Optimal Storage Placement in River Systems

Sophie Abrahams

Flooding is an ongoing risk in the UK and globally, which is increasingly exacerbated by more intense weather events due to climate change. *Leaky-dams* are often used as part of *natural flood management*, a holistic, catchment-wide approach to flood prevention. Leaky-dams prevent flooding by temporarily storing water during heavy rainfall, slowing the propagation of water down the river system.

This project is presented by Arup, with the aim of creating a framework to optimally position leaky dams in a given catchment area. Arup wish to minimise the cost of dam construction while preventing the flooding of a chosen urban area during an extreme (e.g. once in 100 year) storm event.

Model

We represent each *segment* of river in a catchment as a *node* in a network, as shown in the small example in Figure 1. Each node is a mini-system that receives inflow of water from upstream, rainfall and groundwater, and outputs water downstream. When a dam is built on a segment, the behaviour of that mini-system changes, reducing the outflow from that node.

The most downstream node in the network is defined as the *point of interest*: this corresponds to an area that we wish to prevent from flooding. The point of interest in Figure 1 is a town located at node 5. A storm is simulated by inputting a large volume of water to all nodes over a short period of time.

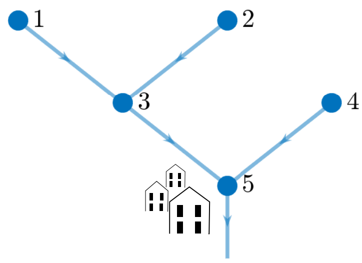

Figure 1 – A small river network. The 5 segments of river are represented by 5 nodes. The point of interest is a town located at node 5.

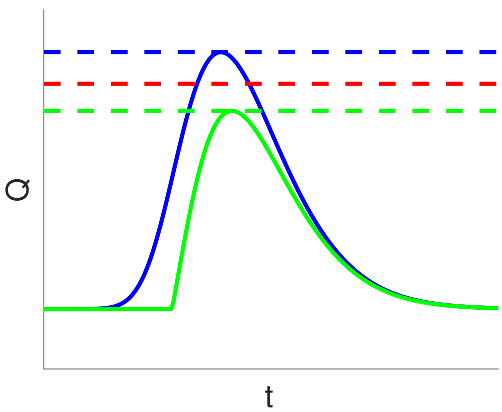

Figure 2 – Plot of flow rate, Q, at the point of interest during a storm for the cases of no dams (blue) and with dams (green). The peak flow rate, Q_{peak}, is shown for the two cases in the same colours. The red line is the critical flow rate, Q_{crit}.

The *flow rate*, Q, out of each node for the duration of the storm is calculated numerically. Plotting the flow rate at the point of interest, we compare the *peak flow rate*, Q_{peak}, to a *critical flow rate*, Q_{crit}. If Q_{peak} rises above Q_{crit}, as in Figure 2, flooding occurs.

The location of dams in the network is represented by a binary string, $\boldsymbol{\alpha}$, where "1" means a dam is located on that segment of river, and "0" means no dam. In the 5 node example in Figure 1, if we choose to build a dam on node 3, we have a string of length 5 with a "1" at position 3: $\boldsymbol{\alpha} = (0, 0, 1, 0, 0)$.

Engineering and Physical Sciences Research Council

Optimisation

The optimal solution is the $\boldsymbol{\alpha}$ that has the lowest *cost* for a given network and a given storm event. The cost of $\boldsymbol{\alpha}$ is defined as construction cost + penalty. The penlty is added for solutions where Q_{peak} is above Q_{crit} and therefore flooding occurs.

A potentially optimal solution is found using a Genetic Algorithm (GA), an optimisation algorithm that can be applied to a wide range of problems. GA begins begin with an initial *population* of randomly generated binary strings called *individuals*, that represent possible $\boldsymbol{\alpha}$. The optimisation steps of GA are based on concepts of natural selection in gene development: individuals undergo *mating* and *mutation* for a number of *generations*. Individuals are chosen at random to mutate, and a number of the entries in their string are *toggled* from "0" to "1" or vice versa. Pairs are selected to mate, with the lowest cost individuals having a higher probability of being selected. Mating is performed by swapping entries in the first half of the binary string to create two *offspring*. Mating allows low cost solutions to be further improved while mutation allows random possibilities to be explored. This process is repeated for a fixed number of generations.

We attempt to improve upon the performance of GA by creating *Salmon Algorithm (SA)*, that is tailored to the specific problem of locating dams in a river network. Instead of mutating by toggling bits – which corresponds to adding or removing a dam – we select a dam in the individual and move it one step upstream. In our example, nodes 1 and 2 are one step upstream of the dam at node 3, and are therefore candidates for the new location. If we select node 2 to be the new dam location, $\boldsymbol{\alpha} = (0, 0, 1, 0, 0)$ mutates to $\boldsymbol{\alpha} = (0, 1, 0, 0, 0)$. We postulate that there is a *sweet spot* for a dam: a location that has enough inflow for a dam to fill during a storm, but not so much that the dam overflows. Pushing dams incrementally upstream may help to find this sweet spot.

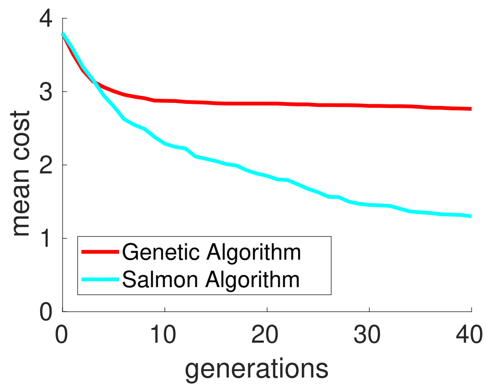

Figure 3 – Plot of cost at each generation averaged over 100 networks.

We test GA and SA by running them on 100 networks of 50 nodes for 30 generations. The average cost over the 100 networks at each generation is plotted in Figure 3. From the plot we see that by generation 40, SA has, on average, found a solution with approximately half the cost of that found by GA.

To further improve the optimisation, we attempt to give the algorithm a running-start by reducing the cost of the initial population. When generating the binary strings, we test the effect of weighting the probability of choosing each node according to different parameters. We find that weighting the nodes by dam size – the potential size of a dam varies greatly between different areas – reduced the average cost of the initial population by over half, in comparison to choosing nodes uniformly.

Conclusions

We have proposed two methods of tailoring GA to the specific problem of finding the lowest cost dam locations. Choosing an initial population based on potential dam size produced a significantly lower cost initial population of solutions, while the mutation method of SA helped to push the cost lower over 40 generations in comparison to GA.

This work will enable the more effective design of natural flood management interventions, helping to reduce flood risk for vulnerable communities, while minimising the impact on the local environment and on the budgets of local and national public bodies. Incorporating optimisation of the dynamic behaviour of the flood catchment represents a significant step forward in our capability in this area, and will directly influence future designs. The work also has the potential to be extended into sustainable flood management designs in the urban realm, an even larger space of opportunity.

Paul Lynch, Senior Engineer, Arup

Engineering and Physical Sciences Research Council