RL4Water: Reinforcement Learning environment for Water Management

Author
Krzysztof Muniak
k.muniak@student.tudelft.nl

Responsible professor
Pradeep Murukannaiah

Supervisor
Zuzanna Osika

TU Delft – Delft University of Technology

1. Introduction

• Water management becomes a more pressing issue in the face of climate change and population growth.

• **Multi-objective Reinforcement Learning (MORL)**- a branch of machine learning where an agent learns to make decisions by interacting with an environment to simultaneously maximize more than one objective.

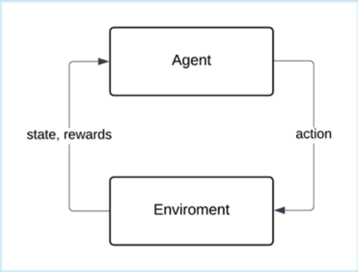

Figure 1: Multi-objective reinforcement learning interaction of the agent with the environment.

• **Water management systems** – model water flow network focusing on reservoir release control and meeting water consumption demands.

• **Water management frameworks** rely on the MORL principles, but do not utilize algorithms developed in that field.

• **Existing literature** - provides insights into the modelling of complex water management systems, but does not provide a unified solution for creating water management simulations.

• **Research question:** How can the water systems simulation be generalized for multi-objective reinforcement learning?

2. Methodology

1. **Study** of the Nile River simulation [1] to identify its core components and similarities with MORL.

2. **Implementation** of the Nile River simulation to be compatible with the Gymnasium API [2] standard and connecting it to MONES.

3. **Verification** of the framework's correctness by verifying original and adapted simulation behaviour.

4. **Generalization** of the framework to make it customizable for different water management problems.

5. **Validation** of the generalization principles developed in the previous step, by implementing Susquehanna River simulation [3].

6. **Comparison** of the RL4Water framework with another RL software for water management [4].

3. RL4Water framework

The RL4Water framework consists of four main classes:

• **Water Management System** - extends the Gymnasium environment allowing for interaction with the agent. It is responsible for controlling the whole water system and integrating the behaviour of the other classes.

• **Flow** - models water propagation through the system, connecting multiple source and destination facilities.

• **Facility** - represents a static component of the system that cannot be controlled by the agent. It is extended by the following classes:

- **Catchment** – models changing inflows to the system.
- **Power Plant**– produces power based on water inflow and physical properties.
- **Demand District** – consumes water and provides rewards based on the proportion of demand met.

• **Controlled Facility** – defines water system components that can be managed by the agent. It is extended by the following class:

- **Reservoir** – can accumulate water and release it based on provided actions.

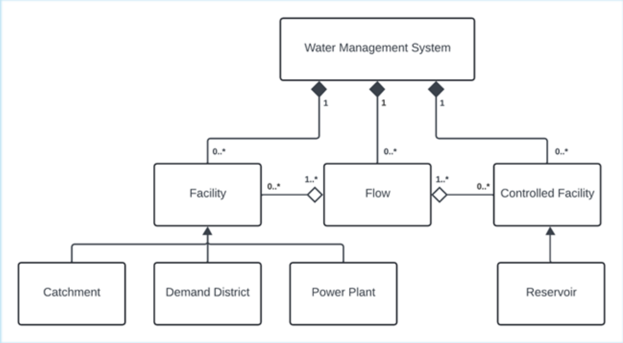

Figure 2: Multi-objective reinforcement learning interaction of the agent with the environment.

Generalization properties

• **Flexibility:** all facility classes can be tailored to simulate various water systems by adjusting their parameters.

• **Modularity:** the framework allows for the creation of complex water systems by seamless connection of smaller components.

• **Extensibility:** the codebase was designed to support extension by other facility classes with minimal effort.

• **Compatibility:** the framework is compatible with various RL and MORL algorithms.

Verification

The correctness of the RL4Water framework was verified by comparing the behaviour of the Nile River simulation to its original implementation. The verification process involved checking the amount of water stored and released from all reservoirs. The results showed a relative difference of less than 1e-10 between the RL4Water simulation and the original implementation, demonstrating the high accuracy of the framework.

4. Case studies

Nile River case study simulated a basin spanning three countries: Egypt, Sudan and Ethiopia. It integrated all of the RL4Water components: reservoirs, power plants, demand districts and catchments, into the structure shown in Figure 3.

Each of the actors in the region has different objectives that were taken into account.

Country	Objective	Direction
Egypt	Irrigation demand deficit	Minimisation
Egypt	Minimum HAD water level	Maximisation
Sudan	Irrigation demand deficit	Minimisation
Ethiopia	Hydropower production	Maximisation

Table 1: Objectives of the Nile River simulation.

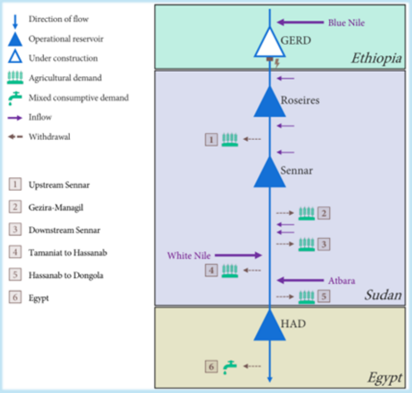

Figure 3: The topology of the Nile River case study detailing all components of the simulation.

The simulation showcased the compatibility of the RL4Water framework with the Gymnasium API and MONES algorithm. It demonstrated the adaptability of RL4Water in managing a large-scale, distributed water system.

Susquehanna River case study simulated a reservoir located in the lower part of the river basin, controlled by the states of Pennsylvania and Maryland. The dam determines the water inflow to four different facilities illustrated in Figure 4.

Since the reservoir is responsible for controlling a large share of river flow in the basin, it impacts multiple shareholders with different objectives.

Location	Objective	Direction
Conowingo reservoir	Minimum water level	Maximisation
Conowingo power plant	Hydropower revenue	Minimisation
Baltimore	Water supply	Maximisation
Chester	Water supply	Minimisation
Peach Bottom	Water supply	Maximisation
Downstream area	Minimum water flow	Maximisation

Table 2: Objectives of the Susquehanna River simulation.

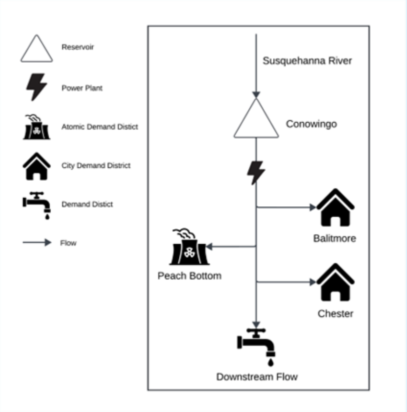

Figure 4: The topology of the Susquehanna River case study detailing all components of the simulation.

The case study proved the generalization properties of the framework. Moreover, it demonstrated its application to centralized systems with a complex reservoir.

5. Limitations

• **Lack of Unit Tests:** only end-to-end test is created, making it harder to locate bugs when encountered.

• **No Assertion of Environment Correctness:** users can incorrectly set up environments, resulting in runtime errors or incorrect results.

6. Conclusions

RL4Water bridges the gap between water management and RL fields, demonstrating practical applications of MORL to real-world problems. It provides a unified platform for developing water management simulations and integrates with multiple RL algorithms through the Gymnasium API. RL4Water's modular design and customizable components confirm its ability to generalize across various water systems, validated through the case studies of the Nile River and the Susquehanna River.

Future work

• **Automatic Network Topology Detection:** implement features for automatic detection of water flow in the system to simplify environment setup.

• **Validation on Additional Simulations:** extend validation to include more water management simulations to increase the framework's applicability.

• **Translation of EMODPS Algorithm:** adapt the Evolutionary Multi-Objective Direct Policy Search (EMODPS) algorithm to fit the Gymnasium API structure to improve compatibility.

[1] Sari Yasin. Exploring trade-offs in reservoir operations through many objective optimisation: Case of nile river basin. 2022.
[2] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
[3] Witvliet Mark. Multi-sector water allocation: The impact of nonlinear approximation network hyperparameters for multi-objective reservoir control. 2022.
[4] Faber Jorian. Bottom-up formulation of water management systems as a reinforcement learning problem. 2024.