

Use of an artificial neural network to capture the domain knowledge of a conventional hydraulic simulation model

Zhengfu Rao and Fernando Alvarruiz

ABSTRACT

As part of the POWADIMA research project, this paper describes the technique used to predict the consequences of different control settings on the performance of the water-distribution network, in the context of real-time, near-optimal control. Since the use of a complex hydraulic simulation model is somewhat impractical for real-time operations as a result of the computational burden it imposes, the approach adopted has been to capture its domain knowledge in a far more efficient form by means of an artificial neural network (ANN). The way this is achieved is to run the hydraulic simulation model off-line, with a large number of different combinations of initial tank-storage levels, demands, pump and valve settings, to predict future tank-storage water levels, hydrostatic pressures and flow rates at critical points throughout the network. These input/output data sets are used to train an ANN, which is then verified using testing sets. Thereafter, the ANN is employed in preference to the hydraulic simulation model within the optimization process. For experimental purposes, this technique was initially applied to a small, hypothetical water-distribution network, using EPANET as the hydraulic simulation package. The application to two real networks is described in subsequent papers of this series.

Key words | artificial neural network, hydraulic simulation model, POWADIMA, replication, water distribution

Zhengfu Rao (corresponding author)
School of Civil Engineering and Geosciences,
University of Newcastle,
Newcastle upon Tyne NE1 7RU,
UK
Tel.: +44 1793 816359
Fax: +44 1793 812 089
E-mail: raoz@halcrow.com

Fernando Alvarruiz
Departamento de Sistemas Informáticos y
Computación,
Universidad Politécnica de Valencia,
Valencia 46022,
Spain

INTRODUCTION

Need for predictive technique

Amongst other things, the first paper of this series on the POWADIMA (Potable Water Distribution Management) research project (Jamieson *et al.* 2007) describes the characteristics of a real-time, near-optimal control system for water-distribution networks. These include the requirement of being able to predict the consequences of changes in the control settings on the network's performance, to determine whether a particular combination is capable of meeting the demands without violating any of the operational constraints. Thereafter, it is a matter of selecting the best combination of feasible control settings so as to minimize operating costs up to the operating horizon, which is the topic of the next paper in this series (Rao &

Salomons 2007). However, this specific paper is restricted to the problem of predicting the consequences of different pump and valve settings for different demand patterns and starting conditions. As such, it forms the first constituent part of the DRAGA-ANN (Dynamic Real-time Adaptive Genetic Algorithm – Artificial Neural Network) control system for the efficient and effective operation of water-distribution networks.

Required attributes

Including optimization in the overall control system requires the predictive technique selected to be computationally efficient in order to keep the computing time within reasonable limits. This is especially important in the case of

real-time operational control where there is a need to adjust the control apparatus at frequent time intervals, since the demand for water is highly variable throughout the day. The predictive technique also needs a high degree of accuracy since some errors, such as those associated with the predicted water level in storage tanks, tend to accumulate over the operating period. As a consequence, it is possible that a different control strategy could be selected which is neither optimal nor for that matter feasible. The third attribute required is that the predictive technique has to be robust. Water demand forecasting is notoriously prone to spectacularly large errors, which initially have to be accepted until such time they can be corrected. Similarly, SCADA (Supervisory Control And Data Acquisition) facilities are capable of introducing measurement errors relating to the current state of the network. Whatever method of prediction is selected, it has to be robust enough to cope with all these different types of errors without 'crashing'.

MODELLING WATER DISTRIBUTION NETWORKS

Empirical models

The hydraulic behaviour of water-distribution networks in general, and their hydraulic dynamics in particular, are time-variant, spatially distributed and highly non-linear which traditionally have been represented either by means of empirical models, as for example, mass-balance, or process-based models, such as hydraulic simulation. In this context, mass-balance models consist of weighted functional relationships between storage, flows and pumping-station discharges. The weights associated with the functional relationships can be determined using linear regression (Sterling & Coulbeck 1975) or from linearization of the nonlinear network (Fallside & Perry 1975). The main advantage of mass-balance models is that the system response can be determined much faster than, say, from a simulation model. They are, however, more appropriate for regional water-supply schemes rather than distribution networks.

Instead of using a simple mass-balance model, the nonlinear nature of the network hydraulics could be more

accurately represented by using a set of nonlinear regression equations. Information required to construct such a model can be obtained in a variety of ways. For example, regression curves can be generated by repeated executions of a calibrated simulation model for different tank-storage levels and loading conditions or by using information from actual operations to form a database relating pump head, pump discharge, tank levels and network demands (Tarquin & Dawdy 1989). Regression models have the advantage of being able to incorporate some degree of non-linearity while providing a time-efficient mechanism for estimating the network response. However, regression curves and databases only contain information for a particular network over a given range of demands. If the network changes appreciably or forecasted demands are outside the range of the database, such an approach provides erroneous results.

Simplified network-hydraulics models

As an intermediate step between empirical models and full hydraulic network simulation, a simplified network-hydraulics model could be considered. In such cases, the network hydraulics are approximated using a macroscopic network model or analysed using a series of linearized hydraulic equations. Macroscopic models represent the pipe layout by use of a highly skeletal network model. Typically only a pump, lumped resistance term (a pipe) and an aggregated demand are included. Both DeMoyer & Horowitz (1975) and Coulbeck (1988) used macroscopic models that had multiple terms relating the effect of various network components, incorporated within a single equation. In certain cases, such as where the boundary conditions are essentially independent of pumping-station discharge, it might be possible to represent the network hydraulics by means of a simple linear model. Jowitt & Germanopoulos (1992) used an approximate linear model for a network dominated by aggregated hydraulic heads, in which small variations in tank-storage levels did not significantly affect pump operations. In a similar application, Little & McCrodden (1989) developed a simple linear model for a supply network in which the hydrostatic head in the controlling tank was held constant. The coefficients for both types of models can be determined following extensive

analyses. As a result, such models are site-specific and have to be judged on that basis to determine their acceptability.

Network simulation models

Network simulation models provide the capability to model the nonlinear dynamics of a water distribution network by solving the governing set of quasi-steady-state hydraulic equations. For a water-distribution network, the governing equations include the conservation of mass and the conservation of energy. In contrast to both empirical and simplified network-hydraulics models, network simulation models are very adaptive to both physical changes and spatial variations in demand. However, although network simulation models are usually more robust than either empirical or simplified network-hydraulics models, they generally require a considerable amount of data in their formulation. They also need significantly more effort to calibrate them properly. In addition, since such models need extra computational effort, they are generally restricted to applications that require the minimum number of individual evaluations.

Over the past thirty years or so, a significant investment has been made in developing generic hydraulic simulation models. Software packages such as WESNET, INFOWORKS, GINAS, STONER/SynerGEE, EPANET, AQUIS and WATNET, have been widely used in recent years for a number of purposes, ranging from planning and design to operational analysis and the development of control strategies for water-distribution networks. However, their use for real-time, near-optimal control is somewhat impractical for large water-distribution networks because of the computational burden optimization imposes. If it were necessary to run a simulation model for each iterative change of pump/valve settings, it is more than likely that the optimal setting would not be found before the next update was due.

Alternative approach

In recent years, significant progress in the fields of nonlinear pattern recognition and system control theory have been made possible through advances in a branch of nonlinear modelling called artificial neural networks (ANN). An ANN

is a nonlinear mathematical structure, which is capable of representing arbitrarily complex, nonlinear processes that relate the inputs and outputs of any system. Mathematicians have shown that multi-layer, feed-forward ANNs have the necessary capability to be a 'universal function approximator'. In their landmark papers, Kolmogorov (1957); Sprecher (1965), and Lorentz (1976) proved the existence of universal function mappings based on simple mathematical structures. Later, Hecht-Nielsen (1987) showed that a three-layer, feed-forward ANN meets the requirements to be a universal function approximator and that any multivariate function can be approximated by an ANN having only a finite number of nodes in the hidden layer. This result is referred to as the 'Kolmogorov Mapping Neural Network Existence Theorem'. Subsequent studies based on the above-mentioned work have shown that a three-layer, feed-forward ANN using sigmoid transfer functions can implement any continuous and bounded multivariate function mapping (Funahashi 1989; White 1990; Blum & Li 1991; Ito 1992; Takahashi 1993).

Methodology selected

For the purposes of real-time operational control, it has already been mentioned that the model selected to predict the consequences of different control settings needs to be computationally efficient, highly accurate and numerically robust. Based on these criteria, hydraulic simulation models may have the necessary accuracy and robustness but lack the computational efficiency for large-scale networks. Empirical and simplified network-hydraulics models are likely to be computationally more efficient but their accuracy and robustness are questionable. ANNs are computationally efficient and robust but require copious numbers of input/output patterns during training and testing, for which it is impractical to use the real network. Therefore, the approach adopted in the POWADIMA research project has been to combine two of these techniques by using an ANN to replicate a conventional hydraulic simulation model of the network. In this way, the complex knowledge base of the hydraulic simulation model is captured in a far more computationally efficient form. Thereafter, the ANN is used in preference to the conventional hydraulic simulation model for predicting

network performance, as part of the optimal-control process.

Whilst this approach is not entirely new, all other known applications to date have related to either planning or design exercises rather than operational control. Examples of the former include the design of groundwater remediation schemes (Rogers & Dowl 1994; Rao & Jamieson 1997) and regional wastewater-treatment planning (Wang & Jamieson 2002), where replicating complex simulation models by an ANN has significantly reduced the computational burden of the optimization process. However, computational efficiency is far more important for operational control in general and real-time operational control in particular, owing to the short time increment between successive updates. The other main advantage of this approach is the high degree of realism, which is imparted by the hydraulic simulation model prior to the replication by the ANN. It will of course be appreciated there is an implicit assumption that, in the first instance, the real network can be accurately modelled using a conventional hydraulic simulation model.

REPLICATION OF A HYDRAULIC SIMULATION MODEL

Artificial neural networks

An ANN is a mathematical representation of interconnected computing elements (or neurons) arranged in layers, which process information by their response to external inputs, in an analogous way to the central nervous system. The attractiveness of ANNs is their potential to learn from input-output data sets and their ability to approximate any continuous non-linear function to any arbitrary degree of accuracy, using a feed-forward process. In the case of a three-layer ANN(I, J, K) which is shown in Figure 1, the input layer has I neurons, the hidden layer has J neurons and the output layer has K neurons, with the network being fully connected between adjacent layers.

Each hidden neuron j receives input from every neuron i in the input layer. Moreover, each input (x_i) is associated with a weight (w_{ji}^h) so that the effective input (Ω_j) to node j is

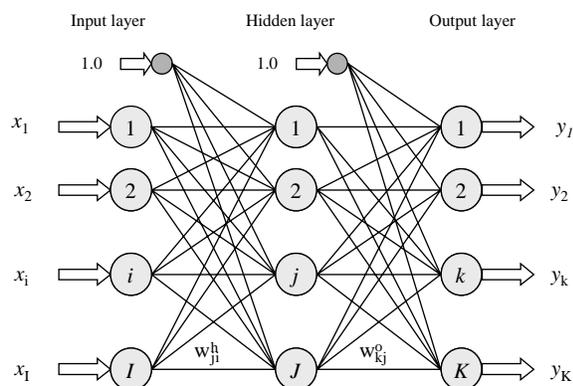


Figure 1 | Architecture of an artificial neural network.

the weighted sum of all the inputs:

$$\Omega_j = \sum_{i=0}^I w_{ji}^h x_i \quad (1)$$

where x_0 and w_{j0}^h are referred to as the bias ($x_0 = 1.0$) and the bias weights, respectively. The effective input, Ω_j , is passed through a nonlinear activation function (sometimes called a transfer function or threshold function) to produce the output (h_j) of the node. The most commonly used activation function is the sigmoid function. The characteristics of a sigmoid function are that it is bounded above and below, is monotonically increasing and is continuous and differentiable everywhere. The sigmoid function generally employed for ANNs is the logistic function:

$$h_j = f(\Omega_j) = \frac{1}{1 + e^{-\Omega_j}} \quad j = 1, 2, \dots, J \quad (2)$$

in which Ω_j can vary on the range $\pm \infty$, but h_j is bounded between 0 and 1. The corresponding output neuron, y_k , is given by

$$y_k = f \left\{ \sum_{j=0}^J w_{kj}^o f \left\{ \sum_{i=0}^I w_{ji}^h x_i \right\} \right\}, \quad k = 1, 2, \dots, K \quad (3)$$

where w_{ji}^h is a weight between the i th input neuron and the j th hidden neuron, w_{kj}^o is a weight from the j th hidden neuron to the k th output neuron, and $f(\cdot)$ is a sigmoid function as defined by Equation (2).

Developing a neural network comprises two major steps, namely 'training' (or learning) and 'testing' (or verification). During the training process, combinations of

known input–output data ('training' sets) are repeatedly presented to the ANN and the weights associated with each neuron (w_{ji}^h and w_{kj}^o) are adjusted until the specified input provides the desired output. Through these adjustments, the ANN 'learns' the correct input–output response behaviour. This training process is usually accomplished by using some particular algorithm in which a cost function, specified as the sum of squared errors between the true output and the output produced by the network, is minimized. When the cost function approaches a minimum, the network is considered to have converged. The minimization of the cost function can be achieved in different ways. The most popular technique is the back-error propagation algorithm proposed by Rumelhart & McClelland (1986). After training, the ANN is then subjected to the verification stage in which other combinations of known input–output data are introduced ('testing' sets) in order to estimate the residual error. Based on the performance of the trained ANN, further adjustments may appropriate to make the model more accurate and/or robust.

Mapping a hydraulic simulation model using an ANN

In the context of operational control, a conventional hydraulic simulation model of a water-distribution network can be described in discrete time as an input–output system, which is depicted in Figure 2.

Here:

\mathbf{P}_t – a vector of control variables representing n_p pump settings at time t ;

\mathbf{V}_t – a vector of control variables representing n_v valve settings at time t ;

\mathbf{D}_t – a vector of variables representing values of demands at the n_d consumer nodes in the network between times t and $t + \Delta t$;

\mathbf{S}_t – a vector of variables representing n_s storage tank water levels at time t ;

\mathbf{E}_t – a vector of variables representing n_p power consumption of pumps during the interval between times t and $t + \Delta t$;

\mathbf{H}_t – a vector of variables representing values of pressures n_h at specific nodes during the interval between times t and $t + \Delta t$;

\mathbf{Q}_t – a vector of containing the values of flow n_q for specific pipes during the interval between times t and $t + \Delta t$;

$\mathbf{S}_{t+\Delta t}$ – a vector of variables representing n_s storage tank water levels at time $t + \Delta t$.

The most widely researched and used ANN structure is the multi-layer, feed-forward network which is ideally suited for modelling input–output relationships. For the purposes of optimal operational control, an ANN can be regarded as a mapping function between an input and output set. In this particular instance, the input set contains the combination of pump/valve setting, demands and initial water levels in storage tanks whilst the output set corresponds to the power consumption of pumps, resulting water levels in storage tanks, pressures and flow rates at critical locations throughout the network. Using a feed-forward ANN(I, J, K), $I = n_p + n_v + n_d + n_s$ is the total number of input values, $K = n_p + n_h + n_q + n_s$ is the total number of output values and J is the number of hidden neurons to be identified during the training and testing process. The value of J is usually found using a strategy of progressively adding neurons to the hidden layer until no further worthwhile improvement in error reduction is achieved.

Training and testing procedure

Before the replication process can be initiated, it is first necessary to apply the process-based hydraulic simulation package to the distribution network in question. Thereafter, critical points within the network have to be selected as these will act as operational constraints for pressures and flow rates in the optimization process. The next stage involves running the process-based model in steady-state mode with different combinations of input values (initial

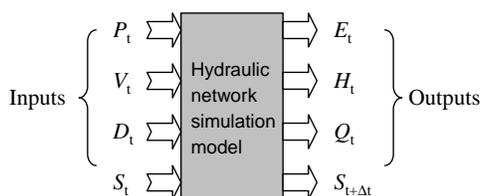


Figure 2 | Input–output model of a water-distribution network.

water levels in the storage tanks, demands for the various district-metering areas, pump settings and valve settings) to ascertain their effects on the output values at the next time-step (water levels in storage tanks, hydrostatic pressures at critical nodes in the network, flow rates in critical pipes and power consumption for the various pumps). The range of demands used in these evaluation runs should exceed the values expected in practice so that the trained ANN will be capable of predicting all possible eventualities. Similarly, bearing in mind that the optimization process (genetic algorithm) needs to encounter some infeasible solutions, the range of initial water levels should exceed the physical dimensions of each storage tank. Depending upon the simulation package chosen, it is possible that some may have simple, in-built operating rules such as switching off the pumps when the storage tank is full. However, for the purposes of real-time, near-optimal control, the process-based model is only required to predict the consequences of the input values, as opposed to taking an operational decision. Therefore, artificially increasing the tank size also eliminates this potential problem.

The number of training sets required depends on the size and complexity of the distribution network but is usually measured in terms of thousands. Therefore, a computer program has been specifically written to generate matching input/output sets automatically with random inputs, all data being normalized. With regard to testing, again it is a matter of experience in deciding the number of additional sets required but typically, this might be about 20% of the size used for training. Training the ANN is performed by adjusting the weighting factors between the neurons in the input, hidden and output layers. The most common criterion used to measure the goodness-of-fit between the predicted and 'observed' is the root mean squared error (*RMSE*). If the error between the conventional hydraulic simulation model and the trained ANN is deemed unacceptable for whatever reason, then the whole training procedure can be repeated, possibly using more training sets and/or modifying the ANN's structure by changing the number of neurons in the hidden layer. A schematic representation of the entire replication process for capturing the knowledge base of a hydraulic simulation model by means of an ANN is given in Figure 3.

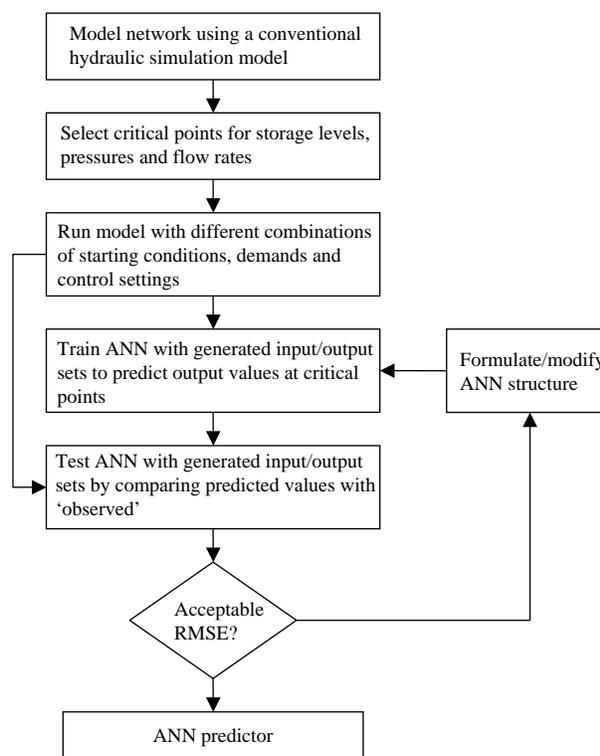


Figure 3 | Replication process for capturing knowledge base of hydraulic simulation model by means of an ANN.

APPLICATION TO A HYPOTHETICAL WATER-DISTRIBUTION NETWORK

Any Town (Modified) water-distribution model

Rather than embark directly on the two case studies, the precaution was taken to first experiment with a somewhat simpler distribution network so as to gain experience of the difficulties that might be encountered in applying the methodology to a real, complex network. To that end, the hypothetical 'Any Town' network (Walski *et al.* 1987) was chosen for this purpose since it has the distinct advantage of being well-documented as a result of having been extensively modelled previously. However, in order to make the task a little more challenging for the subsequent optimization stage, a number of modifications have been introduced including the addition of an extra storage tank and extending the pipe network in the upper-left portion. The opportunity was also taken to convert the measurements to SI units. The resulting network is referred to as the 'Any Town (Modified)' network or AT(M). As a result of

these changes, the AT(M) network has a total of 41 pipes, arranged to form 19 nodes, with 3 fixed-speed pumps and 3 storage tanks, as shown in Figure 4. Most of the pipe lengths and diameters are as per the original network, apart from those subsequently added, which have been given appropriate values within the same range. The precise details can be obtained from the corresponding author, if required.

The AT(M) network has been modelled using the well-known EPANET hydraulic simulation package (Rossman 2000). Again, the selection was made on the grounds that EPANET is fully documented and readily available. Since AT(M) is a hypothetical network without any real data on which to calibrate and verify the model, pre-assigned values for parameters such as the roughness coefficient have been used, consistent with the original network. The time-step adopted for the simulation was 1 hour, representing a compromise between what was desirable from an operating standpoint and the computational burden imposed in the subsequent optimization stage. In generating the corresponding input/output sets for training and testing the ANN, the hydrostatic pressure at 3 arbitrarily chosen nodes and the water levels in the 3 storage tanks were designated as the critical points in the network.

Developing an ANN for the AT(M) network

Several different approaches were tried during the development of a suitable ANN for the AT(M) network,

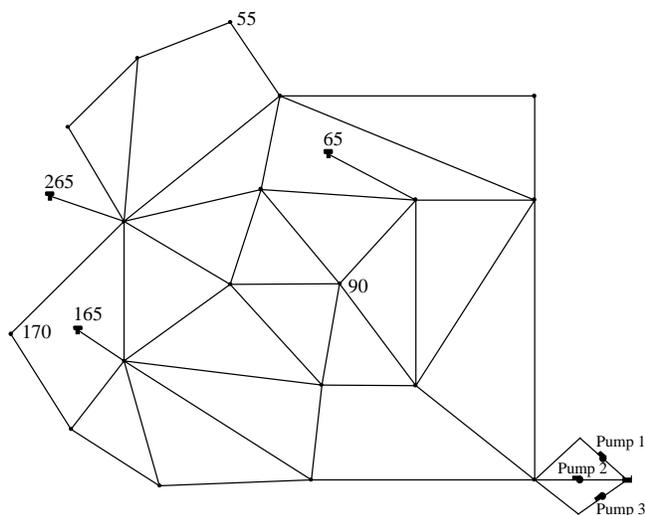


Figure 4 | The Hypothetical Any Town (Modified) water-distribution network.

including the use of a separate ANN for each of the output variables. In the event, it was found that this was not necessary as the accuracy of the individual predictions were no better than using one composite ANN for predicting all the output values at the same time. Therefore, attention was directed at estimating the appropriate number of neurons in the hidden layer of a composite ANN which comprised 5 input values (the first input representing the number of pumps on duty; the second being the aggregated demand for the 19 consumer nodes and the remainder comprising the 3 initial water levels, one for each of the storage tanks) and 7 output values (1 for total pumping power consumption, 3 for hydrostatic pressures at nodes 170, 90 and 55, together with 3 resultant water levels, one for each of the storage tanks 65, 165 and 265, respectively). On a trial-and-error basis, a good representation of the EPANET model was achieved with 20 neurons in the hidden layer, giving a final structure of ANN(5,20,7).

Having decided the appropriate number of neurons in the hidden layer, the next issue to be addressed was the required number of ANN training sets necessary to achieve an accurate representation of the EPANET model relating to the AT(M) network. To that end, different numbers of training sets were used to ascertain the impact on the *RSME* between the predicted and 'observed' output values. It can be seen from Figure 5 that, starting with random initialization of parameters, the *RSME* converges rapidly to 1.65% in approximately 2000 iterations, with little or no improvement thereafter, regardless of the number of training sets

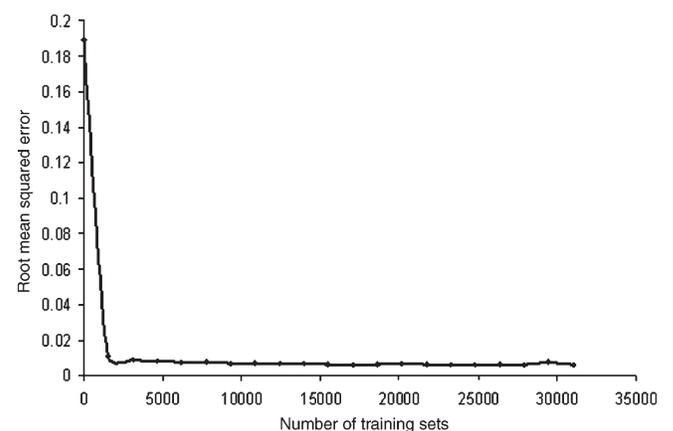


Figure 5 | Relationship between *RMSE* and number of ANN training sets.

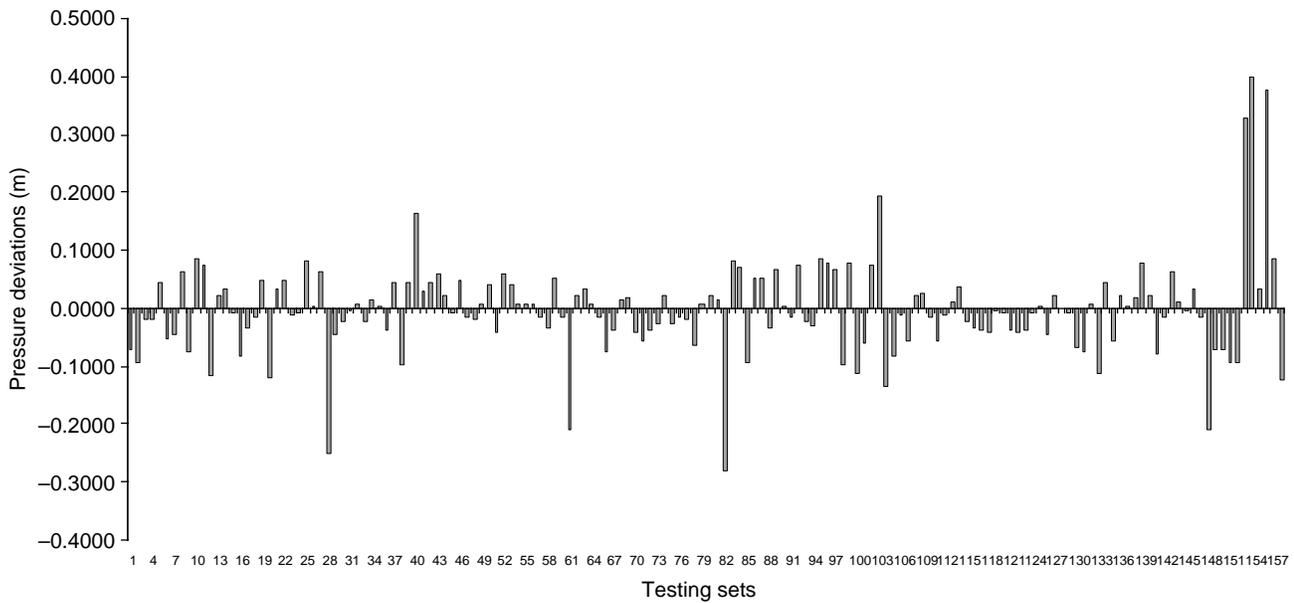


Figure 6 | Comparison of EPANET and predicted ANN results for pressures.

used. The accuracy of replicating the EPANET model by an ANN(5,20,7) is shown in Figures 6 and 7 for the resulting hydrostatic pressures and storage tank water levels for a typical sequence of testing sets over a 24-h period.

Inclusion of water-quality considerations

In addition to the hydraulic variables, an attempt was made to include water quality in the decision-making process by

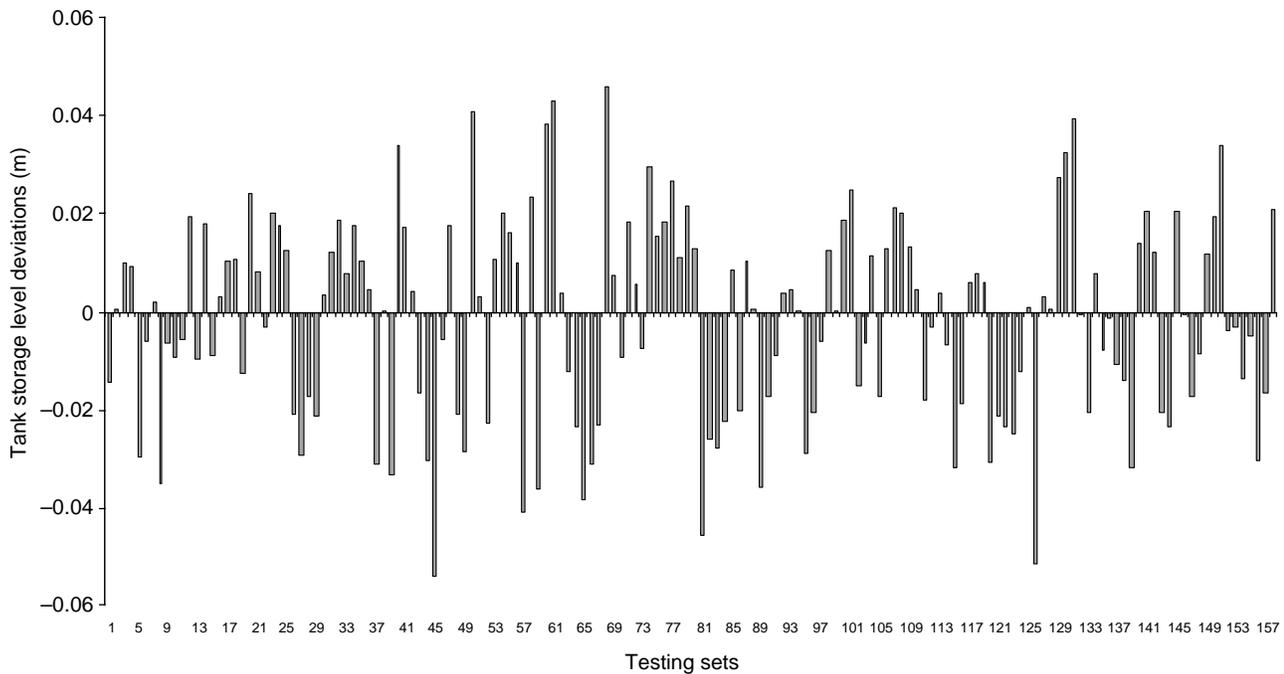


Figure 7 | Comparison of EPANET and predicted ANN results for storage tank water levels.

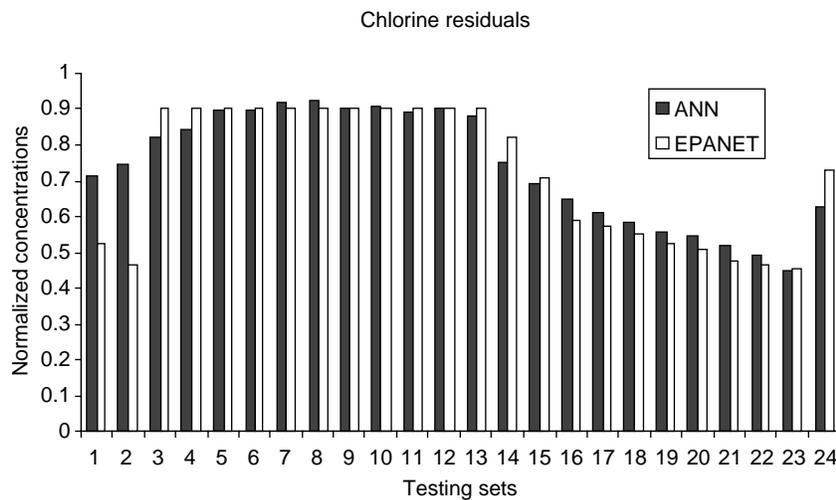


Figure 8 | Comparison of EPANET and predicted ANN results for chlorine residuals.

adopting the same approach of using an ANN to replicate an EPANET model of the AT(M) network which included chlorine residuals. The selection of chlorine residual as the water-quality determinand was an obvious choice since water-distribution engineers endeavour to maintain a minimum prescribed level of concentration to combat any pollution arising from infiltration or re-growth within the network. As previously, a large number of corresponding input/output sets were generated, this time including chlorine residual as an output variable for critical points within the network. Thereafter, an ANN was trained and tested. However, the results were less impressive than those for the hydraulic variables. Firstly, it took an inordinate amount of time for the water-quality simulation model of the network to stabilize and give consistent values for chlorine residuals, even for a very simple network such as AT(M). Secondly, it would seem that the chlorine-residual values are more a function of the size of the network rather than the control settings. If that is the case and there is only a weak relationship with the control settings, then it is hardly the fault of the ANN when somewhat disappointing results are obtained (Figure 8). An alternative but less satisfactory way of indirectly helping to maintain chloride-residual concentrations in the network, which has subsequently been used in this research project, was to assign minimum flow rates to critical pipes where low velocities occur and treat them as operational constraints

in the near-optimal control process, thereby avoiding the possibility of stagnation.

CONCLUSION

Up until recently, hydraulic simulation models have been the only means available to represent the complex, non-linear behaviour of water-distribution networks. However, in predicting the dynamic consequences of different control settings in relation to the initial conditions and short-term fluctuations in demand, they have their limitations owing to the computational burden they impose. Nevertheless, it would seem plausible that there should be a substantial opportunity for improving computational efficiency if the hydraulic simulation model could be approximated by an input/output relationship which, in this case, would be mapped using a multivariate function. With that in mind, this paper advocates the use of an artificial neural network and shows that at least for a simple, hypothetical water-distribution network, a conventional hydraulic simulation model can be replicated with a high degree of accuracy. The computational advantage of doing so is an average 10-fold reduction in the time taken to predict the consequences of different control settings in comparison with a conventional hydraulic simulation model. This computational improvement is expected to increase with the complexity of the

distribution network (see [Martinez et al. 2007](#); [Salomons et al. 2007](#)), thereby enhancing the prospect of making real-time, near-optimal control a practical reality.

ACKNOWLEDGEMENTS

The POWADIMA research project was funded by the European Commission under its Vth Framework thematic programme on Energy, Environment and Sustainable Development (Contract Number EVK1-CT-2000-00084). The authors would like to take this opportunity to thank the Commission and project officers for their support throughout the duration of the contract.

REFERENCES

- Blum, E. K. & Li, K. I. 1991 [Approximation theory and feedforward networks](#). *Neural Nets* **4**, 511–515.
- Coulbeck, B. 1988 A review of methodologies for modelling and control of water supply. In *Computer Applications in Water Supply vol 2 Systems Optimization and Control* (ed. B. Coulbeck & C. H. Orr), Research Studies Press Ltd, Letchworth, UK, pp. 80–109.
- DeMoyer, R. & Horowitz, L. 1975 *Computer Applications in Water Supply* Lexington, MA. Research Studies Press Ltd. Letchworth, UK.
- Fallside, F. & Perry, P. F. 1975 Hierarchical optimisation of a water-supply network. *Proc. IEE* **122** (2), 202–208.
- Funahashi, K. I. 1989 [On the approximate realization of continuous mapping by neural networks](#). *Neural Net* **2**, 183–192.
- Hecht-Nielsen, R. 1987 Komogorov's mapping neural network existence theorem. In *Proc. Int. Conf. on Neural Net.*, vol. **3**, pp. 11–13. IEEE. New York.
- Ito, Y. 1992 [Approximation of continuous functions on \$R^n\$ by linear combinations of shifted rotations of a sigmoid function without scaling](#). *Neural Nets*. **5**, 105–115.
- Jamieson, D. G., Shamir, U., Martinez, F. & Franchini, M. 2007 Conceptual design of a generic, real-time, near-optimal control system for water distribution networks. *J. Hydroinformatics* **9** (1), 3–14.
- Jowitt, P. & Germanopoulos, G. 1992 Optimal pump scheduling in water-supply networks. *J. Wat. Res. Plann. Mngmnt., ASCE* **118** (4), 406–422.
- Kolmogorov, A. N. 1957 [On the representation of continuous functions of several variables by superimposition of one variable and additions](#). *Dokl. Akad. Nauk USSR* **114**, 369–373.
- Little, K. W. & McCrodden, B. J. 1989 Minimization of raw water pumping costs using MILP. *J. Wat. Res. Plann. Mngmnt., ASCE* **115** (4), 511–522.
- Lorentz, G. G. 1976 The 13th Problem of Hilbert. *Proc. Symp. in Pure Maths.* **28**, 419–430.
- Martinez, F., Hernandez, V., Alonso, J. M., Rao, Z. & Alvisi, S. 2007 [Optimizing the operation of the Valencia water-distribution network](#). *J. Hydroinformatics* **9** (1), 65–78.
- Rao, Z. & Jamieson, D. G. 1997 The use of neural networks and genetic algorithms for the design of groundwater remediation schemes. *Hydrol. Earth Syst. Sci.* **1** (2), 345–355.
- Rao, Z. & Salomons, E. 2007 [Development of a real-time, near-optimal control system for water-distribution networks](#). *J. Hydroinformatics* **9** (1), 25–38.
- Rogers, L. L. & Dowla, F. U. 1994 [Optimization of groundwater remediation using artificial neural networks with parallel solute transport modeling](#). *Wat. Res. Res.* **30** (2), 457–481.
- Rossman, L. A. 2000 *EPANET User's Manual*. US EPA, Cincinnati, OH.
- Rumelhart, D. E. & McClelland, D. E. 1986 *Parallel Distributed Processing*, vol **1**. MIT Press, Cambridge, MA.
- Salomons, E., Goryashko, A., Shamir, U., Rao, Z. & Alvisi, S. 2007 [Optimizing the operation of the Haifa-A water-distribution network](#). *J. Hydroinformatics* **9** (1), 51–64.
- Sprecher, D. A. 1965 [On the structure of continuous functions of several variables](#). *Trans. Am. Math. Soc.* **115**, 340–355.
- Sterling, M. J. H. & Coulbeck, B. 1975 A dynamic programming solution to the optimisation of pumping costs. *Proc. ICE* **59** (2), 813–818.
- Takahashi, Y. 1993 Generalization and approximation capabilities of multi-layer networks. *Neural Comp.* **5**, 132–139.
- Tarquin, A. J. & Dawdy, J. 1989 Optimal pump operation in water distribution systems. *J. Hydrol. Engng., ASCE* **115** 158–168.
- Walski, T. M., Brill, E.D. Jr., Gessler, J., Goulter, I. C., Jeppson, R. M., Lansey, K. E., Lee, H. L., Leibman, J. C., Mays, L. W., Morgan, D. R. & Ormsbee, L. E. 1987 Battle of the network models: Epilogue. *J. Wat. Res. Plann. Mngmnt., ASCE* **113** (2), 191–203.
- Wang, C. G. & Jamieson, D. G. 2002 An objective approach to regional wastewater-treatment planning. *Wat. Res. Res.* **38** (3), 4.1–4.8.
- White, H. 1990 [Connectionist nonparametric regression: multi-layer feedforward networks can learn arbitrary mapping](#). *Neural Net* **3**, 535–549.