input.

With this simplification, and introducing appropriate scaling and some abuse of notation, the pool dynamics, represented in continuous time, are given by

$$\dot{y}_{u,i+1}(t) = u_i(t - \tau_i) - u_{i+1}(t) - d_i(t). \quad (5)$$

The index i refers to the regulator i , $i = 1, \dots, N$. The u_i are the input variables, which are actually nonlinear functions of the gate positions and the water levels at the regulator. This representation globally linearises the non-linear behaviour of the canal; although the inputs are constrained in that $u_{\max,i} \geq u_i \geq 0$. The disturbance $d_i(t) \geq 0$ is the total water off-take from pool i (upstream of regulator $i + 1$ and downstream from regulator i). As discussed above, it is (almost) a measurable disturbance, although some aspects of the off-take term cannot be measured directly. Evaporation, seepage and unscheduled off-takes caused by bank erosion or bank collapse are clearly difficult to measure directly. The evaporation and seepage are only a small fraction of the total off-take, and the unscheduled off-takes can be identified from the model.

The control design is to ensure that $u_N = 0$ (no outflow at the bottom regulator) and that all $y_{u,i}$ $i = 2, \dots, N$ are regulated to their reference values $y_{u,i}^{ref}$, $i = 2, \dots, N$ (within a deadband). In this case setpoint regulation automatically delivers water on demand, provided that in steady state demand is within the capacity constraint, that is $d_i(t) + u_{i+1}(t) < u_{\max,i}$ for all $i = 1, \dots, N - 1$.

4.1 Decentralised Control

A decentralized control design is discussed in detail in (Weyer, 2002) see also (Li *et al.*, 2005). The main features are that the upstream regulator input u_i $i = 1, \dots, N - 1$ is used to regulate the water level $y_{u,i+1}$ $i = 1, \dots, N - 1$ based on information that is available at the down stream regulator $i + 1$ and the local regulator i as well as the measured off-take on the pool i . The controller can thus be implemented using a peer-to-peer communication session. u_N is set to zero.

Assuming the off-take is measurable, the decentralized control law takes the form:

$$u_i = T_{ff,i}(s)(d_i + u_{i+1}) + T_{f,i}(s)(y_{u,i+1}^{ref} - y_{u,i+1}) \quad (6)$$

Here $u_N = 0$ and $i = 1, \dots, N - 1$. $T_{ff,i}$ is the feed forward control part and $T_{f,i}$ implements local feedback.

This leads to the closed loop being described by

$$u_i = \frac{sT_{ff,i}(s) + T_{f,i}(s)}{s + T_{f,i}(s)e^{-s\tau_i}}(d_i + u_{i+1}) \quad (7)$$

with $u_N = 0$ and $i = N - 1, N - 2, \dots, 1$ and

$$y_{u,i+1} - y_{u,i+1}^{ref} = \frac{T_{ff,i}(s)e^{-s\tau_i} - 1}{s + T_{f,i}(s)e^{-s\tau_i}}(d_i + u_{i+1}). \quad (8)$$

This representation shows clearly how disturbances propagate upstream; d_i affects all $y_{u,k} - y_{u,k}^{ref}$ for $k = 1, \dots, i + 1$. The decentralised design is by necessity a trade-off between how much the local disturbance d_i is rejected from the water level $y_{u,i+1}$ and how these disturbances propagate upstream, the effect on $y_{u,k}$ for $k = 1, \dots, i$. The transfer function governing the upstream propagation of disturbances is

$$P_i(s) = \frac{sT_{ff,i}(s) + T_{f,i}(s)}{s + T_{f,i}(s)e^{-s\tau_i}}, \quad (9)$$

the local rejection is governed by

$$R_i(s) = \frac{T_{ff,i}(s)e^{-s\tau_i} - 1}{s + T_{f,i}(s)e^{-s\tau_i}}. \quad (10)$$

The critical element in these transfer functions is the delay.

A similar argument can be developed for when the off-take is not measurable. In this case, the control law takes the form

$$u_i = T_{ff,i}(s)u_{i+1} + T_{f,i}(s)(y_{u,i+1}^{ref} - y_{u,i+1}), \quad (11)$$

starting with $u_N = 0$. The corresponding closed loop is described by the iteration ($i = N - 1, \dots, 1$)

$$u_i = P_i(s)u_{i+1} + T_i(s)d_i, \quad (12)$$

which describes the regulator inputs and

$$y_{u,i} - y_{u,i}^{ref} = R_i(s)u_{i+1} + S_i(s)d_i, \quad (13)$$

which describes the error-propagation. Here

$$T_i(s) = \frac{T_{f,i}(s)}{s + T_{f,i}(s)e^{-s\tau_i}}, \quad (14)$$

and

$$S_i(s) = \frac{-1}{s + T_{f,i}(s)e^{-s\tau_i}}. \quad (15)$$

4.2 Decentralised control, without off-take measurements

When the off-takes may not be measured, the feedback $T_{f,i}(s) = H_i(s)/s$ must include an integrator in order to reject the off-take disturbances d_i . This makes the control design much more challenging as the plant also contains an integrator.

In the presence of a double integrator and having to contend with the unavoidable delay in P_i , R_i , S_i and T_i , it can be shown that is now

impossible to avoid that water levels errors are amplified as they propagate upstream along the canal. In (Li *et al.*, 2005) it is demonstrated that for any linear (stable) design $\|P_i(s)\|_\infty > 1$. The way to arrive at a good design is to achieve strong disturbance rejection by way of ensuring that $\|R_{i-1}(s)P_i(s)\|_\infty \ll 1$. In this way, even very long chains of pools can be adequately controlled.³

The design proceeds in two phases, first stability/robustness is guaranteed for the local pool dynamics (design emphasis on S_i); a phase margin of about 40 degrees and a gain margin of about -10dB is considered adequate. A feedback consisting of a PI controller, with low pass filter as to avoid the excitation of the wave dynamics, i.e. $T_f = \frac{K_P s + K_I}{s(T_w s + 1)}$, can achieve this goal. The feed forward term is essentially a low pass filter, again to ensure that waves are not excited. Moreover, the feed forward term must be designed with care to avoid the amplification of disturbances along the canal. To complete the design, it must be verified against the entire system dynamics, this boils down to the non-trivial verification of the transfer functions $R_{i-1} \Pi_k P_k$.

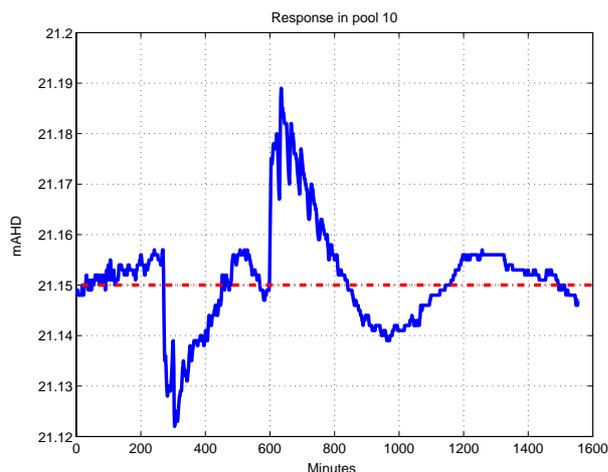


Fig. 8. Decentralised control in action, PI with feedforward

Figures 8 and 10 illustrate the response of the decentralized controller (PI+filter) under a control test scenario. These graphs display field data from the Haughton Main Canal. The set point is the dash-dot line and the solid line is the recorded water level realized by the controller. Several set point changes and disturbances occur over the period of time (a bit more than a day's worth of

³ Worst case scenario is when all pools are identical, as in this situation the consecutive amplifications are re-enforcing each other. This is a rather unusual situation in practice, and one that can always be avoided by adjusting the location of regulators, or better by adding a regulator structure.

data is displayed). The controller maintains the water level to better than 5cm.

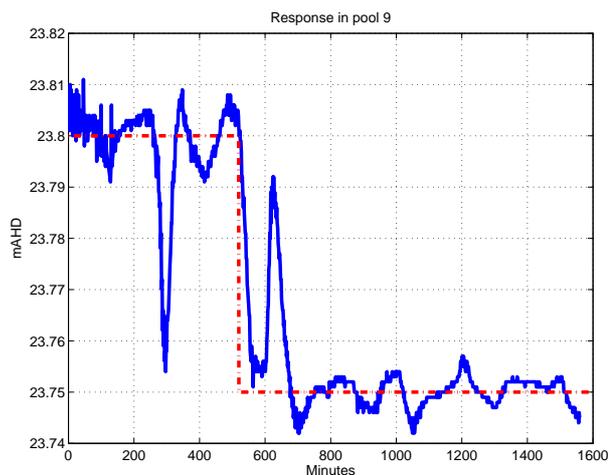


Fig. 9. Decentralised control in action, PI with feedforward

4.3 Decentralised control with off-take measurements

Assuming that the off-takes are measured, it is possible to use a simple lead-lag controller in the feedback path ($T_{f,i}$) and a simple low pass filter in the feed forward path. The design is somewhat simplified in this situation as compared to the design based on the integrator in the feedback path. The only relevant transfer functions are R_i and P_i , which makes it indeed somewhat easier to achieve acceptable gain and phase margins as well as disturbance rejection. The main design problem concerns the feed forward action, as this dominates the error propagation along the canals. As the design is only marginally easier than the PI based design, the PI based design, which has greater robustness with respect to off-takes, is to be preferred.

4.4 Centralised Control

Based on the model (5) linear quadratic methods as well as H_∞ design methods may be pursued to deliver a global controller see (Weyer, 2003c; Li *et al.*, 2004; Li *et al.*, 2005). These serve as benchmarks for the decentralized controllers.

The information infrastructure briefly discussed above does not enable one to implement a central control strategy across an entire irrigation district. Linear quadratic, centralised controllers have been field tested in smallish irrigation systems. Besides serving as a benchmark for the decentralized controllers, a centralized control law can be used to derive an acceptable decentralized controller, through pruning the weaker input-output connections.

To complete an adequate linear quadratic design it is essential to augment the state of the system (water levels y_i) with the integrals of the regulation errors ($\int_0^t (y_i^{ref} - y_i)$) and states to describe high pass filtered input variables (emphasizing the frequency content above the standing wave frequency of the pool that can be effected by the particular input variable). The latter are essential to avoid excitation of standing waves in the pools. A linear quadratic design completed along these lines provides an overall responsiveness better than what can be achieved using the above decentralized controller, with acceptable robustness properties. The design process is however much more demanding. In particular the selection of the quadratic weights so as to achieve an intuitively acceptable response is a time-consuming task.

The intuition derived from the decentralized controller can be used to better advantage in an H_∞ design. The decentralized controller can be used to deliver the loop shaping that is required to achieve a high gain in the low frequency band (to reject the load disturbances), a 20dB/dec roll off around the cross over frequency (which is limited by the wave frequency) to achieve sufficient phase margin, and faster roll-off around and beyond the wave frequencies. Using the decentralized controller to guide the selection of the loop shaping frequency dependent weights, allows the H_∞ design method to deliver a very robust and high performing controller that satisfies all control requirements. The resulting controller is a global controller, with a performance similar to a well designed linear quadratic controller. The advantage is that the design effort is much smaller than what is required for the linear quadratic controller, and entirely comparable to the decentralized controller. Field tests are planned for the near future.

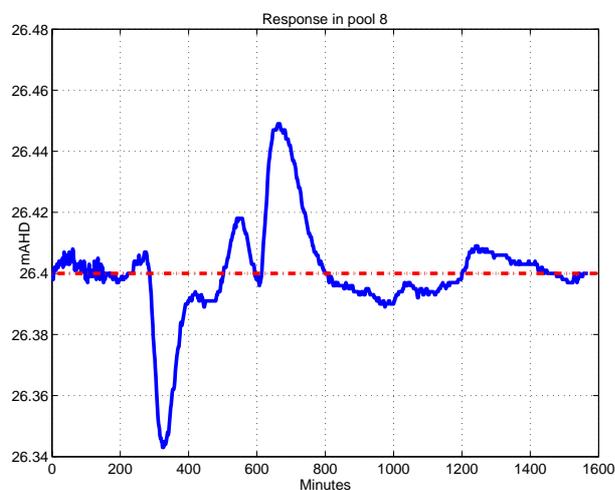


Fig. 10. Decentralised control in action, PI with feedforward

4.5 A law of diminishing returns

In order to appreciate the benefit that may be derived from the various control scenarios, a comparison is made between

- manual control, the present situation, where water orders are averaged, and water regulation is effected once every 24h by manually adjusting the regulator structures.
- decentralized control, without feed forward; PI+filter design using a simple anti-windup scheme.
- decentralized control, with feed forward; PI+filter design using a simple anti-windup scheme.
- global, centralized control (using a linear quadratic regulator).

The comparison is based on a 3-pool simulation, where the scenarios have been implemented on an equal footing. All control implementations encounter identical circumstances. The scenario stretches over a nine day period. It involves:

- A water order cancellation. Half-way through the simulation, a water order (45h long order of 37500 m³) is cancelled. Farmers can order water, typically 4 days in advance, but are allowed to cancel the order with only one hour notice. This is a very demanding scenario from a manual regulation point of view, but realistic.
- A 2.5cm rain event that last for 10h. Under rain conditions, farmers are allowed to stop irrigation. This leads to an emergency stop of the irrigation system, which is virtually impossible to control well under manual operations.
- Ten normal water orders, of various length and magnitude. The water orders are organized as per typical manual operations, so as to create a situation advantageous for manual operations.

The scenario is illustrated in Figure 11, in which the head over the most upstream regulator is indicated. This illustrates the water supply into the canal, which is the main control action to implement the down stream requested water orders.

The results are that any of the automated methods achieve all the water orders as requested and avoid spilling water during the rain condition. The decentralized control strategies achieve 100% water efficiency in that no water is spilled over the most down stream regulator. The Linear Quadratic Regulator controller achieves quicker and better water level regulation (barely noticeable) and an overall efficiency of 99%. (The LQ regulator sacrifices a minimal amount of water in favour of achieving a better water level regula-

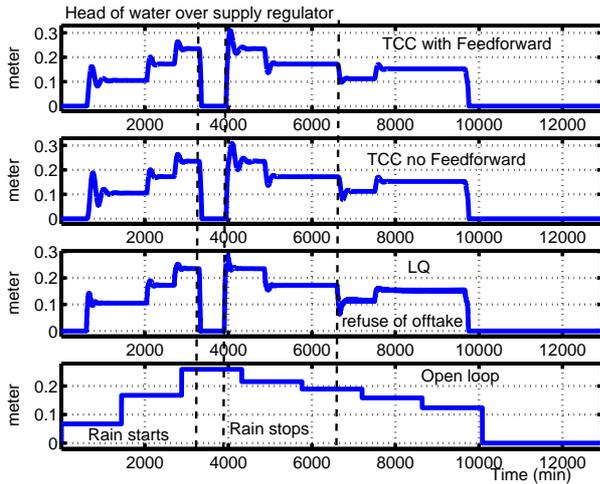


Fig. 11. Water supply over a 9 day period into an irrigation canal

tion.) There is clearly no real incentive to implement centralized control action.

The corresponding water level response in the most down stream pool is illustrated in Figure 12. In this pool it is most difficult to manage acceptable water regulation under manual control. Typically the bottom end of the irrigation system experiences poor quality of service under manual operations. Using the decentralized controllers there is saturation during the rain event (regulators are closed) and at the end of the irrigation period, small excess of water supply, but no outflow. This water is stored in the most down stream pool.

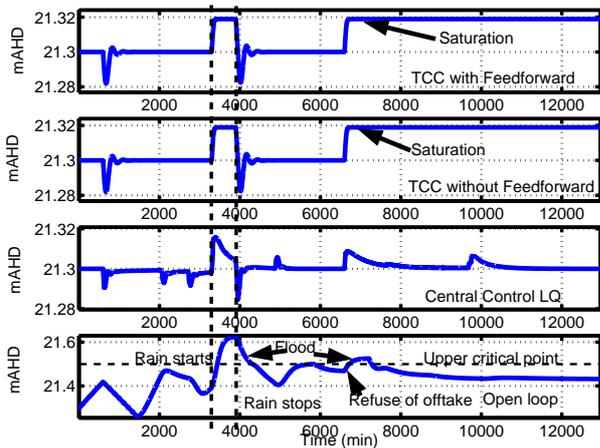


Fig. 12. Waterlevel response in the most down-stream pool

Under manual operations, the overall efficiency was 80% (which is very good compared to historical averages, typical best practice manual operation achieves an efficiency of about 70%). The quality of service is however poor, despite the fact that the water orders were placed as to be advantageous to manual control. Only 90% of the water orders is met. The water levels dropped below

the minimum required to satisfy the requested demand over about 10% of the total duration. Moreover, during the rain period, the most down stream pool flooded its banks for about a day. The rain period only lasted for 10h, but there was a significant water flow imbalance at the beginning of the rain period, which together with the slow response of manual control, aggravated the rain effect to cause significant flooding.

To observe substantial differences between the different forms of automated, real-time, irrigation management, larger canal systems must be compared. Various scenarios have been simulated, and it quickly becomes clear that decentralized control with feed forward (and anti-windup) achieves near global optimal performance. In most scenarios its performance cannot be distinguished from centralized control (which cannot be implemented) and it is substantially better than the decentralized control without feed forward. The latter suffers from a larger variation in water levels, and cannot always supply water orders as requested because water levels fall below the agreed minimum supply point.

4.6 Implementation Issues

There are two main implementation issues to contend with: sampling and saturation.

As pointed out, the information infrastructure does not allow for regular sampling. Globally sampling is performed on an event basis (e.g. only the time that a water level changes by 2cm up or down is reported), although locally, at each site regular sampling is performed. The decentralized controller does not suffer, as it only uses local information. A global controller must be designed as to cope with rather varying information rates and delays across the range of variables it requires to compute the control inputs. Substantial further research work is necessary. This issue is at the heart of so called network control. An alternative approach to communication limited control is exposed in for example (Nair *et al.*, 2004) where an information theoretic point of view is pursued. The latter identifies the minimum data rates required to achieve stability. It turns out that in the case at hand, as the instability in the plant is limited to a chain of integrators that any non-zero communication rate suffices to achieve stability. Of course, this does not imply good performance, but it indicates that the required communication cannot be demanding. This is in-line with the present engineering solution and experience.

Saturation issues can be addressed using anti-windup control ideas. The design can be completed, and consequently adapted as to ensure

smooth feedback without integrator saturation (Zaccarian *et al.*, 2004a; Zaccarian *et al.*, 2004b). Simple anti-windup control is implemented.

4.7 Open Issues

There are a number of interesting open aspects from a theoretical point of view that have received little attention in the control literature. Event based sampling and a better understanding of the disturbance propagation in the irrigation system. The latter is made difficult because of the combination of marginally stable open loop dynamics combined with delays. Although no single pool dynamics is dominated by delay effects, the concatenation of many similar systems in a tree like structure makes this an interesting and difficult problem, see (Li *et al.*, 2004; Li *et al.*, 2005) for first results in the case of a string of identical pools. This work addresses the worst case scenario. What is important is to understand the *average* case, where there is a certain (natural) distribution of delays and pool dynamics. The key ingredient in understanding the propagation of disturbances through the system is to understand the role of the feed forward control in balancing the requirement for providing information to speed up the response, and yet limiting this information so as not to amplify the disturbances.

Another interesting open problem that has received not enough attention in the control literature is to determine an appropriate structure for a decentralized controller. This is a problem of determining the appropriate information structure for control design, or which inputs/outputs are best paired with each other. As, the number of different decentralized controller structures is combinatorially large, exhaustive evaluation is simply not feasible, hence relevant heuristics should be developed. One ad-hoc solution method is to use sequential minimum variance (Weyer, 2003b), based on a particular ordering of the variables to be regulated. Another is to start from a global centralized control design, and to prune the input-output connections based on an appropriate singular value decomposition. In the present application, both approaches lead rather quickly to the proposed solution method. In sequential minimum variance, the ordering of the variables that leads to the above decentralized controller structure is to place the highest priority on the outflow (u_N , next the most down stream water level y_{N-1} , then the next water level y_{N-2} and so on.

5. DISCUSSION

Systems engineering principles as discussed here have been implemented in support of irrigation

system management in a number of pilot projects in Victoria, Australia, with excellent outcomes in terms of achieved overall water savings and quality of service: water distribution efficiency of better than 90% (essentially only the unremovable losses due to evaporation and seepage remain), compared to 70% achieved under manual operations; combined with an ability to meet water requests on-demand in better than 90% of all water orders. All of this is achieved together with much better regulation accuracy, which improves the quality of service to the customers in a significant way. On the basis of these outcomes, Rubicon Systems Australia Pty. Ltd. has been commissioned to start large scale industrial implementation in Victoria.

The white paper on water released by the State Government of Victoria identified that the implementation of advanced control, along the lines expounded here, will save about 400Mm³ of water per year when completed across all of the irrigation districts in Victoria. This volume of savings nearly equals the total annual water consumption for all domestic and industrial use in the State of Victoria. This figure is based on a conservative extrapolation of the actually achieved water savings in the pilot projects to date (only need to achieve half the success achieved in the pilot projects thus far). Similar studies conducted in a New South Wales report even greater saving potentials (*The Business of Saving Water - the Report of the Murrumbidgee Valley Water Efficiency Project*, n.d.). The economic benefits that can be derived from such water savings initiatives are immense. The (*The Business of Saving Water - the Report of the Murrumbidgee Valley Water Efficiency Project*, n.d.) report indicates economic benefits in the order of \$500M when a water saving target of 1Gm³ would be realised.

An important outcome of the information infrastructure is that management, and indeed the community at large, can be much better informed about the capabilities of the irrigation system; how much water was used, when and where it was used. Already, changes in the way irrigators use water can be observed, as a more responsive distribution system enables them to use more water efficient on-farm practices. This results in even more water savings or better economic returns for the farmers.

The present positive results encourage the expansion of the research and development ideas to consider other aspects of water distribution, on-farm as well as urban. The outcomes realized are a small step towards the goal set forth in the (*UNESCO Water Report 2003*) that fresh water management must happen on the scale of entire water catchments.

REFERENCES

- Brutsaert, W. (1971). The Saint-Venant equations experimentally verified. **97**, 1387 – 1401.
- Chaudhry, M. H. (1993). *Open-channel flow*. Prentice-Hall.
- Choy, S. and E. Weyer (2005). Reconfiguration scheme to mitigate faults in automated irrigation channels. In: *44th IEEE Conference on Decision and Control*. Seville, Spain. submitted.
- Cunge, J.A., F.M. Holly and A. Verwey (1980). *Practical Aspects of Computational River Hydraulics*. Pitman Advanced Pub. Program.
- Eurén, K. and E. Weyer (2005). System identification of irrigation channels with undershot and overshot gates. In: *16th IFAC World Congress*. Prague, Czech Republic. to appear.
- Halleux, J., C. Prieur, J.-M. Coron, B. d'Andréa Novel and G. Bastin (2003). Boundary feedback control in networks of open channels. *Automatica* **39**(8), 1365 – 1376.
- Li, Y., M. Cantoni and E. Weyer (2004). Design of a centralized controller for an irrigation channel using loop-shaping. In: *Proc UKACC Control conference*. University of Bath, UK.
- Li, Y., M. Cantoni and E. Weyer (2005). On water-level error propagation in controlled irrigation networks. submitted for publication.
- Litrico, X. and V. Fromion (2004). Analytical approximation of open-channel flow for controller design. *Applied Mathematical Modelling* **28**(7), 677 – 695.
- Litrico, X., V. Fromion, J.-P. Baume and M. Rijo (2003). Modelling and pi control of an irrigation channel. In: *Proceedings of the European Control Conference ECC03*. Cambridge, UK.
- Malaterre, P.O. and B.P. Baume (1998). Modeling and regulation of irrigation canals: existing applications and ongoing researches. In: *Proceedings of IEEE Conference on System, Man and Cybernetics*. San Diego, USA. pp. 3850 – 3855.
- Mareels, I., E. Weyer and D. Aughton (2000). Rubicon research pty ltd, control gates. WIPO Number WO02/16698, PCT Number PCT/AU01/01036 priority date for Australia 21 Aug 2000.
- Mareels, I., E. Weyer and D. Aughton (2001). Rubicon systems australia pty ltd, fluid regulation. WIPO Number 02/071163, Australian Application Number 2002233060, PCT/AU02/00230, priority date for Australia 2 March 2001.
- Nair, G., R. Evans, I. Mareels and B. Moran (2004). Topological feedback entropy and nonlinear controllability. *IEEE Trans Aut Control, Special Issue on Networked Control Systems* **49**(9), 1585 – 1597.
- Ooi, S. K., E. Weyer and M. C. Campi (2003a). Finite sample quality assessment of system identification models of irrigation channels. In: *Proceedings of the IEEE Conference on Control Applications*. Istanbul, Turkey.
- Ooi, S. K., M.P.M. Krutzen and E. Weyer (2003b). On physical and data driven modelling of irrigation channels. In: *Proceedings of the 13th IFAC Symposium of System Identification*. Rotterdam, The Netherlands. pp. 1975 – 1980.
- Ooi, Su Ki (2003). Modelling and control design of irrigation channels. PhD thesis. Department of Electrical and Electronic Engineering, The University of Melbourne.
- Ooi, Su Ki and E. Weyer (2001). Closed loop identification of an irrigation channel. In: *Proceedings of the 40th IEEE Conference on Decision and Control*. Orlando, USA. pp. 4338 – 4343.
- Ooi, Su Ki, M.P.P. Krutzen and E. Weyer (2005). On physical and data driven modeling of irrigation channels. *Control Engineering Practice* **13**(4), 461 – 471.
- Schuermans, J., A. Hof, S. Dijkstra, O.H. Bosgra and R. Brouwer (1999). Simple water level controller for irrigation and drainage canals. *Journal of Irrigation and Drainage Engineering* **125**(4), 189 – 195.
- Securing our Water Future Together* (2004). [//www.dse.vic.gov.au/dse](http://www.dse.vic.gov.au/dse). report prepared on behalf of the Department of Victoria Land and Water Management Policy.
- The Business of Saving Water - the Report of the Murrumbidgee Valley Water Efficiency Project* (n.d.). <http://www.napsqw.gov.au/publications/pratt-water.html>. report prepared by Pratt Water Pty Ltd.
- The United Nations World Water Development Report, Executive Summary* (2003). <http://www.unesco.org/water/wwap/>.
- Water and the Australian Economy* (1999). <http://www.atse.org.au/publications/reports/water9.htm>. report prepared on behalf of the Australian Academy of Technological Sciences and Engineering.
- Water Savings in Irrigation Distribution Systems* (2000). report prepared by Sinclair Knight & Merz.
- Weyer, E. (2001). System identification of an open water channel. *Control Engineering Practise* **9**, 1289 – 1299. based on a prior paper in IFAC Symposium on System Identification, 2000.
- Weyer, E. (2002). Decentralised pi controller of an open water channel. In: *Proceedings of the 15 IFAC World Congress*. Barcelona, Spain.
- Weyer, E. (2003a). Control of open water channels. *IEE Trans on Control Systems Technology*. submitted.
- Weyer, E. (2003b). Controller configurations for irrigation channels derived from sequential

- minimum variance control. In: *In Proceedings of IEEE Conference on Control Applications*. Istanbul, Turkey.
- Weyer, E. (2003c). Linear quadratic control of an irrigation channel. In: *Proceedings of the 42nd IEEE Conference on Decision and Control*. Hawaii, USA. pp. 750 – 755.
- Zaccarian, L., E. Weyer, A. Teel, Y. Li and M. Cantoni (2004a). Anti-windup for marginally stable plants applied to open water channels. In: *Proceedings of the 5th Asian Control Conference*. Melbourne, Australia. pp. 1702 – 1710.
- Zaccarian, L., Y. Li, E. Weyer, M. Cantoni and A. Teel (2004b). Anti-windup for marginally stable plants applied to open water channels. *Control Engineering Practice*. submitted.