JBA Trust Challenge: A Risk-based Analysis of Small Scale, Distributed, "Nature-based" Flood Risk Management Measures Deployed on River Networks

Sheen Cabaneros, Federico Danieli, Giuseppe Formetta, Raquel Gonzalez, Michael Grinfeld, Barry Hankin[*], Ian Hewitt[†], Teague Johnstone, Alissa Kamilova, Attila Kovacs, Ann Kretzschmar, Kris Kiradjiev, Sam Pegler, Graham Sander, Clint Wong

February 27, 2018

Abstract

We develop a network-based model of a catchment basin that incorporates the possibility of small-scale runoff attenuation features ('leaky dams') on each of the edges of the network. The model is forced by a prescribed runoff to each node and predicts the time series of discharge throughout the network. It can be used to analyse the benefit and risk associated with adding dams at specific network locations. We demonstrate the model using idealised one-dimensional and two-dimensional networks, and explore the risk of cascade failure. We discuss the formulation of an optimisation problem to decide on the best dam placements for a given catchment, and give suggestions for future directions.

1 Introduction

Many communities at risk from flooding seek low-cost methods of reducing their flood risk without the capital expense that large-scale engineering solutions demand. There is growing interest in the use of so-called 'nature-based' measures, including tree planting and soil structure management, to prevent fast overland flow. One such measure is small-scale runoff attenuation features (RAFs), typically 'leaky' dams or barriers made from wood that allow low flows to pass under or through but hold back high flows, providing temporary water storage (figure 1). It is hoped that a large collection of such features deployed in a catchment may hold back enough flood water to mitigate flood risk downstream. However, the effectiveness of such systems of RAFs has not been fully tested or explored.

This report details work done at the study group to develop a modelling strategy in which the effectiveness of a collection of RAFs can be assessed. Some work has been done in this direction already by Metcalfe et al. [2017], and we were given a steer towards a network-based analysis, in which individual RAFs could be described

[*]JBA Consulting

[†]Corresponding author: hewitt@maths.ox.ac.uk

Figure 1: An example of a leaky dam.

as lumped elements with an arbitrary stream network. In particular, it is desired to understand the risks of potential failure modes of a system of RAFs, including issues such as

- possible synchronisation or de-synchronisation of peak flows,
- cascade-like failure modes (a collapse of one RAF triggering further collapses downstream),
- inefficient storage provided by some RAFs due to non-optimal placement within a network,
- degradation of the features over time due to ageing / erosion.

Quoting from the problem description, "We are seeking to explore whether there is a mathematical strategy to give insight (and basic rules) into assessing the effectiveness and resilience of many small-scale nature-based flood risk management interventions in complex river networks. This will help us to describe distinctive deployment and maintenance strategies, e.g. concentrated on headwaters, concentrated on lower reaches, maintain every 1/2/10/50 years etc. (leading to different performance profiles through time)."

We begin in section 2 by setting up a network model for an arbitrary stream network, breaking the stream up into segments that may each potentially contain a runoff attenuation feature. Rules for the storage and discharge (flux) in each segment are prescribed based on the slope, stream cross-section and roughness. Modifications of these rules to account for the effect of a leaky dam are developed. The model amounts to a series of coupled ordinary differential equations (ODEs) that are solved numerically given prescribed runoff inflow.

In section 3 we explore solutions for some simple networks forced by typical flood-like runoff, focussing on the response of the discharge at the downstream end of the network. We examine the response to failure of the dams in section 4, including cascade failure. We also discuss appropriate ways to set up an optimisation problem on either the position or design of the dams in section 5.

2 Network model

2.1 Stream networks

We construct a network model in which segments of a channel ('reaches') are described in a lumped fashion. The primary variables are the average cross-sectional area A_i and discharge Q_i, which flows into the next channel segment downstream. The channel segments correspond to nodes of a graph, and the edges that transfer water downstream can be thought of as potential dams (i.e. the positions at which dams might be added). The connections between the channel segments are described using an adjaceny matrix (a_{ij}). The ith row of this matrix is all zeros except for in the jth column, where j indexes the node immediately downstream of the ith node. The positions and connections between the channel segments, as well as their lengths, and slopes, might be determined from studying a two-dimensional digital elevation model (DEM) such as that used by Metcalfe et al. [2017]. In this report we will use idealised network structures, with uniform widths and slopes.

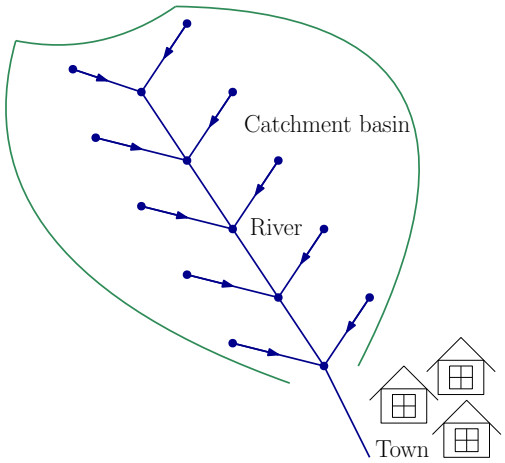

Figure 2: An example channel network. The main trunk is fed by tributaries from either side in a herringbone pattern. Of primary interest is the discharge flowing from the last node, corresponding to the area most at risk from flooding.

Taking ℓ_i to represent the length of the channel segments, volume conservation requires

$$\ell_i \frac{\mathrm{d}A_i}{\mathrm{d}t} = \sum_{j=1}^{N} a_{ji} Q_j - Q_i + q_i, \tag{1}$$

for each node ($i = 1 \dots N$). The sum represents the fluxes from the immediately upstream nodes and q_i represents the in-flow to each segment from rain/runoff from the surrounding land. It may be more convenient to think of (1) in terms of the water volumes $V_i = A_i \ell_i$ stored in each channel segment.

We assume that the $q_i(t)$'s are prescribed, although in a more complete treatment they might be taken from a two-dimensional model (using the shallow water equations for example), or they might be derived from assumed rainfall using a filter to represent the time-delay due to subsurface and/or overland flow.

Given the known slope of each channel segment S_i (which may be related to the bed angle θ by $S_i = \tan\theta$), we could relate the discharge and cross-sectional area. However, it turns out to be more convenient to express the discharge in terms of the water depth h_i behind the potential dam in each reach. In the case that there is no

dam, or when the depth is below the bottom of the dam, this is simply the average water depth and we can relate this to the cross-sectional area.

The relationship depends on the assumed shape of the channel, and on a parameterisation of turbulent flow. If we assume for simplicity that the channel has a rectangular cross-section with fixed width w_i, and use Manning's law, we have

$$A_i = w_i h_i, \qquad Q_i = \frac{w_i^{5/3} h_i^{5/3} S_i^{1/2}}{(w_i + 2h_i)^{2/3} n}, \tag{2}$$

where n is the Manning roughness coefficient and S_i is the slope. Since we can then relate Q_i directly to A_i (by eliminating h_i), we can interpret (1) as a set of coupled ordinary differential equations for the A_i, forced by the inputs q_i. These can be solved numerically using a variety of methods.

More generally, when we include dams, we write

$$A_i = \tilde{A}(h_i; \cdot), \qquad Q_i = \tilde{Q}(h_i; \cdot), \tag{3}$$

where $\tilde{Q}(h; \cdot)$ and $\tilde{A}(h; \cdot)$ are known functions, and the extra parameters $(\cdot)$ will describe the dam as well as the cross-section and slope. We also define $\tilde{h}(A; \cdot)$ to be the inverse of $\tilde{A}$ (which is a monotonically increasing function of h and therefore has a well-defined inverse). Thus we will still have a direct relationship between Q_i and A_i.

2.2 Modification to include dams

We take h to represent the height of the water behind the dam. The dam has its bottom at height b above the stream bed, and its top at height H. The description of the flow past the dam can then be divided into three cases, $h < b$ (corresponding to the water level being below the bottom and the dam doing nothing), $b \leq h < H$ (when the dam is operating normally), $h \geq H$ (when the dam is overspilling).

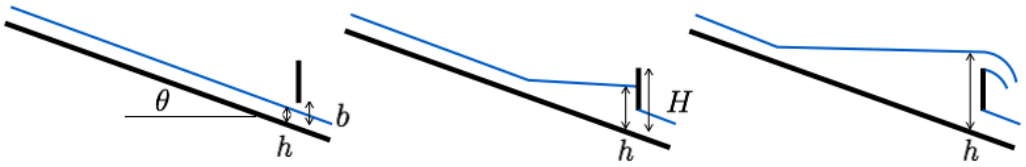

Figure 3: A sketch of the flow past each dam in each of the three cases $h < b$, $b \leq h < H$, and $H \leq h$.

For the first case we use the same Manning relationship as given above to relate Q to h. For the second two cases we adopt relationships from hydraulic theory for the flow beneath sluice gates and over weirs [e.g., Munson et al., 2013]. When the water depth is part of the way up the face of the dam the flow underneath is given by Bernoulli's equation, to be

$$wbh\sqrt{\frac{2g}{b+h}}. \tag{4}$$

An empirical correction factor to account for losses is often included in this formula, but we neglect it for simplicity. The flow through the (leaky) dam is assumed to

similarly vary with the water depth (due to the hydrostatic pressure), and we write this as

$$kw(h-b)\sqrt{2gh}, \tag{5}$$

where k should be interpreted as the dam permeability. When the water depth is above the level of the dam, the over-flow is described as for flow over a weir, giving

$$\frac{2\sqrt{2g}}{3}w(h-H)^{3/2}. \tag{6}$$

The leaky flow through the dam is then

$$kw(H-b)\sqrt{2gh}. \tag{7}$$

In summary, therefore,

$$\tilde{Q}(h;w,\ell,b,H,k) = \begin{cases} \frac{w^{5/3}h^{5/3}S^{1/2}}{(w+2h)^{2/3}n} & 0 \le h < b, \\ w\sqrt{2g}\left[\frac{bh}{(h+b)^{1/2}} + k(h-b)h^{1/2}\right] & b \le h < H, \\ w\sqrt{2g}\left[\frac{bh}{(h+b)^{1/2}} + k(H-b)h^{1/2} + \frac{2}{3}(h-H)^{3/2}\right] & H \le h. \end{cases} \tag{8}$$

When there is a dam, A_i no longer represents the uniform cross-section of the stream, but rather its average over the length. It is most straightforward to calculate the volume $A\ell$ in terms of the water depth h, for each of the three cases mentioned above. This again depends upon the precise geometry; for the rectangular channel we have

$$\tilde{A}(h;w,\ell,S,b,H) = \begin{cases} wh & 0 \le h < b, \\ wb + \frac{w(h-b)^2}{2S\ell} & b \le h. \end{cases} \tag{9}$$

The terms in the second expression here represent the volume of water in the stream up to the depth of the bottom of the dam, plus the volume of water stored in the triangular wedge that forms behind the dam. These relationships are shown in figure 4.

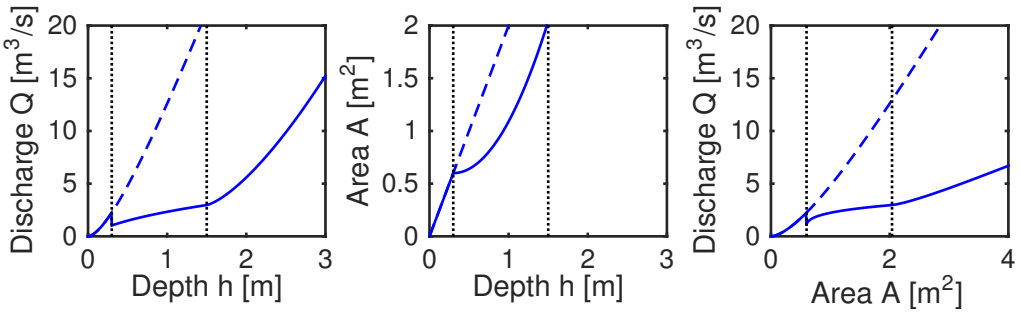

Figure 4: Examples of the relationships between discharge Q, cross-sectional area A, and water depth h, for a rectangular channel of width $w = 2$ m, slope $S = 0.01$, reach length $\ell = 100$ m, and a Manning coefficient 0.01 s m$^{-1/3}$. Dashed lines show the case of no dam (Manning's law for a rectangular channel). Solid lines show the case of $b = 0.3$ m and $H = 1.5$ m, shown by the vertical dotted lines.

2.3 Non-dimensionalised model

We expect that the flow through the dam will be small compared to that under and over it (the fact that it is allowed to be leaky makes it easier to construct, but the

leakiness is not fundamental to its operation). Thus, we assume this can be ignored and set the dam permeability coefficient k to zero for simplicity.

We choose scales, denoted by square brackets, such that

$$[Q] = \frac{[w][h]^{5/3}[S]^{1/2}}{n}, \qquad [A] = [w][h], \qquad [t] = \frac{[\ell][w][h]}{[Q]}. \tag{10}$$

For given typical values of $[Q]$, $[w]$, $[S]$, and $[\ell]$, these determine the scales $[h]$, $[A]$, and $[t]$. Using typical values $[w] = 2$ m, $[Q] = 1$ m^3 s^{-1}, $[S] = 0.01$, and $[n] = 0.01$ s m$^{-1/3}$, $[\ell] = 100$ m, we find $[h] = 0.17$ m, $[A] = 0.33$ m^2, and $[t] = 33$ s.

The dimensionless equations are almost the same as the dimensional ones and are not reproduced here. However, the relevant dimensionless parameters are

$$\alpha = \frac{[h]}{[w]}, \qquad \beta = \frac{[h]}{[S][\ell]}, \qquad \gamma = \frac{\sqrt{2g}[w][h]^{3/2}}{[Q]}. \tag{11}$$

These represent the ratio of depth to width (this is of little importance), the ratio of depth to elevation change across the segments, and the strength of gravity compared to friction. For the values given above we find $\alpha \approx 0.08$, $\beta \approx 0.17$, and $\gamma \approx 0.60$.

The smallness of β is problematic; it indicates that the capacity to hold back a significant volume of water behind the dams is very limited. This is a first indication of why a large number of dams may be required to have even a noticeable effect on the discharge downstream. We suspect that this is an artefact of the assumed uniform width of the channel. It is likely that the locations for the dams may be chosen to dam reservoirs that are wider than the average stream width, where the stream bed is particularly flat, or where there is capacity for significant overflow onto the river banks. We suggest that the contribution to the cross-sectional area due to the volume in the reservoir in (9) is underestimated, and should be increased. Thus we modify the cross-sectional area to

$$\tilde{A}(h) = wb + \lambda \frac{w(h-b)^2}{2S\ell}, \quad H \le h \tag{12}$$

where $\lambda \ge 1$ is this enhancement factor that accounts for a larger volume being stored behind the dam. In practice, this would have to be estimated for each dam location.

2.4 Generalisations for a trapezoidal cross-section

An alternative method to allow dams to store a larger volume of water is to allow for the fact that the river channel likely widens with depth, particularly once the banks are reached and the water can spread to the neighbouring flood plain.

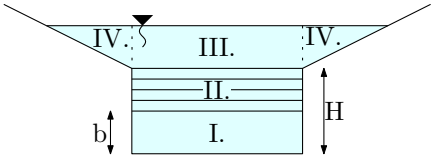

Figure 5: Schematic of a trapezoidal cross-section. The regions I-IV correspond respectively to the region under the dam, the region behind the dam, the region directly above the dam, and the regions to the sides of the dam. The flows through these areas correspond to the four terms in square brackets in (13).

To account for this, we assume that the channel has a cross-sectional shape, comprising a rectangle of width w opening to a trapezium with side slope $1/m$ above the top of the dam at height H. For this geometry the overflowing cases of $\tilde{Q}$ and $\tilde{A}$ are modified to

$$\tilde{Q} = w\sqrt{2g}\left[\frac{bh}{(h+b)^{1/2}} + k(H-b)h^{1/2} + \frac{2}{3}(h-H)^{3/2} + \frac{8}{15}\frac{m}{w}(h-H)^{5/2}\right], \quad H \le h \tag{13}$$

$$\tilde{A} = wb + \frac{w(h-b)^2}{2S\ell} + \frac{m(h-H)^3}{3S\ell}, \quad H \le h. \tag{14}$$

Note that these formulae for a trapezoidal cross-section are not used for the numerical solutions shown in the rest of this report, for which we use the simpler earlier relationships (8) and (9) with the enhancement factor described in (12).

2.5 Alternative versions of the discharge-depth relationship

One potential concern with the above formulation is the discontinuity in the discharge-depth relation when the water depth reaches the bottom of the dam (figure 4). This occurs in the model because the physics used to relate the depth to discharge is different in the two cases of free-stream flow (when we use Manning's law to describe turbulent drag) and flow under the dam (when we use an essentially inviscid formula for flow beneath a sluice gate). There was some discussion at the study group about whether this discontinuity should be there in reality (i.e. does the free surface just touching the bottom of the log dam cause a sudden decrease in flow?). We concluded that this is realistic, although the discontinuity may be smoothed out in reality.

Mathematically, provided the discontinuity in flux involves a reduction as h increases, there should be no problem. As the water depth reaches the bottom of the dam, the flow past it suddenly reduces and the water quickly fills up behind the dam until the depth has increased to allow sufficient flow to balance the inflow from upstream. There could be problems, however, if the discharge suddenly increases when h increases passed b. Examining (8), this would clearly be the case if the slope were sufficiently small, or n sufficiently large. This was not the case with the parameters considered at the study group, and it would require a modification of the formula for flow under the dam when $h \ge b$. In particular, this corresponds to a case when the dominant resistance to flow is still the turbulent drag and the sluice-gate formula is inappropriate. The water depth may be expected to exceed b on both sides of the dam. We do not pursue this eventuality further here.

For the purposes of constructing a very simple network model, we note that the essential behaviour of the dams could be captured using piecewise linear relations of the dimensionless form

$$\tilde{Q}(h) = \begin{cases} wh & 0 \le h < b, \\ w(b + \alpha_1 h) & b \le h < H, \\ w(b + \alpha_1(H-b) + \alpha_2(h-H)) & H \le h, \end{cases} \tag{15}$$

$$\tilde{A}(h) = \begin{cases} wh & 0 \le h < b, \\ w(b + \beta_1 h) & b \le h, \end{cases} \tag{16}$$

for constants α_1, α_2 and β_1. Such a simplification might be warranted if the precise shape and slope of the channels were unknown, or for a theoretical optimisation problem that requires a particularly simple forward model.

Figure 6: An example of a one-dimensional network. The triangles indicate the location of dams which control the discharge Q_i associated with the upstream reach where the average cross-section is A_i.

3 Numerical experiments

3.1 One-dimensional network

In this section we consider a simple example of the model, using the one-dimensional network shown in figure 6. We suppose that each of the channel segments is the same (i.e. equal widths, lengths, and slopes), and the discharge in the final segment is of most interest for the community requiring protection. For these calculations (and all others shown in this report) we use the original flux and area formulas (8) and (9) with the enhancement factor described in (12).

The model is forced with a 'storm' input in the form of a Gaussian,

$$q(t) = q_0 + q_{max} e^{-t^2/\sigma^2}, \tag{17}$$

where q_0 is a baseline inflow (groundwater flow into the channel, say), q_{max} is the peak flood inflow at time $t = 0$, and the flood is spread over a time period σ. In figure 7 we compare the resulting modelled discharge in each channel segment between the case of no dams, and the case of having a dam on each segment (we use only five segments for ease of illustration; using more segments allows for greater potential of reducing the peak discharge).

Figure 8 shows an example when the input has a double peak. In this case, as might be expected, the dams are less effective at reducing the height of the second peak, because they are already holding back a lot of water and have less capacity to store and delay water for the second storm.

3.2 Two-dimensional network

Here we consider a simple two-dimensional network as shown in figure 9. There are more interesting questions to consider about the positioning of dams in this case. For example, if one has funding to build a certain number of dams, which of the channel segments are the best ones on which to put them? Putting them on the central trunk is likely to ensure that they are used, but also means that they may more easily overspill and lose their effectiveness. They may also be more susceptible to cascade failure (as discussed in the next section).

In figure 10 we show two examples of the response to a flood input of the form in (17). In the first case, 4 dams are placed on the main trunk (nodes 1-4), whereas in the second case 4 dams are placed on the upper branches (nodes 5,6,9,10). The discharge from the final segment (node 4) is plotted, along with its maximum value. Both dam placements have the effect of slightly delaying and reducing the peak discharge, with the second design being marginally more effective. This is because the dams near the bottom of the central trunk are overspilling and losing their effectiveness, whereas the dams on the side branches are all having a significant effect. However,

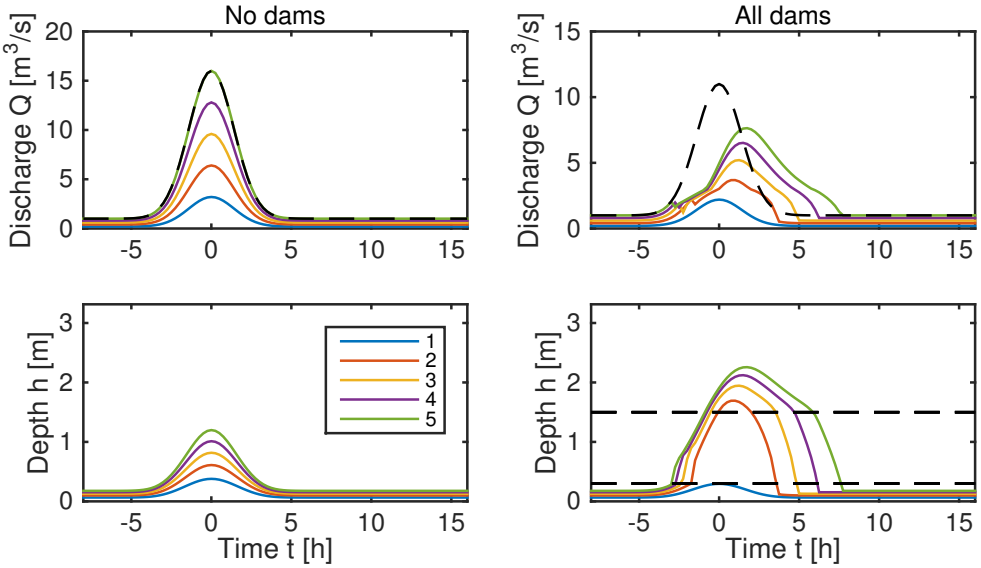

Figure 7: Solutions for a one-dimensional 5-node network as in figure 6, forced by uniform inflow to each node $q(t) = q_0 + q_{max}e^{-t^2/\sigma^2}$, with $q_0 = 0.2$ m^3 s^{-1}, $q_{max} = 3$ m^3 s^{-1}, and $\sigma = 2$ h. Parameter values are as given in section 2.3, together with $\lambda = 50$, $b = 0.3$ m and $H = 1.5$ m. Left-hand panels show the response with no dams, when the peak discharge is almost identical to the peak cumulative inflow (which is shown by the black dashed line in the upper panels). Right-hand panels show the response if a dam is included on each of the 5 reaches. The dashed lines in the lower panel shows the heights of the bottom and top of the dams.

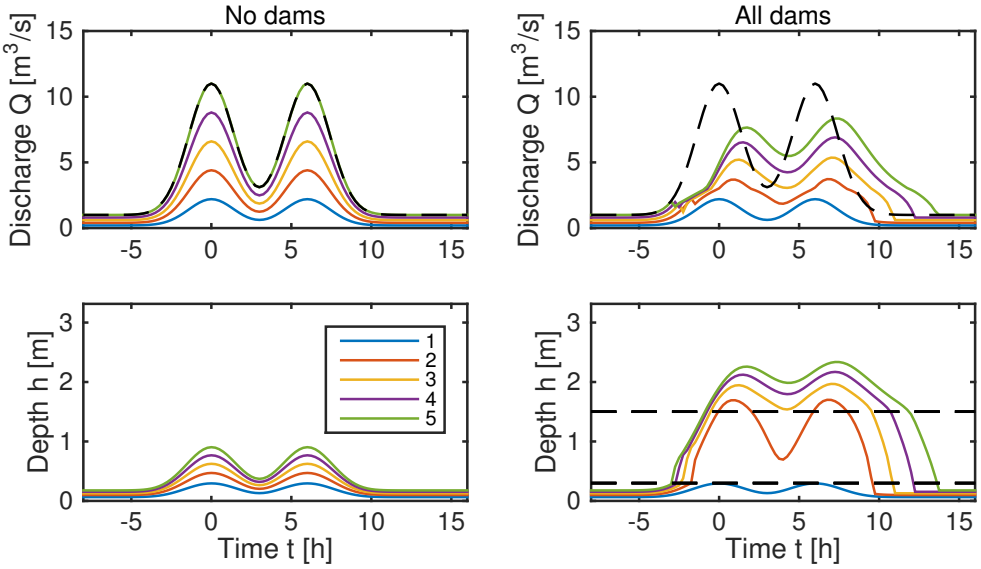

Figure 8: Solutions for a one-dimensional network as in figure 6, forced by a double peaked input to each node. Parameter values are as in figure 7. The dashed line in the upper panels is the cumulative inflow to all nodes, and teh dashed lines in the lower panel show the height of the bottom and top of the dams.

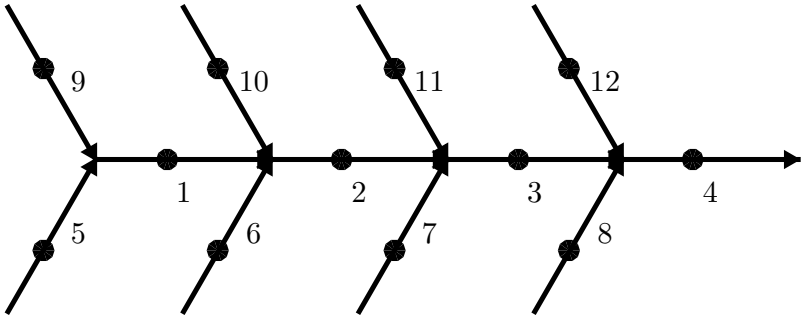

Figure 9: An example of a two-dimensional 'herringbone' network.

for different sized floods, the optimal arrangement can vary, and unfortunately there does not appear to be a clear rule for the most effective dam placement, even in this simple example.

4 Failure mechanisms

4.1 Fragility curves

Fragility curves show the probability of a dam failing when the water exceeds a given level. Different fragility curves have been provided for different conditions of the dam, and show that the worse the condition of the dam, the more likely failure is to occur at lower water levels. The dam's condition is described in terms of 5 different categories. It should be expected that over time, the condition of the dams will worsen so that newly installed dams are in the lowest category, while older dams are in the higher categories, for which the probability of failure at a given water level is higher.

We can investigate the effects of failure by either randomly or deterministically assigning each dam a failure depth h_c based on the probability distribution in the fragility curves.

4.2 Cascade failure

One of the potential risks of installing many dams in a catchment is the possibility that they all collapse at once, creating a flood surge that is larger than would have occurred if no dams had been installed at all. Provided each dam stores only a small reservoir of water, the collapse of one dam on its own should not be too catastrophic. But if the collapse of one dam causes others further downstream to collapse too there is the obvious danger of the surge escalating. This risk may be an important factor in deciding the best placement of dams (perhaps outweighing the efficiency of peak-flow reduction under 'normal' operating conditions).

The main method suggested for analysing this risk is to run an ensemble of simulations of flood events, assigning a failure depth to each dam using the probability distribution suggested by the fragility curve. Ideally, this ensemble should include a range of storm conditions too. Such analysis should in principle use local weather forecasts or rainfall records to sample the expected distribution of rainfall forcing.

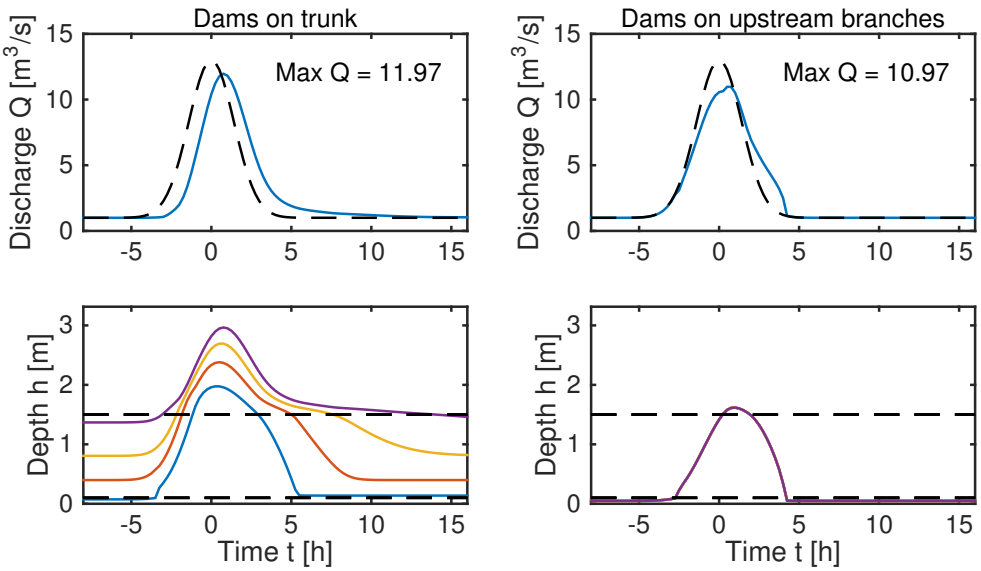

Figure 10: Solutions for a two-dimensional network as in figure 9, forced by uniform inflow to each of the 8 branch nodes $q(t) = q_0 + q_{max}e^{-t^2/\sigma^2}$, with $q_0 = 0.125$ m^3 s^{-1}, $q_{max} = 1.5$ m^3 s^{-1}, and $\sigma = 2$ h. The upper panels show the discharge downstream (node 4), and the lower panels show the water depth at the 4 dams (nodes 1-4 for the case on the left, and the identically-behaving nodes 5,6,9,10 for the case on the right). Parameter values are as given in section 2.3, together with $\lambda = 20$, $b = 0.1$ m and $H = 1.5$ m.

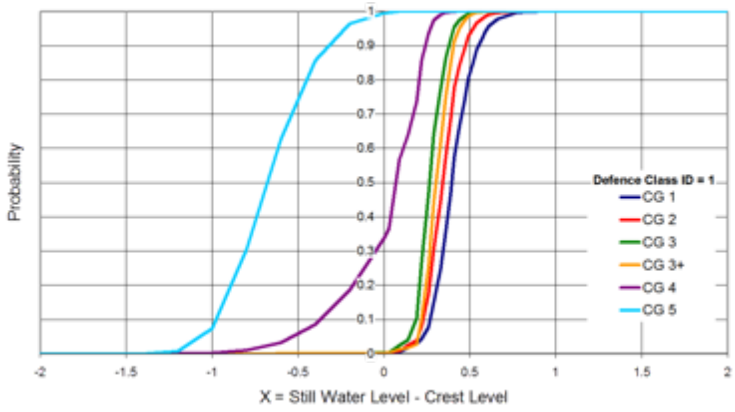

Figure 11: Example fragility curves for dams in different categories, CG1 being the strongest and CG5 the weakest. The curves show the probability of the dam failing above the given water level. Units of water level are meters.

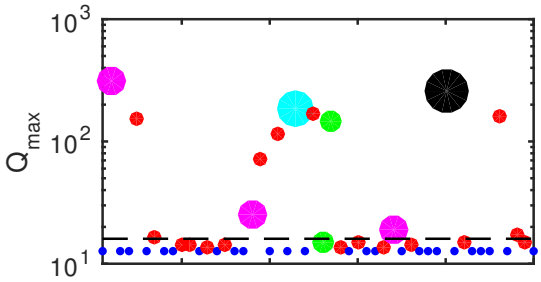

Figure 12: Maximum discharge at the downstream node for an ensemble of runs on the one-dimensional network in figure 6, forced by the same upstream inflow $q(t) = q_0 + q_{max}e^{-t^2/\sigma^2}$, with $q_0 = 1$ m^3 s^{-1}, $q_{max} = 15$ m^3 s^{-1}, and $\sigma = 2$ h. Two examples from this ensemble are shown in figure 13. The size and colour of the dots indicates the number of dams that failed during each realisation (blue=0, red=1, green=2, pink=3, cyan=4, black=5), and the dashed line shows the peak discharge in the case that no dams are installed.

This may be costly in general, and so it may be desirable to establish some rules of thumb about which dam placements are more, or less, at risk from cascade failure.

Knowledge gained from this type of analysis might be used to plan for the size and strength of dams that should be built at different locations. For example, it could be that certain locations are particularly prone to collapse (downstream of merging tributaries for example), and building one stronger 'buffering' dam could significantly reduce the risk of a cascade.

As an example of cascade failure in the network model, we return to the one-dimensional example shown in figure 6. We impose a regular storm inflow to the upstream node of the form given in (17), and examine an ensemble of 50 runs. Each of the dams is assigned a critical water depth h_{ci} such that when $h_i > h_{ci}$ the dam collapses; the critical depth is sampled from a normal distribution with mean 3.5 m and standard deviation 0.5 m (the top of the dam is at 1.5 m so dam collapse usually occurs when the dam is already submerged).

The results are shown in figures 12 and 13. Figure 12 shows the peak discharge Q_{max} at the downstream node for each simulation. This is coloured by the number of dams that collapse; small blue dots indicate that no dams failed, and since the forcing is identical in each case, the peak discharge in this case is always the same. It is lower than what the peak would have been in the absence of any dams, so the dams are proving effective in these cases. Larger dots correspond to more dams having failed. In most cases only one dam collapses, and the peak discharge recorded downstream is strongly dependent on which one fails (the red dots in figure 12). A larger peak occurs if the collapsed dam is further downstream, since if an upstream dam fails (and importantly does *not* precipitate a cascade of downstream failure) the sudden release of water from that dam is buffered by the dams further downstream.

If, on the other hand, a single dam failure leads to further collapse of two or more dams, the peak discharge can be much larger. In one example, all five dams collapse in quick succession, and the time series of this example (the black dot in figure 12) is shown in figure 13, where it is compared to an example with no failure. We have found that the pattern of failure in this one-dimensional model, including the likelihood for cascading failure, depends heavily on the assumed dam sizes, critical water depth distribution, and the magnitude of rainfall events.

As a second more instructive example of cascade failure, we revisit the herring-

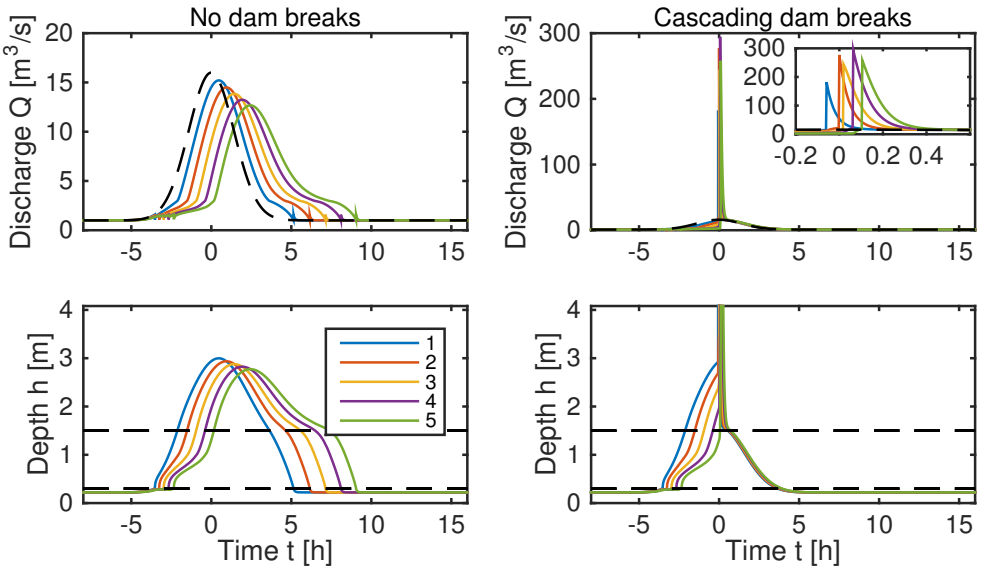

Figure 13: Example from an ensemble of runs on the one-dimensional network in figure 6, forced by the same upstream inflow $q(t) = q_0 + q_{max}e^{-t^2/\sigma^2}$, with $q_0 = 1$ m^3 s^{-1}, $q_{max} = 15$ m^3 s^{-1}, and $\sigma = 2$ h. The dams have critical failure water depths h_{ci} drawn from a normal distribution with mean 3.5 m and standard deviation 0.5 m. In the case on the left, no dams fail, whereas in the case on the right (when the uppermost dam is particularly weak), they all fail in a cascade. Parameter values are as given in section 2.3 except with $\ell = 1000$ m, $S = 0.005$, and $\lambda = 20$, together with $b = 0.3$ m and $H = 1.5$ m.

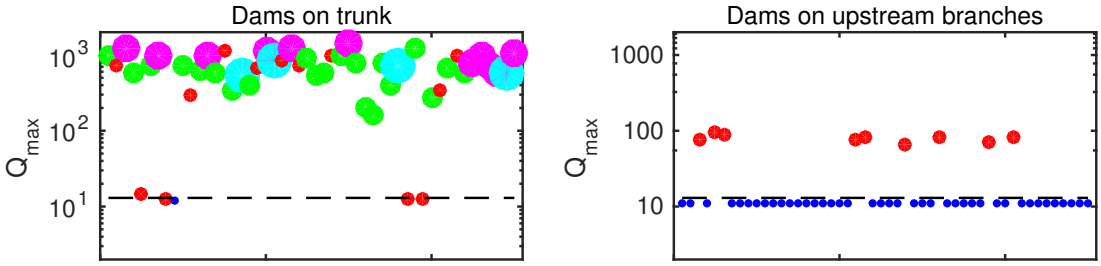

Figure 14: Maximum discharge at the downstream node for an ensemble of runs on the two-dimensional network in figure 9, forced by uniform inflow to each of the 8 branch nodes $q(t) = q_0 + q_{max}e^{-t^2/\sigma^2}$, with $q_0 = 0.125$ m^3 s^{-1}, $q_{max} = 1.5$ m^3 s^{-1}, and $\sigma = 2$ h. Each dam is assigned a failure depth h_{ci} drawn from a normal distribution with mean 2.5 m and standard deviation 0.5 m. In the first case, dams are placed on the four trunk segments (nodes 1-4), while in the second case they are placed on the upstream side branches (nodes 5,6,9,10). The size and colour of the dots indicates the number of dams that failed during each realisation, and the dashed line shows the peak discharge in the case that no dams are installed. Parameter values are as in figure 10.

bone network in figure 9. We consider the two possible placements of 4 dams that were discussed earlier; either along the main trunk (nodes 1-4) or on the upstream side branches (nodes 5,6,9,10). In figure 10 we found that there was relatively little difference in the peak discharge measured at the downstream node with these different placements. However, in figure 14 we see that the first case is much more at risk from cascade failure. This figure shows the peak downstream discharge in an ensemble of simulations, with the failure depths for each dam being different each time. The failure depths are sampled from the same distribution in each case.

In almost every run with dams on the trunk, we see cascade failure occurring so that 3 or 4 of the dams collapse; this leads to extremely high (though short-lived) peak discharge. In contrast, when the dams are placed on the side branches, there is no possibility of cascade failure (the individual branches do not communicate with each other) and it is unlikely that more than one dam collapses. Thus, although each of these dam placements is similarly effective at reducing the peak discharge, there may be strong reason for preferring the second design that places them on the upstream tributaries.

4.3 Erosion

A possible model for scouring beneath the dam is to say that there is an eroded depth d that evolves according to a prescribed erosion rate that depends upon the flow rate beneath the dam. We then write the height of the dam bottom and top above the stream bed as

$$b = b_0 + d, \qquad H = H_0 + d, \tag{18}$$

where b_0 and H_0 are the initial values of the dam heights when the dam is installed. As erosion occurs, the dominant effect is to increase the gap beneath the dam, which allows an increased flux to pass before the dam comes into operation.

Of course another effect of the scouring is to loosen the foundations of the dam which increases the chances of failure. Such dam collapse was considered in the previous section, and the gradual weakening could be accounted for by reducing the critical water depth h_c over time.

One possible law for the bed erosion rate under the dam is to base it on a modification of the well know Meyer-Peter and Müller equation [Gyr and Hoyer, 2006, Fowler, 2011], as

$$\frac{\mathrm{d}d}{\mathrm{d}t} = e_0 \left(1 - \frac{d}{d_{max}}\right) \max\left[\frac{\rho_w f u^2}{(\rho_s - \rho_w) g d_p} - \tau_*, 0\right]^{3/2}, \tag{19}$$

with

$$d = 0 \qquad \text{at} \quad t = 0. \tag{20}$$

Here τ_* is the critical Shields street below which no erosion occurs, ρ_s and ρ_w are the densities of bed material and water respectively, d_p is the diameter of the particles making up the bed, and f is a dimensionless friction factor, while e_0 is an erosion rate scale than depends upon the cohesiveness of the bed material and d_{max} is the maximum depth of erosion. The flow velocity u is the velocity under the dam, which is given in terms of h by

$$u = \sqrt{\frac{2gh^2}{b+h}}. \tag{21}$$

5 Optimisation

5.1 Minimising cost

A number of different optimisation problems were proposed, with the aim of establishing a formal procedure to decide both the placement and the design of the dams (described in the model with the two parameters b and H that represent the bottom and top of the dam). The goal is to reduce the peak water level at some particular location, which we take to be the most downstream node of the network (node N, say). Given that water depth is related to discharge in our model, we can equivalently think about reducing the peak discharge there. In order to provide a well-posed optimisation problem we must also introduce some element of cost associated with each dam.

There are two complementary ways to frame the optimisation problem. In one formulation, we attempt to minimise the peak water depth for a given amount of resources. In this case, the cost is the constraint, and the objective to be minimised is the water depth h_N. Alternatively, we could decide on a maximum allowable water level h_{flood} (or discharge) and attempt to minimise the cost required to ensure that this constraint is achieved. Most discussion at the study group followed the latter approach.

The simplest assumption about the costs of the dam is that they are proportional to the amount of material used, so we take a cost function

$$C = \sum_{i=1}^{N} c_i = \sum_{i=1}^{N} \gamma(H_i - b_i), \tag{22}$$

for some constant γ. An additional fixed cost for each dam location could also be added (to reflect the added logistical costs associated with each site). Thus, we can pose the problem of choosing the control parameters b_i and H_i to minimise C subject to the constraint that $h_N < h_{flood}$, where h_{flood} is the maximum allowable water depth.

The water depth h_N depends on the parameters through the solution of the network model, driven by some prescribed rainfall forcing. Ideally, the optimisation problem should account for the full range of forcing that is expected for any location, and this may involve sampling a large range of possibilities. In practice it is likely that the hard constraint on the water depth is not achievable for every eventuality, and we may instead choose the constraint to be a certain reduction in the *probability* of it being exceeded, from 1 in 10 years to 1 in 100 years for example.

There was not sufficient time at the study group to solve the full optimisation problem. As a proof of concept, we performed the calculation for a one-dimensional network with steady-state rainfall. Clearly this is not of direct use for flood planning, but it demonstrates that a tractable mathematical problem can be posed.

The steady-state optimisation problem can be formulated as follows:

$$\underset{H,b\in\mathbb{R}^N}{\text{minimise}} \quad C = \sum_{i=1}^{N} \gamma(H_i - b_i) \tag{23a}$$

$$\text{subject to} \quad b \geq 0, \tag{23b}$$

$$H > b, \tag{23c}$$

$$h_N \leq h_{flood}, \tag{23d}$$

$$\sum_{j=1}^{N} a_{ji} Q_j - Q_i + q_i = 0, \tag{23e}$$

where the elements Q_i are taken from the flux function (8). In this formulation, the flux balance is incorporated as a set of nonlinear equality constraints, while the condition on the downstream water depth is added as an inequality constraint.

In order to solve this minimisation problem, we use the MATLAB routine `fmincon`, which is appropriate for optimisation problems with nonlinearities. We provide an initial guess for H and b, which satisfies all of the constraints, in this case chosen to be $b_i = 0.5$, $H_i = 1$ for all $i \in N$. Then during each iteration, we check the h_N corresponding to the steady state solution stemming from our choice of H and b. The algorithm then calculates a feasible descent direction and step length, which provide new values for H and b. The routine stops when the difference between the current iterate and the previous one is smaller than some fixed tolerance.

The results of running this algorithm are shown in figure 15. There is not a unique solution to this problem, since in steady-state the discharge is imposed and so the water depth at the downstream node depends only on the dam at that node. To force the solution to include upstream dams we imposed an additional constraint that the thickness of each dam $H_i - b_i$ must exceed a minimum value (by changing (23b) to $H > b + 1$). The optimisation algorithm then correctly identifies that the minimum cost is achieved by choosing all the dams to have this minimum thickness.

An additional cost not included in (22) is that associated with dam failure. The potential impact of this risk is likely to be an important consideration in design, and to account for this in the optimisation procedure we should ascribe a cost to it. For example we might calculate the probability of cascade failure (as in the previous section) and add a cost term proportional to this probability.

5.2 Time-delay model

An alternative way to think about the problem is in terms of the time delay, or flow speeds, over each section of the network. Such an approach might be useful for thinking about the synchronisation and desynchronisation of different parts of the stream network. Little progress was made with this consideration at the study group.

6 Conclusions

We have formulated a network model for a catchment area that allows for simple exploration of the effectiveness of different dam placements and designs, and is sufficiently cheap to solve that it may be useful in analysing risks that require a large ensemble of simulations. The formulation treats each reach of the stream as a lumped

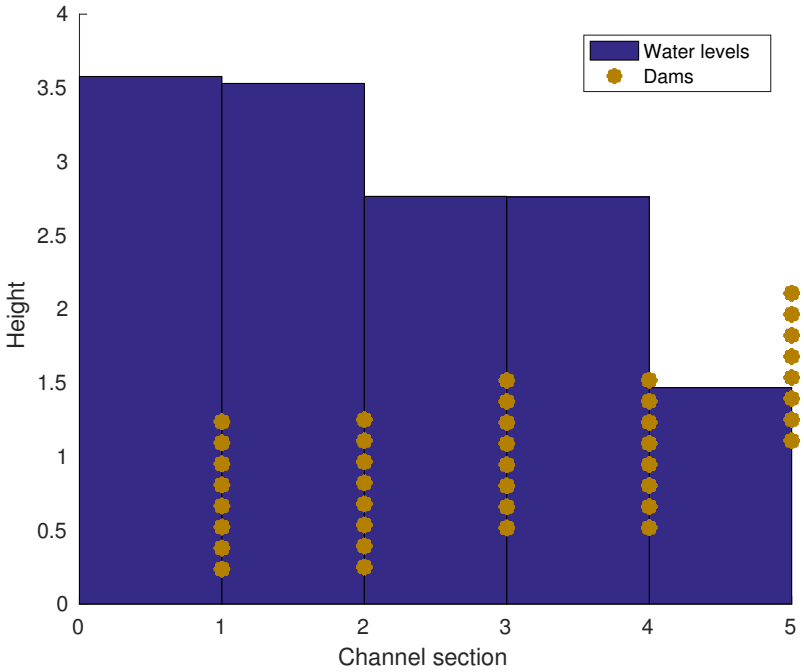

Figure 15: An optimal solution to the cost minimisation problem for steady-state flow with five channel segments. In this case, it is possible to satisfy the constraint $h_N < h_{flood}$ by choosing the final dam to have the minimum allowable thickness $H - b = 1$ provided its bottom b is sufficiently high to allow the prescribed discharge to flow underneath. For the steady-state problem, the upstream dams turn out to be irrelevant, but the optimisation algorithm correctly identifies the minimum-cost solution as requiring all dams to have the minimum thickness. The solution is not unique (the depths b are unconstrained) and different initial guesses converge to different optimal solutions.

element (a node and an edge), with a discharge-area relationship that can be modified to account for the design of a dam on that reach. The model amounts to a system of ordinary differential equations representing mass conservation at each node. Example solutions of the model illustrate its potential for understanding the behaviour of different dam placements.

An immediate conclusion from analysing the scales involved is that a large number of dams are needed to have any significant effect on the peak discharge downstream. The dams should be located in places with the potential to store a large volume of water (flat and wide reaches of the channel). Simple back-of-the-envelope calculations are sufficient to estimate the amount of storage required. The stored water volume must equate to the difference between integrals of the inflow hydrograph and the desired discharge hydrograph.

We have also explored possible methods to account for the failure of a dam, and for the degradation over time due to erosion. These confirm that cascade failure is a risk when dams are placed along a main artery and the risk may be lessened by spreading dams around tributaries. In all cases, the effectiveness of a design must be tempered with the costs, and we have suggested a strategy for setting up a formal optimisation problem that could be used at the planning stage.

There is a lot of scope for further work using this approach. In particular, we can foresee

- Using a two-dimensional surface elevation model to construct real channel networks and to calculate individual catchment basins for each segment that can be used to provide estimates of runoff to each node,
- Using real rainfall records or reanalysis products to force the model with realistic runoff estimates (in particular, with appropriate statistics for storm events),
- Investigating the use of different sizes of dams so that individual dams are not under- or over-used.
- Using a large ensemble of simulations to quantify the risk of cascade failure for any given design.
- Further developing the optimisation problem, to account for the expected spectrum of runoff forcing and for realistic costs.

Acknowledgements

We thank Onno Bokhove for his preliminary work on formulating this challenge together with the JBA Trust so that it could be brought to the study group.

References

A.C. Fowler. *Mathematical Geoscience*. Springer, London, 2011.

A. Gyr and K Hoyer. *Sediment transport. A geophysical phenomenon*, volume 82. Springer, Netherlands, 2006.

P. Metcalfe, K. Beven, B. Hankin, and R. Lamb. A modelling framework for evaluation of the hydrological impacts of nature-based approaches to flood risk management, with application to in-channel interventions across a 29-km^2 scale catchment in the United Kingdom. *Hydrological Processes*, pages 1–15, 2017. doi: 10.1002/hyp.11140.

B.R. Munson, T.H Okiishi, W.W. Huebsch, and A.P. Rothmayer. *Fluid Mechanics*. Wiley and sons, Singapore, 7th edition, si version edition, 2013.