Bottom-up Formulation of Water Management Systems as a Reinforcement Learning Problem

Author: J.E.C. Faber[1]
Supervisors: P.K. Murukannaiah[2] & Z. Osika[2]
[1] j.e.c.faber@student.tudelft.nl
[2] Intelligent Systems, Interactive Intelligence, Delft University of Technology

1 Introduction

- **Water Management Systems (WMS)** are complex systems, where often **multiple conflicting objectives** are at stake.
  - Hydropower production
  - Irrigation demands
  - Water supply
- **Reinforcement Learning (RL)** particularly well-suited for tackling such problems.
  - Handling multiple objectives as **linear combination sub-optimal**
  - Instead return **vector of rewards**
- This research analyses **three case studies of WMS in the context of RL.**
  - Nile River in Ethiopia, Sudan and Egypt [2]
  - Lower Volta River in Ghana [1]
  - Susquehanna River in USA [3]

RL Problem Definition

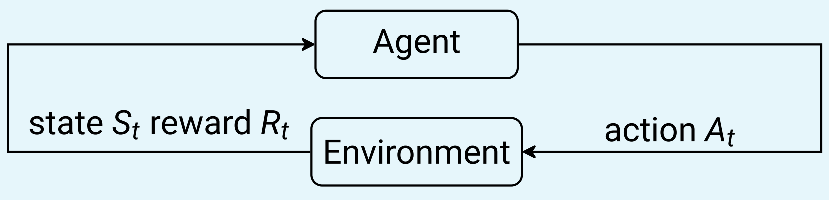

Figure: Agent-environment interaction

An **agent** interacts with an **environment**, a set of **states**. Based on a (partial) **observation** of the environment, the agent makes a decision in the form of an **action**. Through this action, it transitions to a different state and receives a **reward**. The agent 'learns' by trying to maximise the cumulative reward.

2 Research Question

"What are **core properties of a WMS** based on three case studies, and how can they be **modelled as a RL problem?**"

3 Methodology

1. **Analyse** three case studies, consider RL concepts
2. **Identify** similarities and differences
3. **Define** core properties of general WMS
4. **Formulate** RL problem
5. **Implement** as custom Gymnasium Environment

4 Results

Core properties basic WMS –> water moving from A to B

- **Graph structure** where **nodes** represent water points and **directed edges** represent flows.

Basic WMS as RL problem formulation

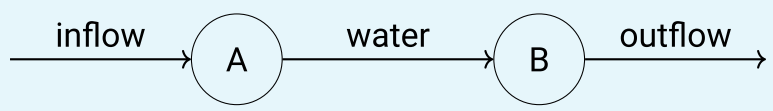

Figure: Graph representation of most basic WMS

Nodes have a few key properties: `volume`, `capacity`, `demand`, `inflow` and `outflow`. Edges have a `source`, `destination` and `flow`. Nodes can have an incoming- and outgoing edge.
Environment: all nodes and edges.
State: state of nodes and edges.
Observation: `volume` at each node.
Action: how much water to release from a node.

Adding features through generalisation

With **basic WMS as starting point**, features are added to simulate more mechanics from the case studies. Features are **generalised** to **handle differences** and provide **greater flexibility**.

Mass-balance equation: Instead of individually implementing volume changes such as **evaporation** and **leakage**, users can add **"water update functions"**, which take the state's `info` and return a float. Equation 1 shows this - with V volume, IF inflow, OF outflow and $\sum WUF$ the sum of water update functions at time t.

$$V_{t+1} = V_t + \sum IF_{t+1} - \sum OF_{t+1} + \sum WUF_{t+1} \tag{1}$$

Custom reward function: Problem of **different reward functions** across case studies. Implementation allows for **custom reward functions**. Just like water update functions, they take the current state's `info` and return a float. Below is a very simple example.

```
1: Function volumetric_reliability(info)
2:    return info["inflow"]/info["demand"]
```

Larger in-degree and out-degree: Larger in-degree implemented by simply **summing incoming flows**. Out-degree more complex: **proportional allocation** to scale total **release within bounds**.

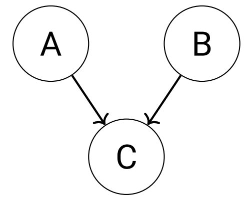

Figure: Node with in-degree of 2

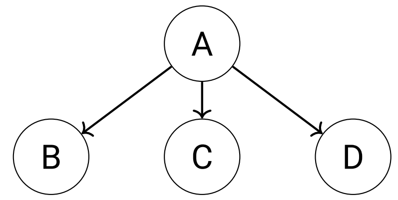

Figure: Node with out-degree of 3

Comparison with Nile River simulation

Same actions: proposed implementation produces **same results as Nile River** simulation. It is therefore **just as accurate**.

5 Conclusion

- Provides **generalised formulation and implementation** of WMSs as a RL problem.
- Applicable to **wide range of WMSs, flexible** and **versatile**.
- Should be **tested on more WMSs**.
- Contributes to research on WMS in context of RL.
  - **Removes the need** to write WMS simulation from scratch.

References

[1] A. Owusu, J. Z. Salazar, M. Mul, P. van der Zaag, and J. Slinger. "Quantifying the trade-offs in re-operating dams for the environment in the Lower Volta River". In: *Hydrology and Earth System Sciences* (May 2023).
[2] Y. Sari. "Exploring Trade-offs in Reservoir Operations through Many Objective Optimisation: Case of Nile River Basin". Master's Thesis. Delft University of Technology, Aug. 2022.
[3] M. Witvliet. "Multi-sector Water Allocation: The impact of nonlinear approximation network hyperparameters for multi-objective reservoir control". Master's Thesis. Delft University of Technology, Apr. 2022.

TU Delft — Delft University of Technology