Water Pressure Optimisation for Leakage Management Using Q Learning*

line 1: 1st Ahmed Negm
line 2: *Engineering Department*
line 3: *Lancaster University Energy Group*
line 4: Lancaster, UK
line 5: a.negm@lancaster.ac.uk

line 1: Xiandong Ma
line 2: *Engineering Department*
line 3: *Lancaster University Energy Group*
line 4: Lancaster, UK
line 5: xiandong.ma@lancaster.ac.uk

line 1: George Aggidis
line 2: *Engineering Department*
line 3: *Lancaster University Energy Group*
line 4: Lancaster, UK
line 5: g.aggidis@lancaster.ac.uk

***Abstract*— The recent global urbanisation problem has set the industry and researchers sights to the importance of safe, effective water distribution due to the unprecedent demand placed on our aging water networks. Our current water practices often increase the degradation of assets through heightened pressures causing more failures and leakage. Whilst the higher network pressures assure customer demands are met; they cause detrimental failures to the system, long-term expenses, higher carbon emissions and energy consumption. This paper uses a baseline reinforcement learning algorithm to optimise valve set point for active pressure control. Using optimised Q-learning in an EPANET-Python environment, the agent learns to modify valve set points to decrease the average pressures whilst remaining within the OFWAT mandated pressure limits of 10m. This code is tested on the d-town test network. The agent shows continuous improvement finding an optimised set point of 26m and dropping the average system pressure by 2% by making simple changes to two pressure reducing valves. The agent learns the optimal actions to take for different states however further improvements can be made through the use of deep neural networks.**

***Keywords*— Reinforcement Learning; Urban Water; Pressure Optimisation**

I. Introduction (*Heading 1*)

Pressure management is quickly becoming a necessity for the operation of urban water network as the demand increases and leakage infecting most water distribution systems (WDS). In earlier reports, the loss of water was estimated to be 30% which is assumed to increase [1]. Leakage in potable water distribution systems is costly in both economic and environmental terms. Any process that can minimise this wastage is seen as a positive practice to employ by both water companies and regulatory bodies alike. Additionally, leakage leaves WDSs vulnerable to contaminants and a degrading water quality. Several factors heighten the leakage level including inherent factors such as pipe infrastructure, age-induced corrosion, pipe fittings or operational factors such as high pressures, cyclic loading, and transient surges [2].

Several pressure management algorithms draw on the use of genetic algorithms (GA) [1], [3] or genetic algorithm hybrids [4] in order to optimise the pressure reducing valve (PRV) settings and therefore control the system pressure. These articles have highlighted the effects of pressure regulation on leakage prevention and carbon emissions. In this case study, we use a Q-learning algorithm to train an agent to choose the best PRV settings for the 'D-town' test network [5].

II. Methodology

Reinforcement learning is a fast-growing emerging field of machine learning that imitates the natural human perception method by implementing a trial-and-error strategy that explores the optimisation space and learns the optimal policy. This AI strategy is founded in Markov Decision Processes (MDP) and is widely applicable in fields such as robotics, games, stocks trading and more [6].

In this study, we create a virtual environment consisting of the network model on EPANET [7] and a complementary python class that extracts, processes and changes network data. This environment defines the states of the model and outputs the current observation, state and reward to trigger the agent's upcoming action. The schematic can be shown below in Figure 1. The EPANET software and python terminal communicate using the epynet library by Vitens [8].

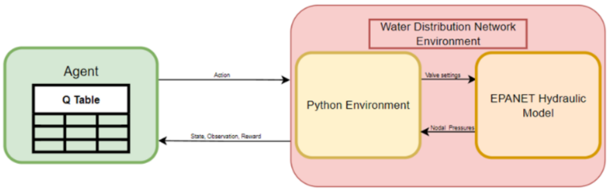

Fig. 1. Q-learning schematic for water distribution networks

The agent learns the optimal policy using a Q learning algorithm that can be summarised in the equation below.

$$Q^{new}(s_t, a_t) = Q(s_t, a_t) + \alpha \cdot (r_t + \gamma \cdot maxQ(s_{t+1}, a_t) - Q(s_{t+1}, a_t)) \quad (1)$$

The equation reads the new Q value for a given action and state in the Q table will equal to the old Q value plus the learning factor (α) multiplied by the current reward (r_t)added to the discount factor (γ) multiplied by max future Q factor and subtracting the current Q value. This algorithm is executed iteratively after every step the agent takes to populate the Q table for every state (rows) and action (column). This trial and error behaviour expose the agent to global minima errors which is a common problem in RL called the 'Exploration/Exploitation" dilemma. In order to bypass this a random action is taking through the epsilon parameter that ranges between 0-1. The closer the value is to one the more likely the agent will perform a random action. As the agent learns the model, this epsilon value decays to allow the agent to behave based on the learned policy and decrease the amount of random actions. Another important parameter is the discount factor, γ, which decrees how important future rewards are over current actions to the agent. The learning rate dictates how aggressive the optimisation search is.

In our experiment we use the average network pressure and pressure violations less than 10m as our observations. Rewards were calculated based on whether the step lowers the average pressure (negative) or increases the pressure (positive) and whether more node are in violations (negative) or less (positive). The optimisation parameters are dictated in the table below.

TABLE I. Q-LEARNING OPTIMISATION HYPERPARAMETERS

Netw-ork	Learning rate	Discount factor	Epsilo-n	Episode length	Epis-des
D-town	0.5	0.95	0.7	48 steps	15000

III. Discussion and Results

The RL algorithm has shown to improve the agent's understanding of the network over time as it begins with random actions and slowly optimises its actions to receive less negative rewards. This result has been optimised by modifying the hyperparameters which showed us that a higher epsilon was needed to overcome local minima. Figure 2 shows the rewards with a moving average of 500 episodes whilst Figure 3 shows a comparison between the best (max), average and worst (min) performances.

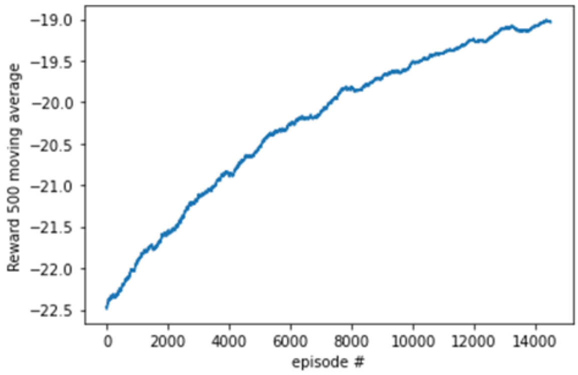

Fig. 2. Rewards over episode number using a moving average of 500 episodes

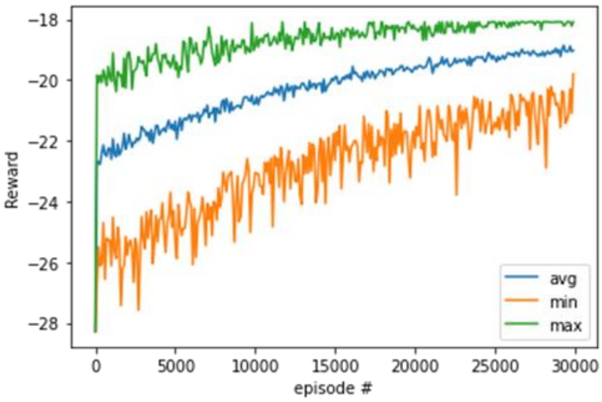

Fig. 3. A comparison between the minimum average and maximum rewards

Figure 3 shows that the maximum reward levels off at -18 and this is due to the negative rewards collected as the agent approaches the optimum but it now learns how to reach the optimal valve set point as fast as possible of the algorithms.

IV. Conclusions

Our results prove the usefulness of Reinforcement Learning as an optimisation tool for pressure management. However, this tool requires hyperparameter optimisation to be useful and can be computationally demanding. Other optimisation techniques would outperform Q-learning unless we introduce deep neural network to the technique. The use of 'Deep Reinforcement Learning' for pressure optimisation should be a promising field full of novelties.

Acknowledgment *(Heading 5)*

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] A. S. R. Araujo et al., "Pressure Control for Leakage Minimisation in Water Distribution Systems Management," Water Resour. Manag. 2006 201, vol. 20, no. 1, pp. 133–149, Apr. 2006, doi: 10.1007/S11269-006-4635-3.

[2] H. Rezaei, B. Ryan, and I. Stoianov, "Pipe Failure Analysis and Impact of Dynamic Hydraulic Conditions in Water Supply Networks," Procedia Eng., vol. 119, no. 1, pp. 253–262, Jan. 2015, doi: 10.1016/J.PROENG.2015.08.883.

[3] H. Monsef, M. Naghashzadegan, R. Farmani, and A. Jamali, "Pressure Management in Water Distribution Systems in Order to Reduce Energy Consumption and Background Leakage," Water supply Res. Technol., vol. 67, no. 4, pp. 397–403, 2018.

[4] S. Nazif et al., "Pressure Management Model for Urban Water Distribution Networks," Water Resour. Manag. 2009 243, vol. 24, no. 3, pp. 437–458, May 2009, doi: 10.1007/S11269-009-9454-X.

[5] A. Marchi et al., "Battle of the Water Networks II," J. Water Resour. Plan. Manag., vol. 140, no. 7, p. 04014009, May 2013, doi: 10.1061/(ASCE)WR.1943-5452.0000378.

[6] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," IEEE Signal Process. Mag., vol. 34, no. 6, pp. 26–38, Nov. 2017, doi: 10.1109/MSP.2017.2743240.

[7] L. A. Rossman, "EPANET 2 USERS MANUAL," 2000.

[8] Vitens, "Epynet: Object-oriented wrapper for EPANET 2.1," 2017. https://github.com/Vitens/epynet (accessed May 06, 2022).

This work is funded by Lancaster University through European Regional Development Fund, Centre of Global Eco-Innovation, and industrial partners DNS Ltd.