02 - MULTI-OBJECTIVE OPTIMISATION AND RESERVOIR MANAGEMENT

Eva Kordomenidi, JBA Consulting

Abstract

Multi-objective reservoirs are often used to serve multiple demands for domestic, industrial, irrigation, environment, hydropower production, and flood control. Determining how to best meet these multiple needs is a complex problem because of the nonlinear storage-inflow relationship, conflicting objectives, dynamic properties, and nonlinear constraints. The answer is set of solutions that define the best trade-off between competing objectives. This paper presents the findings from the multi-criteria optimization modelling exercise for a dam in Southeast Asia. The dam has outlets for irrigation and hydropower generation as well as a spillway with seven gates that control flood releases. In this study, a methodology for undertaking multi-objective optimization for reservoir operations has been developed and applied to Dam A. The study is intended to be a proof-of-concept pilot demonstrating the feasibility of applying multi-objective optimization to fully harness the capacity of dam/reservoir assets. In this pilot, the reservoir rule curve is adjusted so that three performance objectives are simultaneously minimized irrigation deficit, hydropower deficit and area of flood inundation.

The performance of the reservoir rule curve is assessed in terms of expected annual values for the three objectives, by simulating the long-term (23-year) continuous behaviour of the reservoir inflow, storage, and outflow. Parametrizing the rule curve provides a means for allowing a degree of freedom. An optimal reservoir rule curve can be found by searching for the rule curve parameters which optimize the three performance objectives. The results from the study show that it is possible, with a parametric rule curve, to undertake multi-objective optimization for dam operations in a practical, efficient, and comprehensive manner.

This approach illustrates wider application of multi-objective optimization to reservoir management as it is practicable and computationally achievable without the need for specialist hydraulic modelling software, the need to code an optimization algorithm or the need for high-performance computers. The study has not been extended to the next step of selecting a single preferred reservoir rule curve (dam operating strategy) from the Pareto set of optimum solutions. This would require consideration of trade-offs between the three objectives, which would involve all stakeholders.

1. INTRODUCTION TO MULTI-OBJECTIVE OPTIMISATION

Most of the water resources optimization problems involve conflicting objectives. Most of the reservoir systems serve multiple purposes and they are multi-objective in nature. Multi-objective optimization problems require simultaneous optimization of several incommensurable and often conflicting objectives. Because of presence of the multiple conflicting objectives, it is not possible to find a single optimal solution, which satisfies all goals. Instead, the interaction of multiple objectives yields a set of efficient or non-dominant solutions. These optimal solutions belong to the space of objective functions, and the set of the optimal solutions is called the optimal Pareto set. The optimal Pareto set identifies a series of optimal solutions, which consider the existing relation between the different criteria of optimality, and that are characterized by the fact that starting from a solution within the set, one objective can only be improved at the expense of at least one other objective.

The goal of solving a multi-objective problem is to find the optimal Pareto set for the decision-maker to choose the preferred solution. A solution selected by the decision-maker always represents a compromise between the different objectives. Therefore, one cannot speak of a unique solution, but of a series of compromise solutions, which take into account the existing relation between the different objectives, giving the decision maker more flexibility in the selection of a suitable alternative.

Developing a trade-off between long and short-term objectives is one of the challenging issues in reservoir operation management. Flood control and water supply are two common examples of short-term and long-term reservoir operation objectives.

2. PROJECT OBJECTIVES

The analysis involves the following:

- Review of historical reservoir release operating rules and actual practice.
- Dam operation objectives: The following objectives are relevant; (i) irrigation demand, (ii) hydropower production and (iii) downstream area flooded.
- Quantifying objectives: Each objective is quantified using a metric of its expected annual value over a long-term (23-year) hydrological simulation of the reservoir's behavior.
- Simulating the historical reservoir release operating rule, and quantifying historical performance on the three metrics, as a baseline.
- Undertaking multi-objective optimization to search the solution space of potential reservoir rule curves to find the Pareto (non-dominated solution) surface.
- Considering possible areas for further improving reservoir operating performance, e.g., by gauging upstream water-levels and/or reservoir catchment rainfall in conjunction with runoff forecasting.

3. RESERVOIR AND DAM DESCRIPTION

Dam A, in Southeast Asia, is used for the purpose of irrigation water supply and hydropower generation. The exact location remains confidential following our client request. The dam, built in the 1980s, has a spillway with gates that control the release of water.

3.1 Catchment

The reservoir drains a catchment of approximately 3826 km^2. The catchment is primarily rural, and elevation is relatively steep. Soils in the catchment are a mixture of red brown and mountain brown forest soils. Highland regions of Southeast Asia are covered with highly leached, iron-rich, dark red and reddish brown soils. These soils tend to erode quickly once the forest has been cleared. The catchment is mostly covered in woodland: a mixture of deciduous and evergreen forest, with some scrub in the lower-lying portions closer to the dam.

3.2 Uses

- The primary purpose of the dam is to store water which can then be supplied for irrigation during the dry season, enabling the production of two crops per year within the irrigated areas

(Monsoon paddy and summer paddy). Water for irrigation is abstracted from a Weir, located 3.8km downstream of the dam. The system of irrigation canals and weirs dates back many centuries before construction of the dam. There are canal offtakes on the south and north sides of the river, upstream of the weir, and these feed 2 canals. The total irrigated area was approximately 37,000 hectares (93,000 acres) in 2017-18, which represents 73% of the total irrigable area. The irrigated area makes up about 1.7% of the total irrigated land in the country.

- Hydropower: Hydropower generation is provided by two units at Dam A1, with a combined capacity of 25 MW. The average annual energy production is 134 GWh, which implies that the hydropower units run at 61% of their rated capacity on average through the year.
- Domestic supply: The dam also serves the purpose of providing domestic water. This is limited to 3.4 m3/s (120 cusecs). In this study, we will assume that this amount is negligible, and the purpose of domestic supply will be excluded from the multi -objective optimization.

4. DATA

4.1 Hydropower flow data

Hydropower flow records have been provided at daily intervals for 1 January 1996 to 31 December 2019. Most flows are above 10 m^3/s and below 75 m^3/s.

4.2 Irrigation data

Irrigation abstraction records have been provided at daily intervals for the period 1 January 1996 to 5 August 2019.The irrigation abstraction flow is mostly below 60 m^3/s. Demand for irrigation is lower between mid-November and early February. For the rest of the year irrigation abstraction seems constant.

4.3 Reservoir level data

Analysis of the reservoir levels undertaken suggests a strongly seasonal annual cycle of reservoir level, with high levels at the end of the wet season (from October onwards), declining levels through the dry season and then increasing levels once the monsoon arrives, typically from June onwards. There appears to have been a tendency towards a more gradual drawdown in the early months of the year (January to March). It is not known whether this is due to variations in weather or in the amount of water supplied from the reservoir for irrigation, hydropower, and water supply. However, it is a fact that there is less water demand for irrigation and hydropower during these months than the rest of the year.

The minimum level can be reached any time between about early May and late July, largely dependent on the timing of the onset of the monsoon. Note that the dead storage level is quoted as 111.25m. The water level apparently dropped below this in 1999, perhaps aided by evaporation.

The rate of drawdown through the winter and spring months, and the amount of storage created in the reservoir, does not appear to affect the eventual peak water level reached during the wet season.

This type of finding is quite common on reservoir systems where the volume of annual inflow is large compared with the volume of storage available. At Dam A, the average annual inflow volume is about 5100 million m^3. In comparison, the reservoir capacity (between dead storage level and full supply level) is about 340 million m^3. So even if the reservoir is drawn down to the dead storage level, in an average year it could refill several times over.

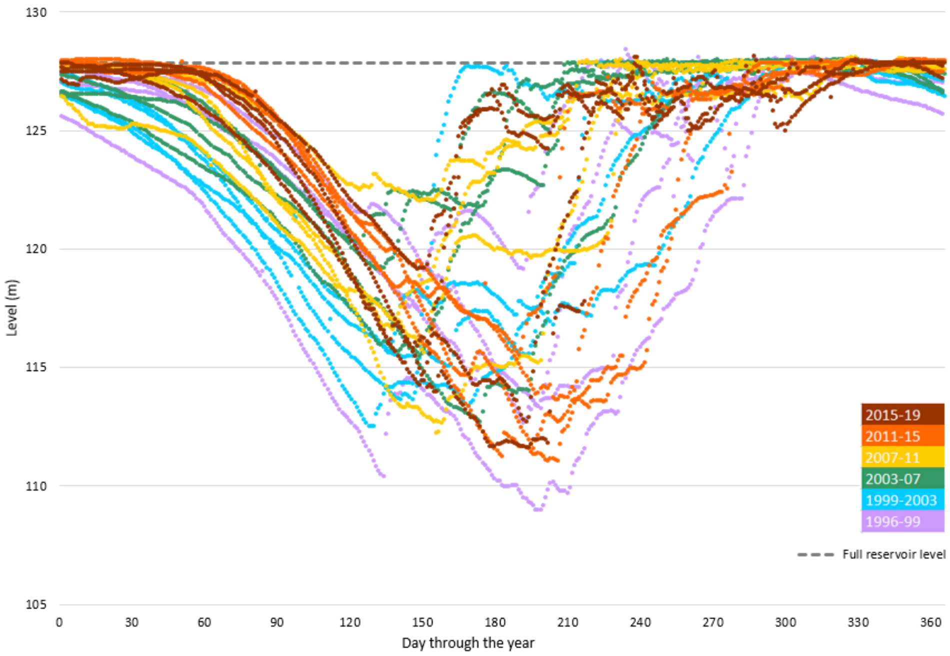

***Figure 1**: Annual (Jan-Dec) variation in water level, colour-coded by groups of years*

5. METHODOLOGY

5.1 Approach to optimization

Multi-criteria, or multi-objective, optimization problems require simultaneous optimization of several often-competing objectives. Because of this, it is not possible to find a single optimal solution, which satisfies all goals. Instead, the solutions exist in the form of alternative trade-offs, also known as the Pareto optimal solutions. In the cases of multi-objective optimization, typically there is no single optimal solution which can simultaneously satisfy all the goals, but rather a set of non-dominated or Pareto optimal solutions exists. Generation of Pareto optimal solutions is not a trivial task. The optimal Pareto set identifies a series of optimal solutions, which consider the existing relation between the different criteria of optimality, and that are characterized by the fact that starting from a solution within the set, one objective can only be improved at the expense of at least one other objective.

The goal of solving a multi-objective problem is to find the set of Pareto optimal solutions from which the decision-maker can choose the preferred solution. A solution selected by the decision-maker always represents a compromise (trade-off) between the different objectives.

Most common methods used in reservoir optimization require a large number of decision variables, which are typically the sequences of releases from all reservoirs (if more than one reservoir) and for all time steps of the control period[2] [3]. This method is often referred to as high-dimensional or perfect foresight method. In this method, the decision variables are the complete series of time series, inflows, or releases (or transformations of them) from the reservoirs, depending on the problem formulated. It is noted that the values of decision variables in this method depend completely on the inflow series. Therefore, the assumption behind this method is that the inflow series are perfectly known for the entire control horizon (hence the name ''perfect foresight' method). One criticism of this method is that the number of decision variables becomes too large in the perfect foresight method and therefore it gets difficult to locate their optimal values even using efficient search algorithms like evolutionary algorithms.

The approach outlined above is for determining the optimal sequence of outflows for a predefined sequence of inflows over some control period and is valid only for that sequence of inflows. The approach taken in this study is completely different and starts from the premise that operators will not be comfortable deviating too far from the historical rule curve and that a perturbation of an historical rule curve is likely to be more easily acceptable in practice.

The framework used in this study is based on a parameterization-simulation optimization approach. The main idea of the parameterization-simulation optimization approach consists of parameterizing the operating rules for hydropower water releases for the reservoir. This methodology does not use the step-by-step releases of the reservoir as decision variables. It therefore avoids an extremely large number of variables. Therefore, the total number of decision variables in the system reduces and becomes independent of the number of simulated time steps. Parameterization of the rule curve is linked to simulation of the reservoir system, which enables the calculation of a performance measure of the system for given parameter values, and optimization, which enables determination of the optimal parameter values. This approach permits wider application of multi-objective optimization to reservoir management as it is more practicable and less computationally heavy.

The general idea of the parameterization-simulation optimization is not new. It has been applied before to reservoir systems and other water related systems, such as wastewater systems. [4] [5] [6] [7] [8]

If the step-by-step releases are used for the optimization, the number of decision variables is very high which increases the computational effort of the optimization and renders this a difficult task to model.

---

[2] Koutsoyiannis, D. and Economou, A., Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems, Water Resources Research, 10.1029/2003WR002148, 39, 6, (2003).

[3] Xiaomei Sun, Jungang Luo and Jiancang Xie Multi-Objective Optimization for Reservoir Operation Considering Water Diversion and Power Generation Objectives, Water 2018, 10(11), 1540 (2018).

[4] Guariso, G., S. Rinaldi, R. Soncini Sessa: The management of Lake Como WP-82-130, IIASA, Laxenburg, Austria. (1982).

[5] Lund, J. R., and I. Ferreira, Operating rule optimization for Missouri River reservoir system, J. Water Resour. Plann. Manage., 122(4), 287 – 296 (1996).

[6] Lund, J. R., and J. Guzman, Derived operating rules for reservoirs in series or in parallel, J. Water Resour. Plann. Manage., 125(3), 143 – 153 (1999).

[7] Nalbantis, I., and D. Koutsoyiannis, A parametric rule for planning and management of multiple reservoir systems, Water Resour. Res., 33(9), 2165 – 2177 (1997).

[8] Oliveira, R., and D. P. Loucks, Operating rules for multireservoir systems, Water Resour. Res., 33(4), 839 – 852, 1997.

Parametric optimization can be easily combined with simulation procedures. Also, by using the parametric methodology the optimum solution depends only on the statistical properties of the data series available. A disadvantage of the method is that the form of the operation rules is predefined; and actual data will deviate from this relationship. Also, in multiple reservoir systems a parameterized low dimensional approach can yield a suboptimal solution, as the solution depends on how good the adopted parameterization is and therefore high dimensional optimization is often preferred.

Depending on the problem and its formulation, the two approaches can produce outputs of similar quality.[9]

6. MULTI-CRITERIA OPTIMIZATION MODELLING FRAMEWORK

The methodology follows the basic steps below, also shown in Figure 2:

a. Historic data are reviewed and relationships between reservoir water level and reservoir flow releases are explored.

b. A rule curve is defined, which is a relationship between reservoir water-level and reservoir flow release (in this case hydropower release which is then used as irrigation water). The relationship is parametrized, with a small number of parameters that can be relaxed within a range and used as decision variables.

c. A reservoir routing model is developed driven by historic inflow data and the reservoir release outflow controlled by the parametric expression. Initial values of the parameters that will be used as decision variables are assigned.

d. Objective functions are formulated for: irrigation deficit, hydropower deficit, and flooded area. The three objectives can be calculated from the results of the reservoir routing model.

e. The decision variables to be optimized are driven by an external optimization search algorithm. Values are assigned to the decision variables, the reservoir routing model is simulated, the objective functions are evaluated. New decision variable values are chosen (in this case using the Non-dominated Sorting Genetic Algorithm III – NSGA-III) and the algorithm cycles in this way until convergence to an optimal solution is reached.

[9] Koutsoyiannis, D. and Economou, A., Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems, Water Resources Research, 10.1029/2003WR002148, 39, 6, (2003).

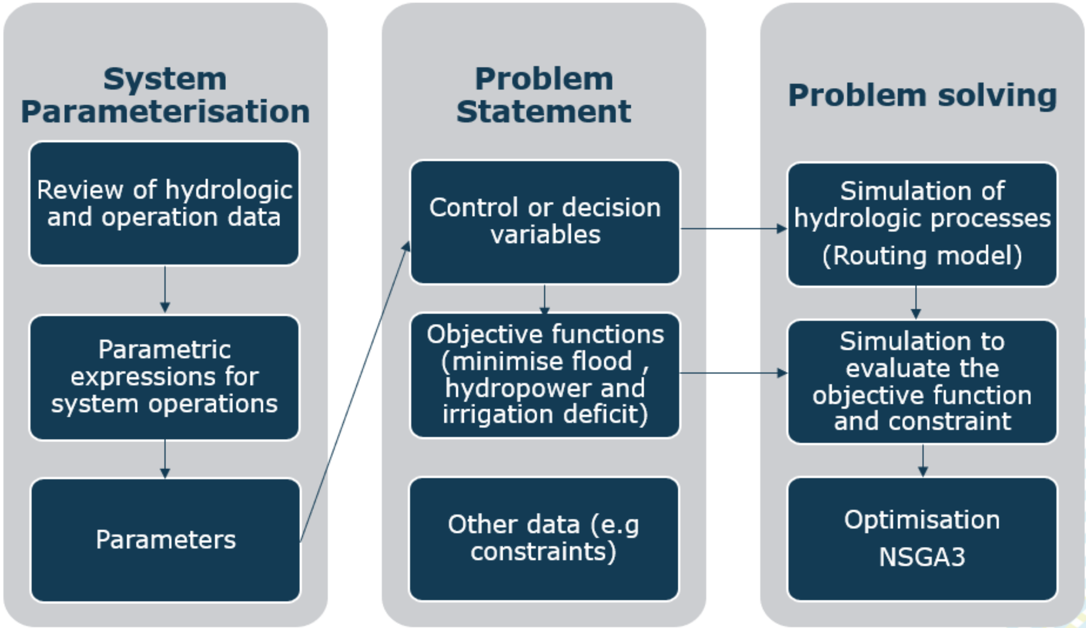

***Figure 2**: Multi-Criteria Optimization Framework*

7. PARAMETRIC EXPRESSION OF SYSTEM OPERATIONS

The total flow release from the reservoir potentially flows along three flow paths:

- through the dam hydro-electric station Q_{HE}
- through the dam irrigation gates Q_{IG}
- over the dam spillway gates Q_{SG}.

The dam irrigation outlets are almost always shut if the hydroelectric turbines are working, and the gates are opened only in flood conditions, according to a fixed procedure that has not been included as a variable in the optimization because the relative contribution under flood conditions is too small.

Therefore, the control rule for optimization is governed by the hydropower flow, Q_{HE}.

The irrigation abstraction at the downstream irrigation weir Q_{IW} can be calculated as the total flow release from the reservoir ($Q_{HE} + Q_{SG}$) minus the environmental flow (Q_{EF} which must be allowed to pass the irrigation flow offtake weir), up to a maximum irrigation demand (Q_{ID}) for the month concerned i.e.

$$Q_{IW} = min(Q_{HE} + Q_{SG} - Q_{EF}, Q_{ID})$$

The historic data suggests that the abstraction at the irrigation weir is higher than the reported irrigation demand and normally more in line with the irrigation canal capacity, but this has not been accounted in this exercise as the weir and canals system was not represented in the reservoir routing model.

The dam also serves the purpose of providing some domestic water supply to Mandalay. This is limited to 3.4 m^3/s (120 cusecs). Some data has been supplied which partially documents domestic supply over the last decade. There are several days where zero supply has been recorded. In this study, we have assumed that this amount is negligible in comparison with other releases and domestic supply has been excluded from the calculations.

The general formulation of the relationship between reservoir water-level versus reservoir release flow is shown in Figure 3. This is loosely based on an approach described by Hurford and Harou (2014)[10]. There are two components, (i) hydropower release, Q_{HE} and (ii) spillway release, Q_{SG}.

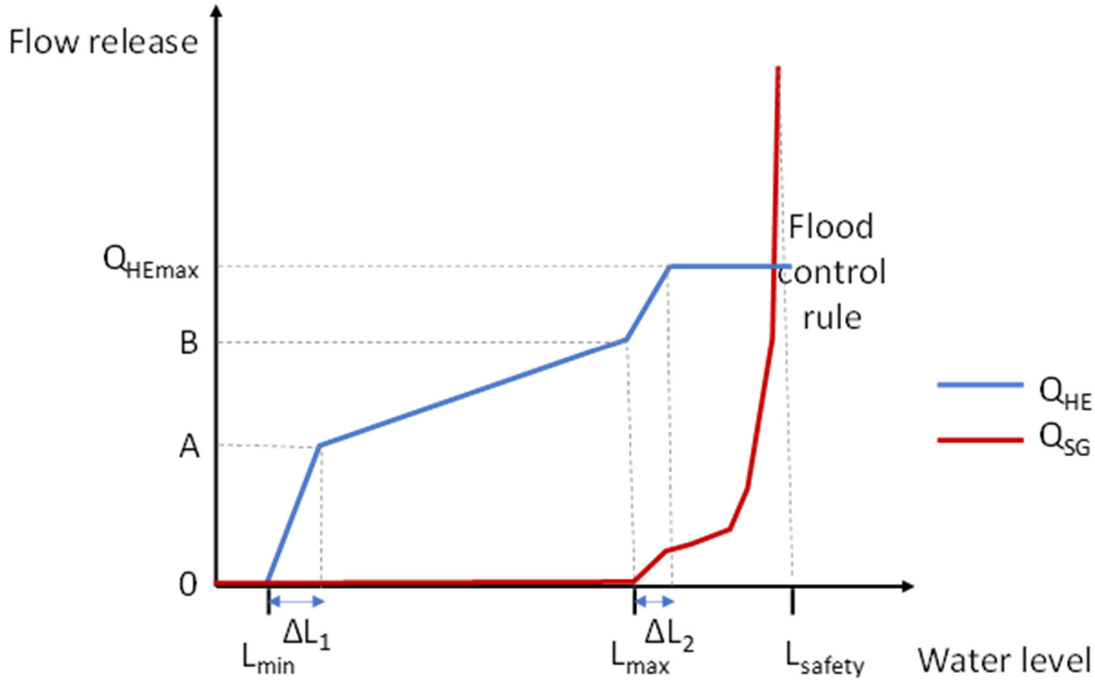

***Figure 3**: Generic relationship of reservoir release as a function of reservoir water-level.*

For the hydropower release, L_{min} and L_{max} are the minimum and maximum normal operating levels for the turbines. L_{min} is the dead water storage level based on the historic data. When L drops to L_{min} or lower then Q_{HE} is set to zero. L_{safety} is the maximum level allowed for dam safety purposes, which for this reservoir is set at 129.54mRL.

The data review found that there are three distinct seasonal patterns of reservoir releases throughout the year, probably in response to irrigation requirements. Therefore, to reproduce this seasonal response, it has been decided to use three seasonal rule curves, for the three seasons identified. The three seasons are (i) December to May, (ii) June to September and (iii) October to November. For each season, the rule curve has two variable parameters, A and B, therefore six parameters in total. A linear variation of flow release between A and B has been assumed. Although it would be logical for B to exceed A this has not been imposed as a constraint.

The rule curves also have two fixed parameters, ΔL_1 and ΔL_2. These have been introduced to avoid oscillations, keeping the flow transitions smooth. After testing ΔL_1 was set to 1.9m, and a sensitivity test was conducted using a smaller value.

[10] Hurford, A., and Harou, J. (2014). Balancing ecosystem services with energy and food security-assessing trade-offs for reservoir operation and irrigation investment in Kenya's Tana basin. Hydrol. Earth Syst. Sci. 11, 1343–1388. doi: 10.5194/hessd-11-1343-201 (2014).

Q_{HEmax} is the maximum possible flow through the turbines. At high water levels, Q_{HEmax} is approximately 74 m^3/s. During flood conditions, it would be sensible to set Q_{HE} as high as possible. This is done gradually so that by the time the level reaches $L_{max} + \Delta L_2$, $Q_{HE} = Q_{HEmax.}$ However, the observed data does not necessarily support this as it has been observed that for reservoir levels between 126-128mRL the Q_{HE} ranges widely between 10-80 m^3/s, which is not in line with what would be expected in terms of hydropower generation. No more anecdotal information was received to explain this, but it can be attributed to the hydropower plant not requiring to be operated due to demand and /or short term flood operations being of greater importance in these higher reservoir levels therefore not prioritizing hydropower generation during these events. Another possible explanation is that, perhaps having some knowledge that rainfall is falling in the catchment on the same day, water is being retained in the dam for future irrigation requirements. After testing ΔL_2 was set to 2.5m.

The initial values for the simulation model are derived by reviewing the monthly hydropower release data. Parameters A and B for each season are the decision variables in the optimization-simulation model. The initial values for parameters A and B for each of the three seasons are shown in Table 1.

The parameters A and B require a range (maximum and minimum value) for the optimization search algorithm. Based on the 5th and 95th percentiles of recorded hydropower release values, 25 m^3/s and 75 m^3/s were used as the lower and upper limit respectively for both parameters.

***Table 1**: Hydropower release: Initial parameter selection (parameters used as decision variables (i.e. not fixed) are shown in red)*

Reservoir level (m)	Hydro Release (m^3/s) (Dec-May)	Hydro Release (m^3/s) (June-Sept)	Hydro Release (m^3/s) (Oct Nov)	Notes
111.3 (L_{min})	0.0	0	0	Dead water level
115 ($L_{min} + \Delta L_1$)	**41.7**	**44.3**	**42.3**	Parameter A
126 (L_{max})	**55.1**	**70.6**	**63.5**	Parameter B: spillway gates start to open
128.5 ($L_{max} + \Delta L_2$)	74.0[11]	74.0	74.0	Spillway gates fully open
129.5(L_{safety})	74.0	74.0	74.0	Dam safety max

***Figure 4**: outlines the process of parameterizing the hydropower release rule curve.*

[11] Max Q_{HE} is approximately 74 m^3/s

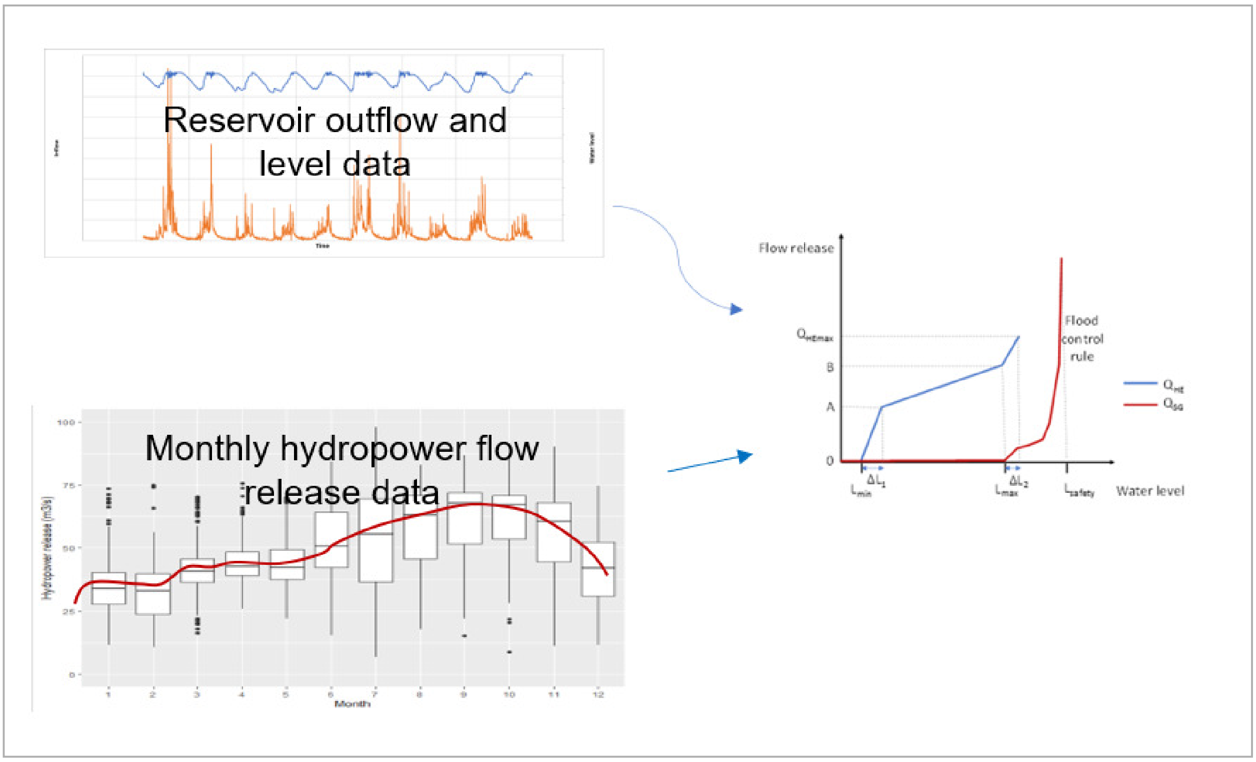

Figure 4: *Process of reservoir release rule parametric expression*

8. MODEL SETUP

8.1 Software platform

The modelling platform has been set up using Microsoft Excel and the add-on XLOptimizer. This is a combined simulation-optimization model. XLOptimizer is a generic optimization tool that uses Microsoft Excel as a platform for the definition of the problem at hand. Practically any problem that can be formulated in a spreadsheet can be tackled by this program. XLOptimizer uses the NSGA-III algorithm.

8.1.1 Optimization algorithm

The NSGA-III algorithm has been used to find a trade-off between the three objectives of irrigation, hydropower and flooding. A number of solutions form the initial population, which is extended to include a number of offspring solutions (created by crossover) and a number of mutant solutions (created by mutation of the population). These form the extended population, which is sorted, based on dominance, and categorized by rank (corresponding to each front). The extended population is then truncated to its original size and the process is repeated. The parameters of the NSGA-III algorithm have been selected based on sensitivity analysis to recommended ranges: the population size of 50, crossover probability and mutation probability of 0.5 have been used in this model.

8.1.2 Calculation procedures

Daily data has been used in the optimization model, covering the 23-year period from January 1996 to December 2018. This is the most complete set of data received for the project.

The spillway gates come into operation when there are floods, and the reservoir is already largely full. In practice the gates are operated to maintain the reservoir level at a level, around 126 m, lower than stated in the formal dam operating procedures.

8.2 Objective Functions

The objective function in an optimization problem is the real-valued function the value of which is to be either minimized or maximized over the set of feasible alternatives. There are three objective functions in the current simulation-optimization model.

8.2.1 Irrigation deficit

Monthly irrigation demand totals have been supplied. Figure 5 shows the process followed to derive the irrigation deficit estimates. The daily irrigation supply (which is derived from adding spillway and hydropower releases and subtracting 2 m^3/s assumed to be the minimum environmental flow requirement) has been used to estimate monthly supply. The deficits below the demand were calculated. Based on the initial parameterization, it has been noted that the only month when an irrigation deficit has been observed during the period of record is April. The observed data compared against the irrigation demand confirms this. Given that April is towards the end of the dry inter-monsoonal season, observing deficits to irrigation in this month is intuitive.

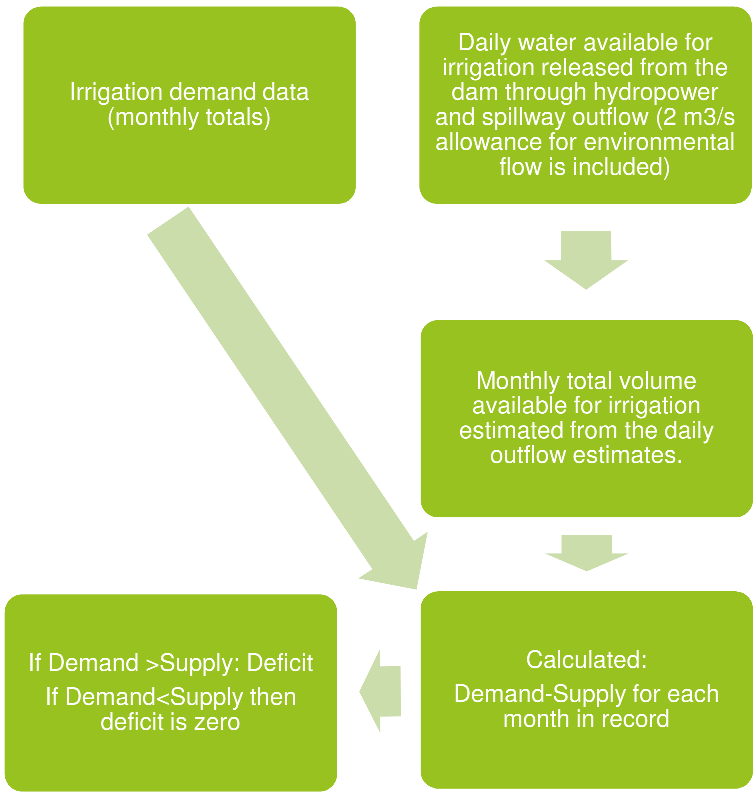

***Figure 5**: Irrigation deficit development schematic*

The sum of the difference of irrigation supply vs demand volume for each month for the 23 years of data available is the total deficit volume.

8.2.2 Flood damage

In order to quantify this objective, information on the impact of flooding downstream of the dam is required. Unlike the other two objectives, the flood damage objective depends not only on the nature of the annual reservoir rule curve but more importantly on the rules that control the operation of the spillway gates on a much shorter time scale, during the progression of a flood. We have not attempted to parameterize the flood rule curve as, in order for the simulation to represent flood damages it would be necessary to run it at a sub-daily time step. This could be computationally expensive, given the need to run numerous scenarios. Rather, existing information from another study undertaken for the downstream river was used. Flood damage is estimated by using the relationship between total daily outflow and the area downstream of the dam inundated under this outflow. The estimated areas flooded for different dam outflows are shown on Table 2.

There are some limitations to this approach, including the inability to account for the variable inflow from ungauged tributaries that join the river downstream of the dam, and the assumption that flood damages are directly proportional to the extent of the area flooded. It would be possible to trial different versions of this objective function in future work.

Table 2*: Downstream area flooded (information extracted from 'Strengthening Integrated Flood Risk Management (TA-9634 REG)', JBA and Landell Mills. July 2020)*

								1000 year	10000 year	PMF
Flow(m^3/s)	0	390	400	600	800	1000	1200	5150	6680	13200
Type of land use	Flood area (km^2)									
Bare ground	0	0	1.8	2.12	2.27	2.42	2.87	3.95	4.12	4.27
Field	0	0	15	29.83	41.6	55.29	70.49	150.78	162.39	195.8
Forest	0	0	0	1.04	1.5	2.22	2.91	7.1	8.17	10.92
Shrubland	0	0	1.38	2.02	2.39	2.76	3.01	4.27	4.55	5.15
Urban	0	0	0.68	1.86	2.6	3.74	5.75	13.39	14.64	17.71
Total area (Including river and floodplain area)	0	0	28.54	47.23	61.02	77.47	96.16	191.57	206.12	246.46

Based on the daily outflow, the daily area flooded is estimated and the total sum of area flooded each day is used as a proxy for flood damage.

8.2.3 Hydropower deficit

Unlike for irrigation, no data are available on hydropower demand. Therefore, hydropower demand cannot be a component on the deficit estimation .Information on max turbine output against head has been made available and this has been used for the formulation of the hydropower deficit. The hydropower deficit is calculated as the difference between the maximum turbine capacity for each day, which depends of the static head, and the energy produced, which is a function of the daily discharge through the hydropower station.

The difference between the maximum potential generation and the daily power output for each day of the 23 years is the hydropower deficit estimate.

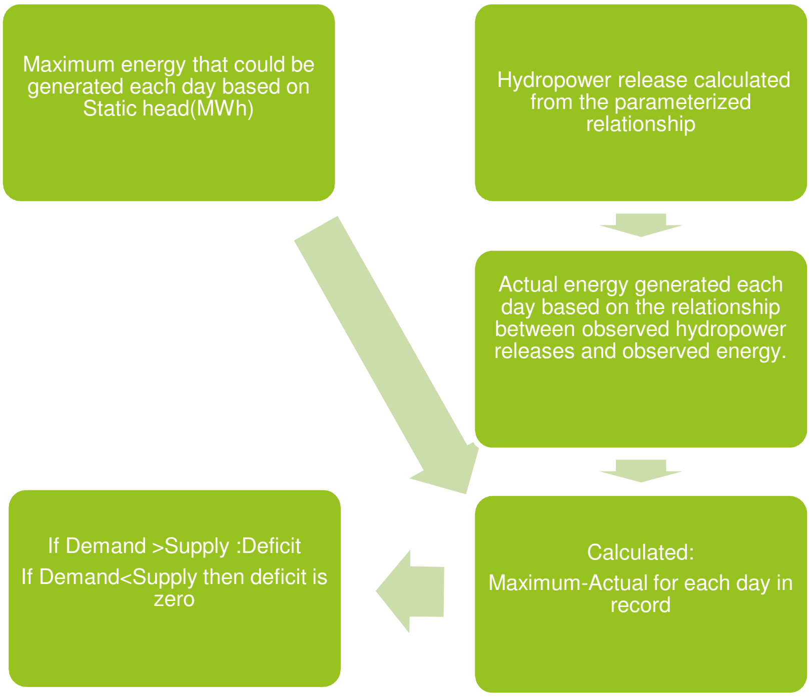

Figure 6: *Hydropower deficit development schematic*

It is noted that no information has been made available for hydropower demand. It is preferable for the deficit to be driven by demand data but in the absence of these, the maximum capacity is deemed as an acceptable way to represent the hydropower deficit.

8.3 System constraints

- Dam level : One constraint is currently imposed on the optimization model is that the dam water-level cannot exceed the dam safety level of 129.54m. However, with the current 23-year inflow series the water-levels do not exceed about 128.5m, which is to be expected given the way in which the rate of gate opening accelerates as the water-level rises. If operations during flood conditions were to be part of the optimization, it would be useful to test this constraint more rigorously, if stochastic inflow data were to be used.
- Reservoir discharge limits : The total outflow discharge limits are constrained by the capacity of the hydropower turbines and the spillway capacity. These are embedded in the model.
- Power generation limits: These are based on information on turbine output received. This has been discussed under the data analysis section.

- Environmental flow : In the absence of specific information with regards to environmental flow requirements, the assumption was made that there is a minimum flow required to sustain ecological flows to the river downstream of the dam, including downstream of the irrigation offtake weir. The minimum outflow from the dam has been recorded as 1.98m3/s and this has been used as the environmental flow required by the system. This environmental flow carries on past the irrigation weir and is not used as irrigation water.

9. RESULTS

The simulation-optimization model has been run for three different cases:

a) Baseline case: This model uses the A and B parameters set by adopting the theoretical minimum turbine operating reservoir level determined from the analysis of the relationship between turbine output and static head. The level corresponding to the A parameter is 115m. The initial parameters for this model version are set out in Table 1.

b) Sensitivity analysis - two objective run: This uses the baseline, but only two objective functions have been used in the optimization procedure: hydropower and irrigation. The third objective, flooding, has not been considered.

c) Sensitivity analysis - a lower A parameter reservoir level: This uses the baseline model, with the exception that the reservoir level assigned to the A parameter is 113 m. This is to account for the discrepancy between the estimated theoretical minimum turbine operation level and the lowest level that the turbine has been operating based on the data available.

In this paper we present the results for the baseline case.

9.1 Baseline

Figure 9 shows the results from the baseline simulation-optimization run. Since there are three objectives, the results ideally need to be viewed in 3D. Figure 9 shows three different 2D visualisations of the 50 points that form the Pareto set, with the third objective function represented with colour.

The symbols used for the Pareto set in Figure 9 distinguish between points which are superior to the historical operating pattern in terms of all three objectives (shown as triangles) and those that are not (shown as circles). The value of the three objectives corresponding to the historical operating pattern is marked with a large open square. As would be hoped, all points in the Pareto set show an improvement on the historical point according to at least one objective. The subset of 9 out of 50 points that are improved in terms of all objectives may be more attractive scenarios for the dam operators to consider implementing.

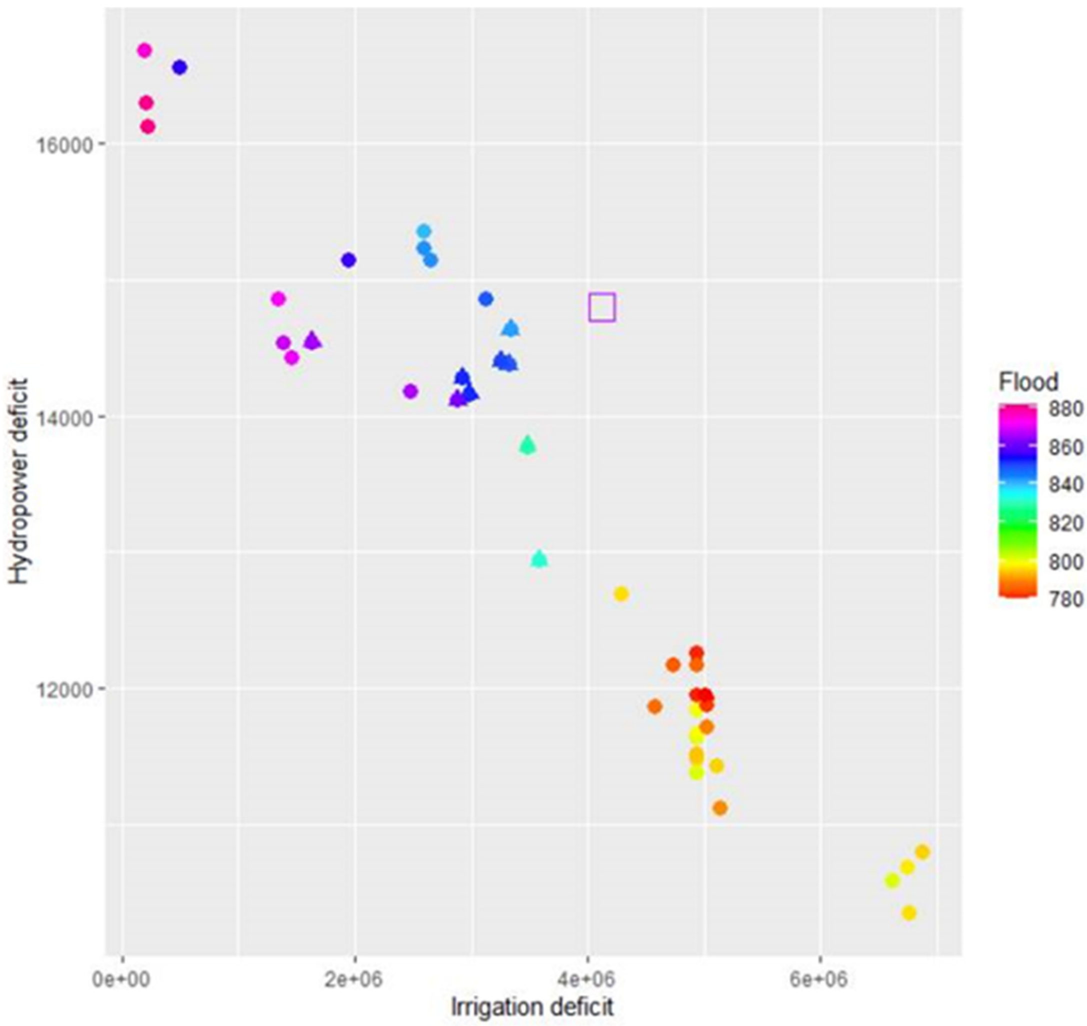

Figure 7*: Baseline optimization – Pareto set, viewed as 2D plots for each pair of objectives, with the third objective shown in colour. Triangles show the subset of points in the Pareto set which are superior to the historical operating pattern (marked with the large square) in terms of all three objectives.*

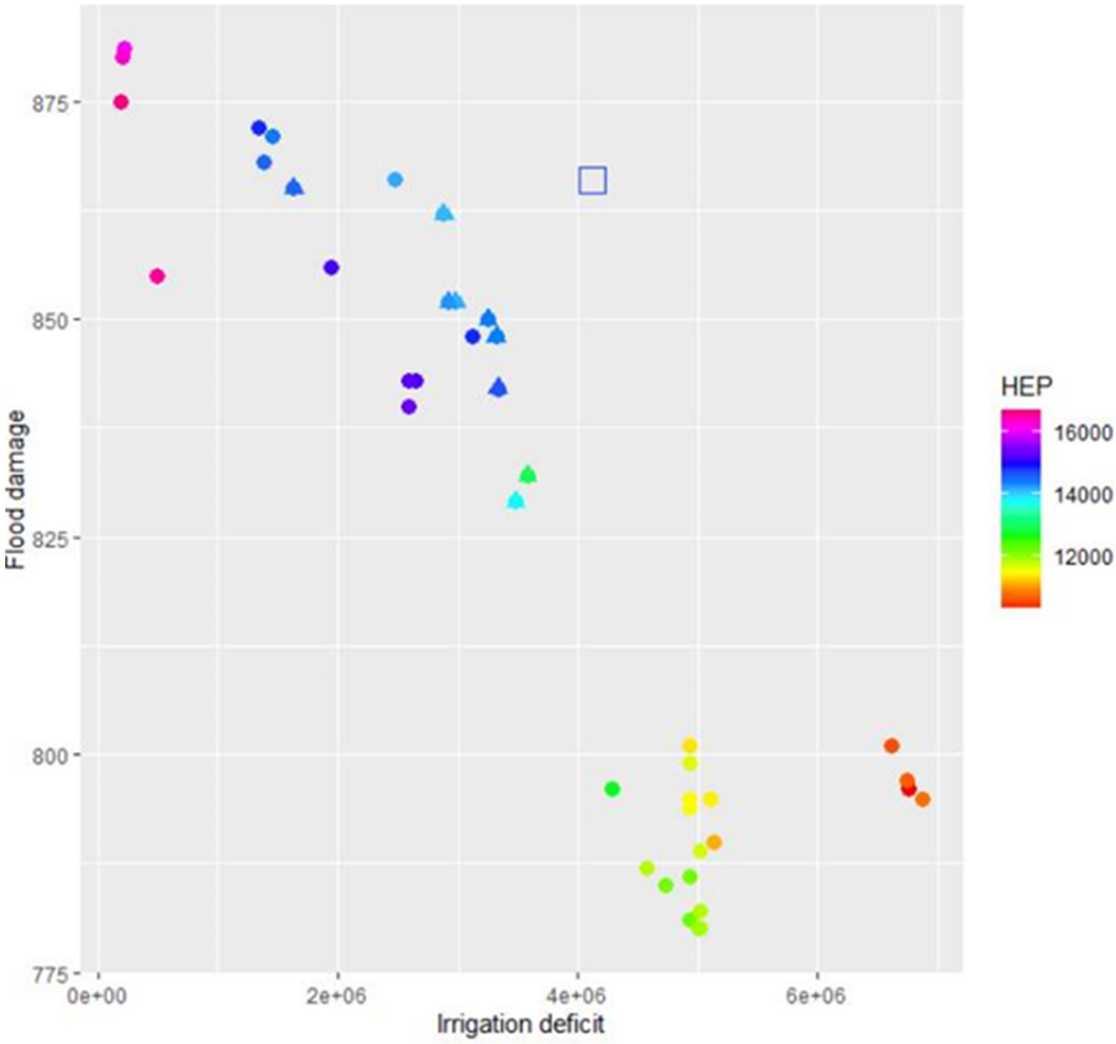

Figure 8: *Baseline optimization – Pareto set, viewed as 2D plots for each pair of objectives, with the third objective shown in colour. Triangles show the subset of points in the Pareto set which are superior to the historical operating pattern (marked with the large square) in terms of all three objectives.*

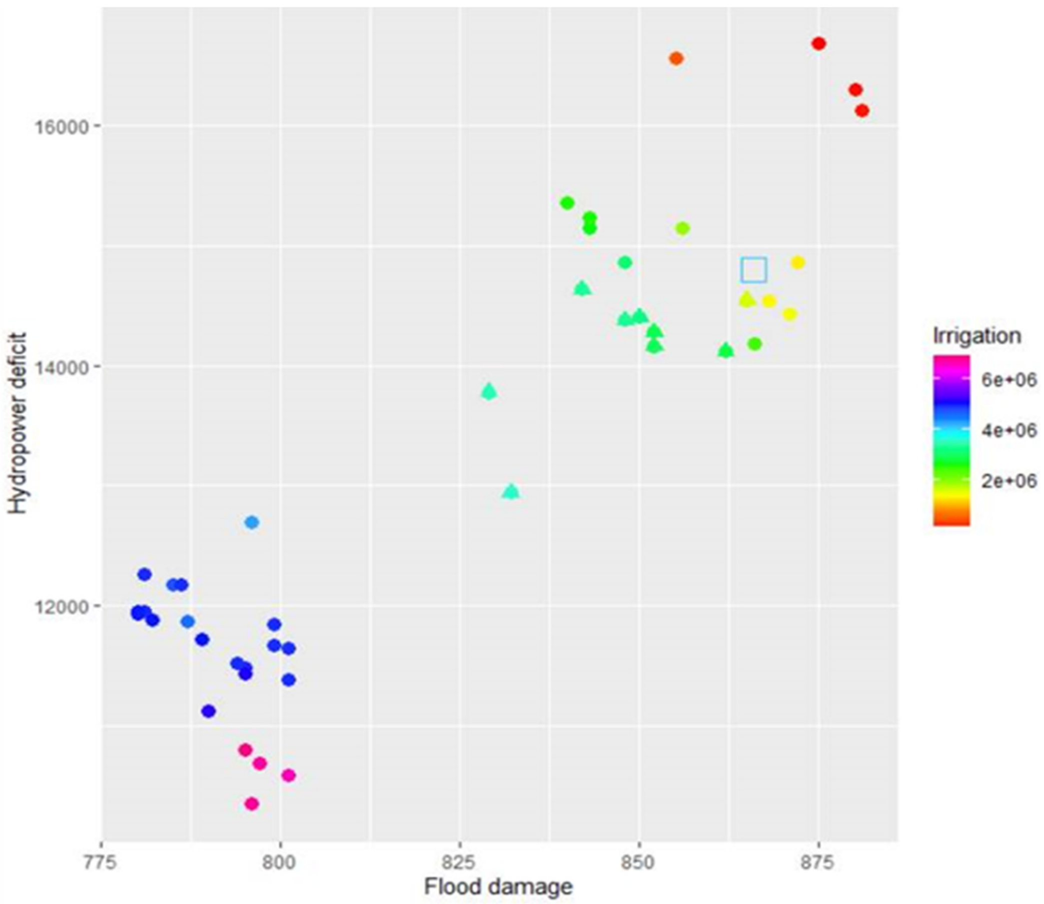

***Figure 9**: Baseline optimization – Pareto set, viewed as 2D plots for each pair of objectives, with the third objective shown in colour. Triangles show the subset of points in the Pareto set which are superior to the historical operating pattern (marked with the large square) in terms of all three objectives.*

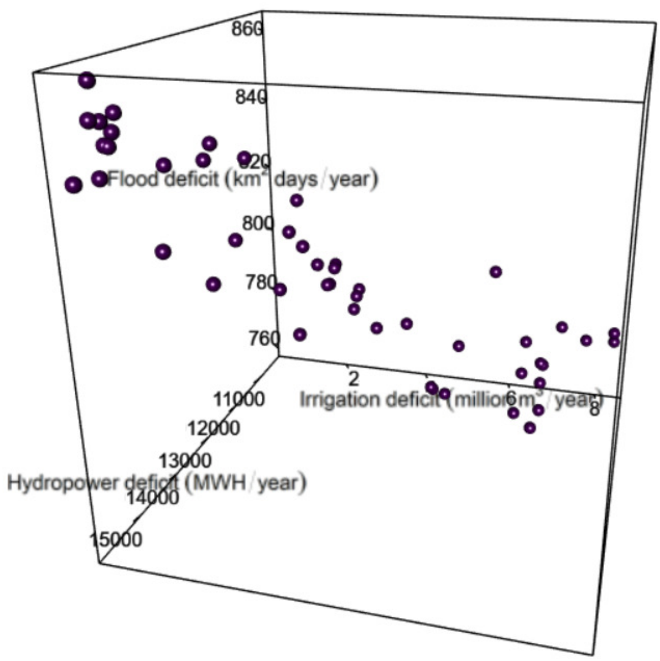

***Figure 10**: Snapshot of the 3D rotatable visualization of the Pareto set.*

Figure 11 compares the final Pareto set with earlier solutions in the optimization process. This shows that the optimization process works as the non-dominated solution set is producing collectively reduced objective function values in comparison with the starting objective function values.

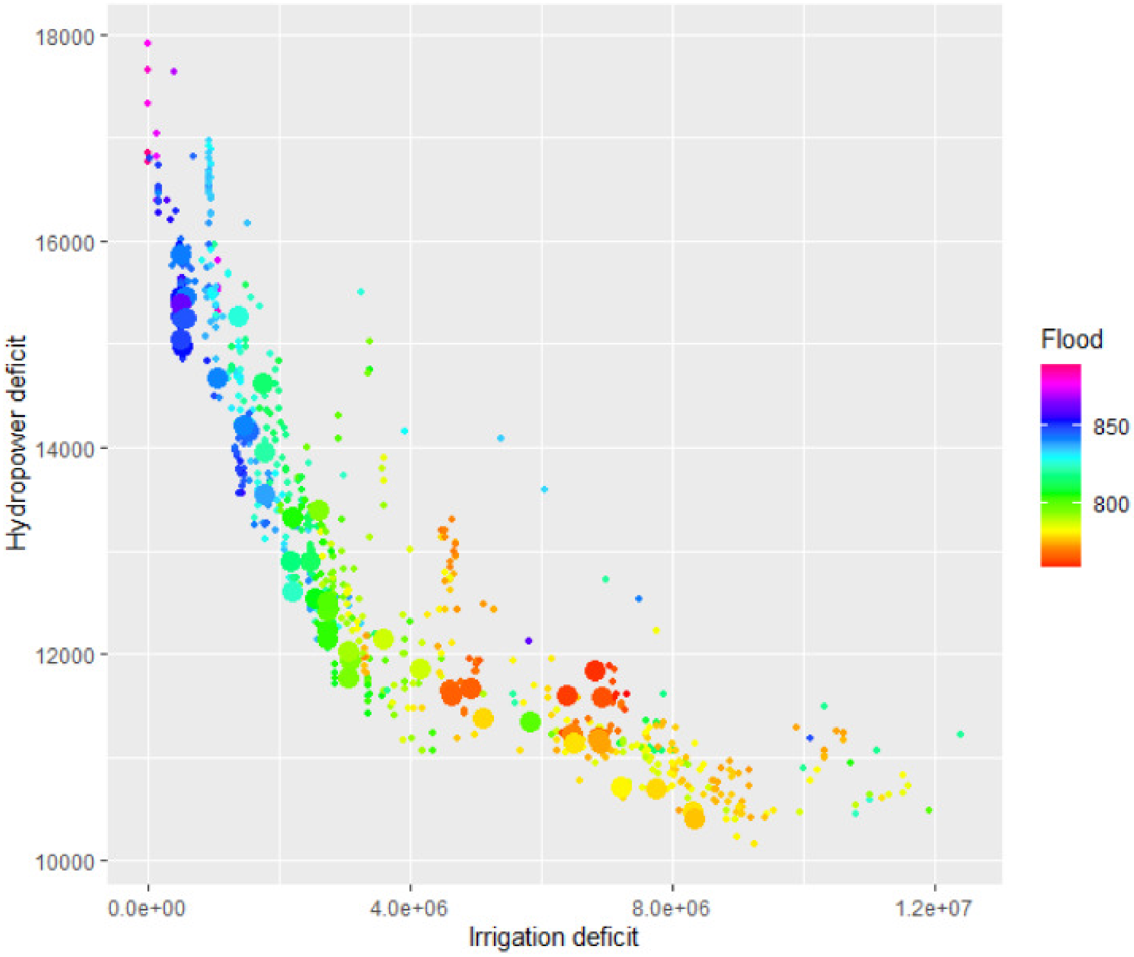

Figure 11: *Baseline optimization - comparison of Pareto set (large points) with earlier solutions in the optimization process (small points)*

Figure 12 shows the set of rule curves produced for each of the three seasons for all the solutions in the Pareto set, including the initial values.

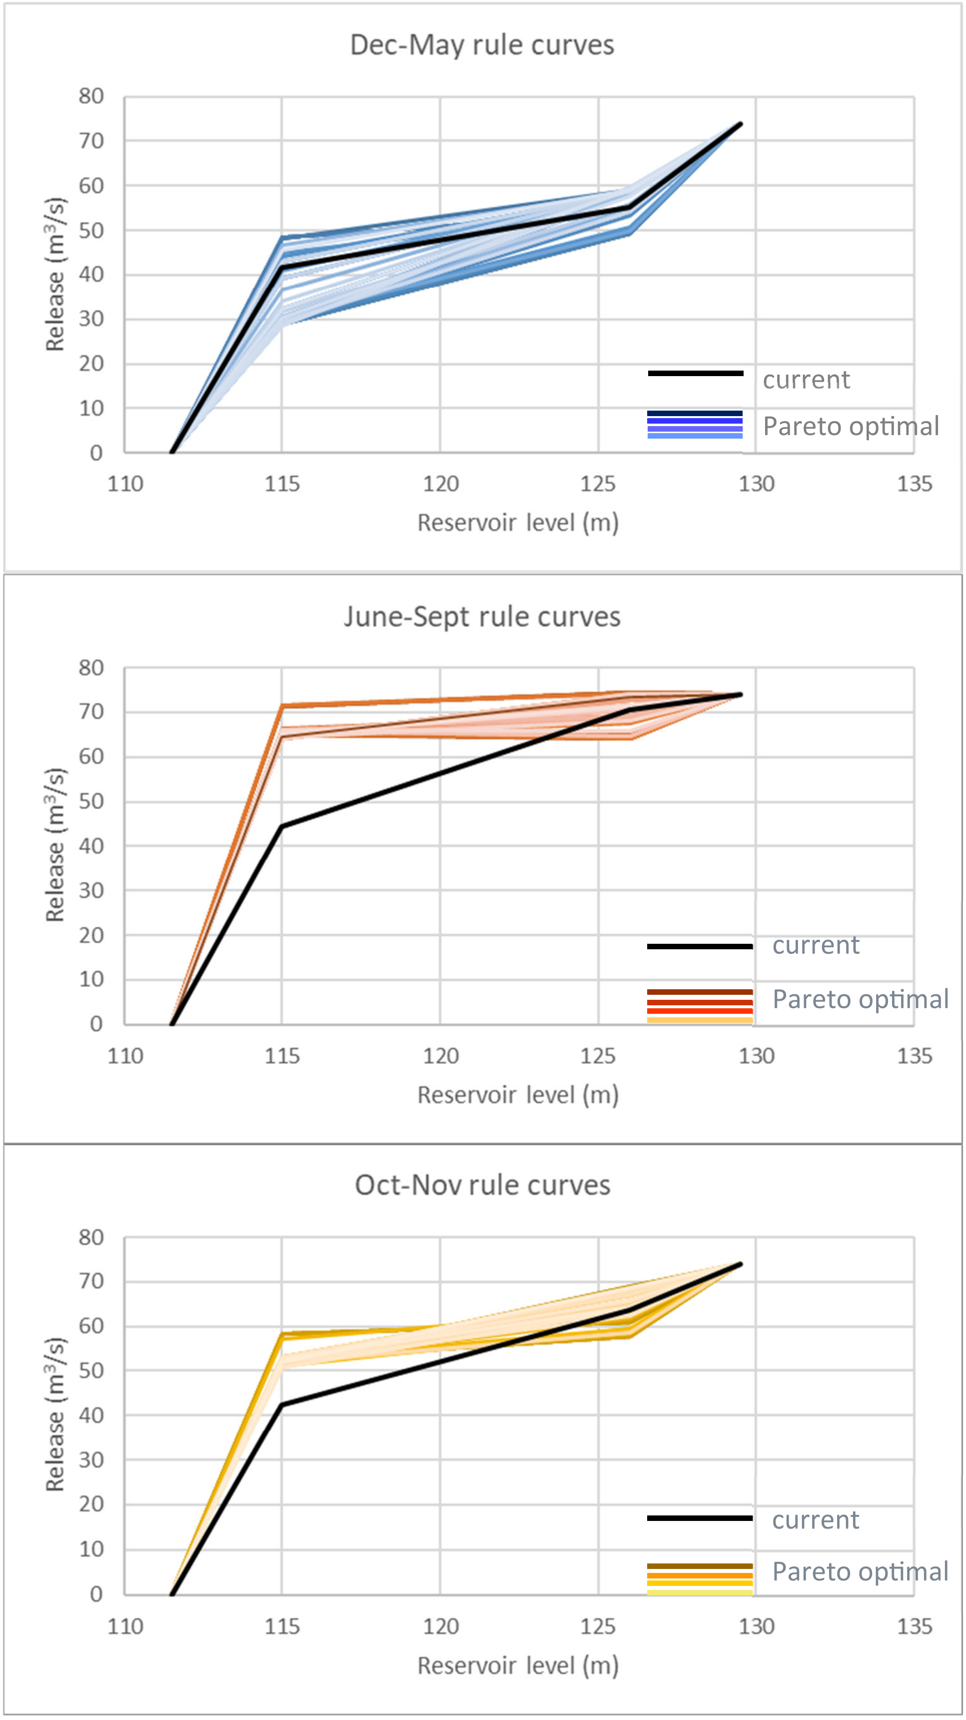

***Figure 12**: Rule curves for all points in the Pareto set, compared with initial rule curve.*

The main conclusions from reviewing the results are:

- None of the three objectives could be minimized to zero, although irrigation deficit could be minimized to close to zero at the expense of high hydropower deficit and high flood damage.
- The extent to which the three objectives compete can be seen from the patterns in the 2D plots (Figure 9). None of the points in the Pareto set have low values for both hydropower and irrigation deficits, or for both irrigation deficit and flood damage. In contrast, it is possible to achieve low values for both hydropower deficit and flood damage, at the expense of a high irrigation deficit. On the hydropower deficit-flood damage plot, for any given irrigation deficit (e.g., as shown by the green points), there is a conflicting relationship between hydropower and flood damage.
- The amount of scatter in the plot indicates the degree to which the pair of objectives is linked. For example, there is relatively little scatter on the hydropower-irrigation plot, indicating that improvement in one objective generally leads to worsening of the other, irrespective of the value of the third (flood) objective.

Additional in-depth investigation would be needed to fully understand the reasons for the relationships between the objectives. They are probably a function of the fact that irrigation demand and spilling through the flood gates are strongly seasonal whereas hydropower generation does not show strong variability throughout the year. Peak irrigation demand is in March-April, which may be why the rule curve that provides the best benefit for irrigation (out of the points in the Pareto set) has relatively low releases during the dry season (December to May). The effect that this can be seen in Figure 13 and Figure 14, which show the annual pattern of water level and dam release, as an average over all years in the simulation, for three points in the Pareto set.

- The point that achieves the best result for the irrigation objective leads to high water levels through the whole year. Water is conserved in the dam during December to March, meaning that by April to June, irrigation releases can be higher than they would have been under other rule curves because the water level is so much higher.
- In contrast, the “best hydropower” rule curve leads to high releases through the turbines during November to March; by April the water is starting to run out and so releases have to be curtailed, but this temporary reduction in generation is more than compensated for by the higher generation at other times of year.
- The “best flood” rule curve follows a similar pattern to the best hydropower rule curve leads to high generating releases during December to April, so that by late April the reservoir is starting to run dry but then the flood season starts and replenishes the reservoir. With the current model setup, the flood damage reduction is constrained by the way the spillway opening rules are currently represented in the model and the reduction to the flood damage is dependent on changes to hydropower releases. Therefore, there is limited ability for the optimization to reduce the flood damage, which, in a 3-objective optimization context, ranges from approximately 761 to 862km^2 days/year.
- All three rule curves are very similar over the Sept/Oct/Nov period which is due to the flood season being so dominant in filling up the reservoir, regardless of what the water-level was at the onset of the flood season.

- Another way to compare the effect of the different rule curves is in terms of the volume of water spilled through the spillway gates. This varies little between the points on the Pareto front. The "best hydropower" and "best flood" points lead to almost identical mean annual spill volumes. The "best irrigation" point is associated with a 15% higher spill volume than the other two points. This is to be expected because the spilled water can potentially be abstracted for irrigation whereas it will be lost from hydropower. The limited variation in spill volumes is a consequence of the way that the spillway gate rules have been fixed, and of the fact that the reservoir tends to fill up during every monsoon season, regardless of the preceding pattern of hydropower releases.
- With regards to the rule curves (Figure 12), the optimization shows that generally for the December-May period the A and B parameters are higher than the ones originally assigned to the model as current operations.
- For June-September the A parameter from the optimization for all points appears to be much higher to than the one originally assigned (64-72 m^3/s as opposed to an initial value of 44.3 m^3/s). The rule curve then is very flat and derived B values from the optimization range between 60-71 m^3/s (Figure 12).
- The rule curves that achieve the best results for flood damage and hydropower deficit reductions are very similar and result in similar annual reservoir water-level profiles.

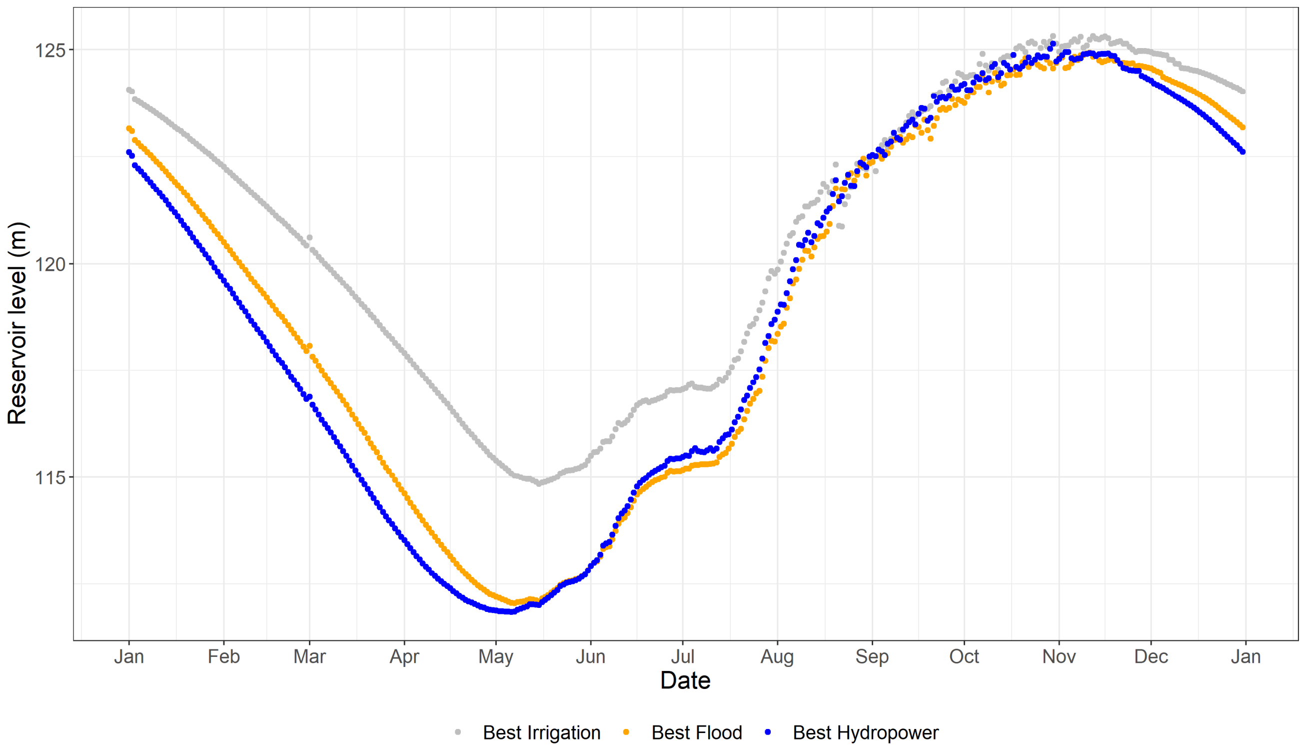

***Figure 13**: Typical annual water level profile for the three-best single-objective solutions in the Pareto set.*

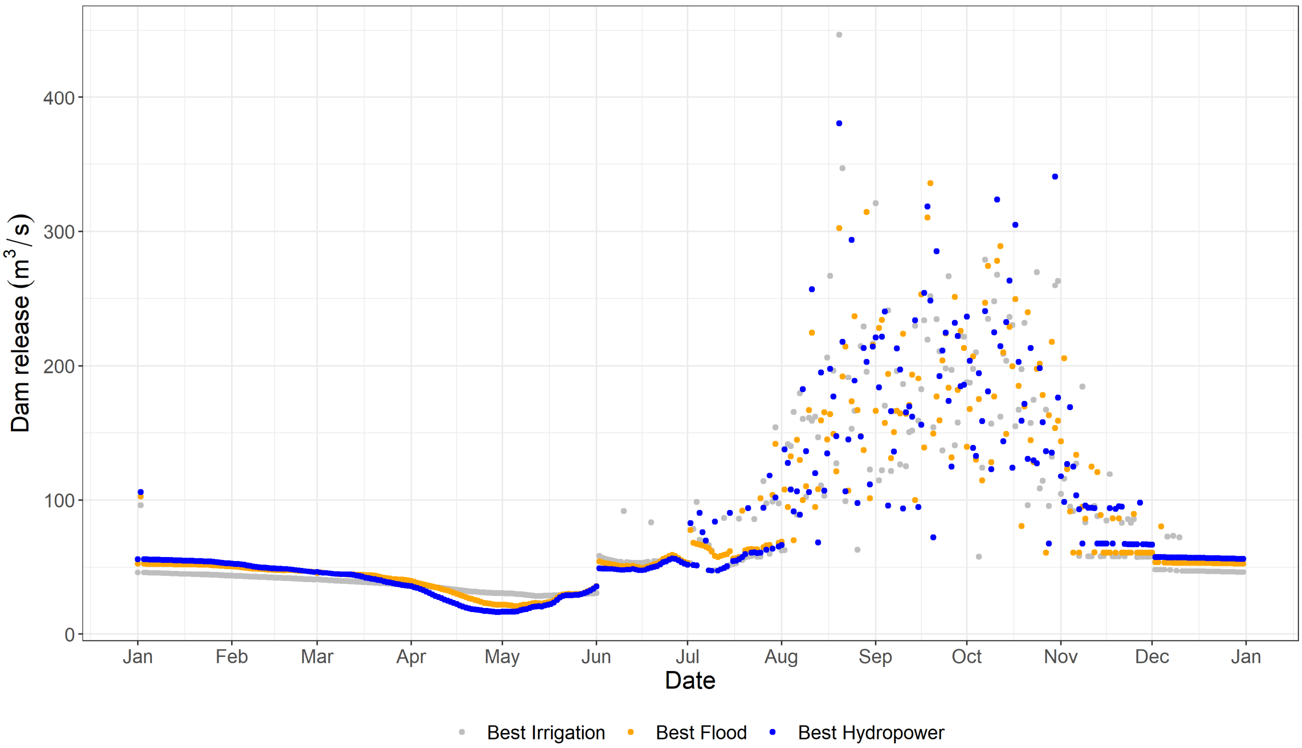

***Figure 14**: Typical annual variation in total reservoir release for the three best single-objective solutions in the Pareto set.*

10. CONCLUSIONS AND RECOMMENDATIONS

10.1 Conclusions

The primary purpose of this study is to provide a proof-of-concept, demonstrating the application of multi-objective optimization to a reservoir operating policy. The study has shown that it is feasible, with a parametric approach to rule curve definition, to undertake multi-objective optimization for dam operations in an efficient and practical manner.

The main challenges encountered during the study were:

- Fully understanding the historical reservoir operating policy, i.e., the way in which releases from the reservoir have historically been managed, in the absence of comprehensive written procedures or an adopted rule curve.
- Parameterizing the rule curve in a way that reproduces the historical reservoir releases and allows enough flexibility for the optimization algorithm to work, including with seasonal variation, while avoiding an excessive number of relaxation parameters.
- Setting up the objective functions in a way that creates meaningful links between model output and societal outcomes. Both the hydropower and flood objective functions would benefit from further examination.
- Representing the reservoir routing effect and implementing the release rules in a spreadsheet, so that the optimization could be carried out using XLOptimizer. In particular, it was necessary

to compromise and use a daily time-step, given the resolution of the inflow data and the constraints of model run-time. This limited the ability of the model to represent sub-daily spillway gate movements during flood events.

The XLOptimizer add-on proved very straightforward to use. It is low cost, including a free version that offers sufficient functionality.

Multi-objective optimization produces a wealth of results, and it is worth allowing ample time and budget for exploring the results. As well as visualizing the solution space defined by the objectives, it is also informative to explore the Pareto optimized parameter combinations and to investigate how these parameters affect reservoir operation, i.e., reservoir levels and releases. Exploring the results can lead to an improved understanding of the problem and may lead to a better definition of the objectives, parameterized rule curve and constraints.

Visualizing the solution space is important. In 2-dimensions the Pareto set is a front which is simple to visualize. In 3-dimensions the Pareto set is a surface which is much more difficult to visualize. For reporting 3D results it is often necessary to reduce to 2D plots showing the third dimension in color, or as a family of contours. In 4 or more dimensions the Pareto set is a hyper-surface which is a lot more difficult to visualize. Six suggestions for visualizing results in many dimensions are given by Ibrahim et al. (2016), who also propose a new approach, 3D Radial Coordinate Visualization.

Within the limited budget available for this project, it has been possible to produce meaningful results. A full multi-criteria optimization project would include additional work, some suggestions for which are mentioned in the recommendations section below. Further work would benefit from involving the reservoir operators more closely.

10.2 Recommendations for any future next steps

Further exploration of results

- The Pareto set developed in this study is worthy of further exploration, which may lead to more insights into improving the set-up of the optimization problem.
- Visualization of results plays a key role to the solution set selection process, which is typically the next step following the current analysis. This paper currently presents results by showing 2D and 3D objective space plots. This is the most commonly used visualization of the obtained solutions. From these, both convergence of the solutions and diversity and spread can be seen.
- Inflow series available: The routing model which is embedded in the optimization-simulation model uses a 23-year daily timeseries. During that period, the reservoir has not reached a level higher than 128.4m. More extreme conditions, whether very wet or very dry, could provide a more robust test of the various control rules that are trialed during optimization. Modelling a period of extremely high inflows could help to examine the consequence of the candidate control rules for dam safety, as a complement to the current approach of imposing fixed rules for gate openings once the water level exceeds a threshold. One way to test more extreme inflow conditions would be to generate artificial inflows using a stochastic model of rainfall or river flow.
- Since the model uses a daily timestep, it can offer guidance on overall flood alleviation, but little guidance on a sub-daily basis. When the dam has reached a level that requires the spillway gates

to be operated, this must be done on a sub-daily basis and further work on incorporating short-term flood alleviation operations would be required to model this comprehensively. It would be necessary to ensure that any changes to the operations of the flood gates were able to safely pass the design flood.