Improving the performance of fast flood inundation models by incorporating results from very high resolution simulations

S. Néelz[1], J. Hall[2], G.Pender[1]

1. School of the Built Environment, Heriot-Watt University, Edinburgh, UK.
2. School of Civil Engineering and Geosciences, Newcastle University, Newcastle-upon-Tyne, UK.

Abstract

In order to accurately predict flood flows over complex topographies and in urban areas it is usually necessary to resolve hydraulic simulation models on very fine grids that can, for example, resolve the effects of individual buildings. The associated computational cost is high, which makes the approach infeasible for flood risk analysis in which many thousands of flood inundation simulations may be required. However, much coarser resolution models can predict the general tendencies of the flow, so are in some respects informative. The study described in this paper has sought to improve the predictions of coarse resolution flood models by making use of information from a few very high resolution model runs. The approach is based upon interpolating the residuals between the coarse and fine model runs and using them to improve prediction. It is shown to improve the accuracy of predictions at the Thamesmead site in the Thames Estuary at relatively modest additional computational cost.

1. Introduction

There is increasing interest in flood prediction in urban areas, which can be concentrations of flood risk. Flood risk analysis involves using hydraulic models to predict variables that are relevant to flood damage, such as flood depth or velocity. To calculate the expected damage involves integrating over a wide range of flow conditions. In urban areas protected by dikes or other engineering systems, analysis involves calculating the expected flood damage for a very wide range of dike system states. The computational cost of risk computation can therefore be enormous, unless the run times of hydraulic models can be kept to a minimum. One approach to resolving this problem is to use rapid flood inundation models that can resolve the main characteristics of the flow (Krupka et al. 2007). This paper seeks, at some additional computational cost, to incorporate some results from a few fine resolution model runs in order to improve the accuracy of prediction.

The fusion of models from results from fine and coarse resolution models is an approach that has received some attention in the statistical literature in the context of statistical emulators of computer codes. Oakley and O'Hagan (2000) use a hierarchical emulator approach in which code runs from models at different levels of resolution are fused in the framework of Gaussian processes. However, Gaussian process emulators are not well suited to highly nonlinear dynamical process models, which is the situation observed in, for example, flood inundation modelling in which the transition from dry to wet gives a marked step. Here therefore we explore approaches that are intended to deal with the non-linearity in water depth associated with the arrival of the wetted front and also make use of a very small number (typically two) of high resolution model runs in order to improve the performance of coarse resolution models.

Fine resolution models are attractive not only in their predictive accuracy (due to better resolution of topography and improved numerical accuracy) but also due to the higher spatial resolution at which predictions are output. This is illustrated in Figures 1 and 2, which show velocity fields from flood inundation models in an urban area running on 2m and 50m grids. Of particular significance is the concentration of flows in gaps between buildings that are resolved in the 2m model, which mean that considerably higher velocities can be predicted at a given point in this model.

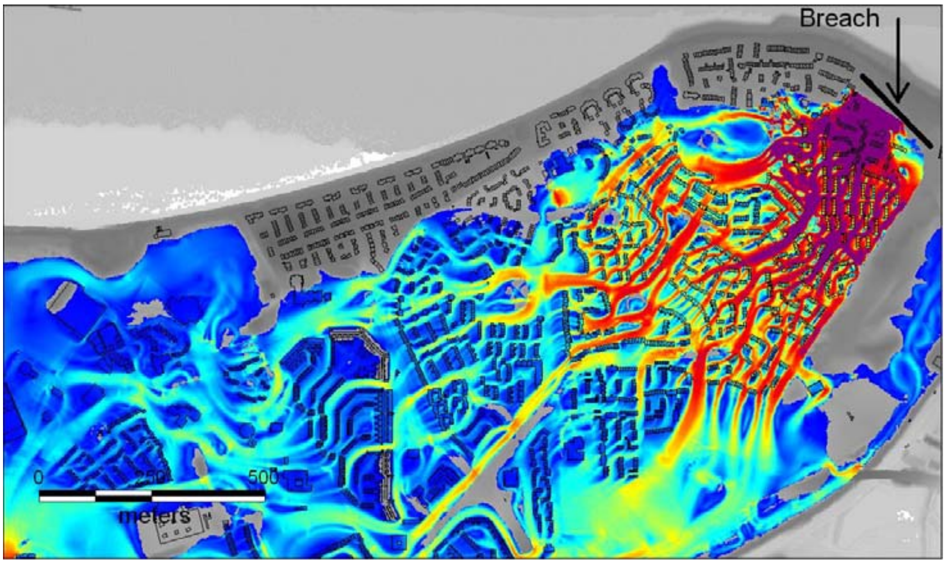

FIGURE 1. Peak velocity field for the Thamesmead site from a 2m grid model

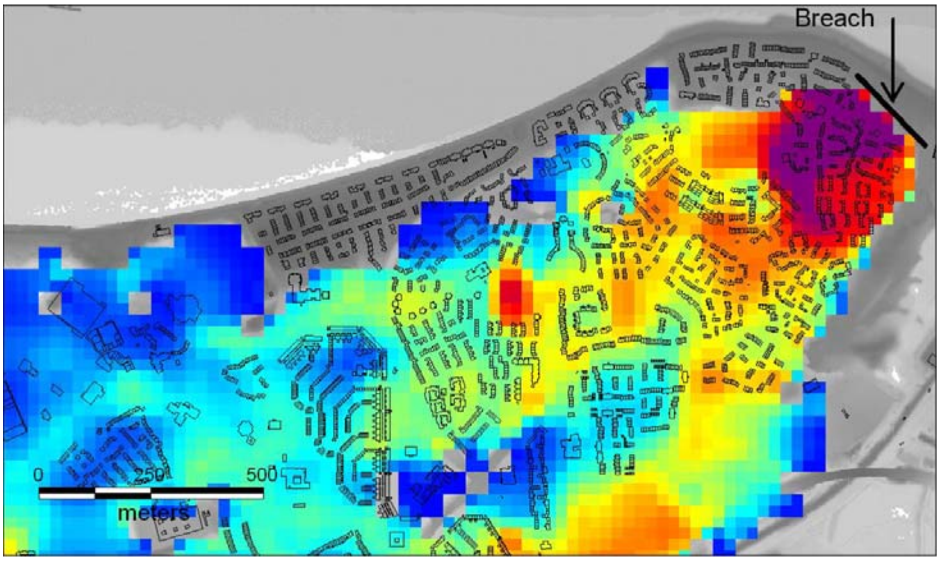

FIGURE 2. Peak velocity field for the Thamesmead site from a 50m grid model

2. Experimental set-up

2.1 Study case

The study site is a ~5 km^2 area of lowlands along the banks of the Thames estuary in England. It is generally very flat, with most areas lying at only 1m above Ordnance Datum Newlyn (the topographic elevation datum for Great Britain). This floodplain is below sea level at high tide, so that the possibility of tidal flooding is physically realistic, although this would imply the formation of a breach through one of the very high flood defences (walls and embankments) that currently defend the floodplain against such disasters. The floodplain is mostly urbanised, with new developments being built at a high rate. These are a combination of small apartment blocks and, predominantly, detached bungalows.

1m (resolution) LiDAR (Airborne laser scanning) Digital Terrain Model (filtered to remove vegetation and building features) supplied by the Environment Agency of England and Wales provided the underlying topography used in the study. This DTM was further processed by the authors to remove any “false blockages” such as bridges or elevated buildings, and to fill in any areas of missing data (mainly caused by the presence of artificial ponds). In addition, exact building locations and geometries were extracted from Ordnance Survey.

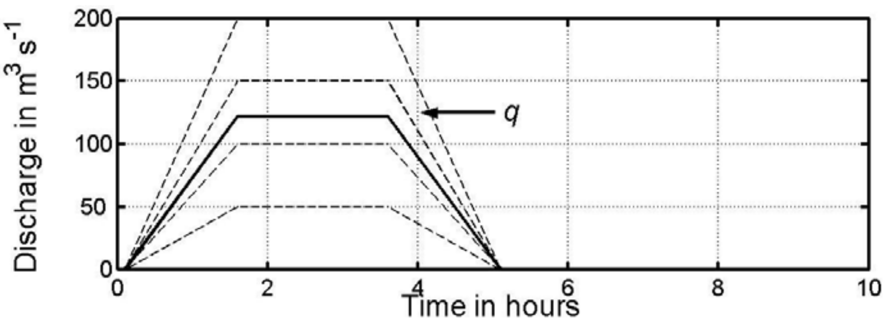

FIGURE 3. Shape of the inflow hydrograph used in the study.

The flood event modelled is represented by a single inflow/time upstream boundary condition representing a breach in the flood defences at the location indicated in Figure 1. The shape of the hydrograph (Figure 3) is similar to what could be expected for tidally driven inundation through a breached embankment. Inflow lasts for a few hours and ends as the tide recedes. This same shape was used in all simulations, but the magnitude of the inflow varied and this was parameterised using the value taken by the inflow at its peak, referred to as q in the rest of this paper.

2.2 Models

The study case described in Section 2.1 was modelled using TUFLOW, a computational engine that provides dynamically linked 2D and 1D solutions of the free-surface flow equations to simulate river flow and flood propagation. TUFLOW solves the full 2D shallow water equations on a regular square mesh. Its solution is based on the Stelling finite difference, alternating direction implicit (ADI) scheme (Stelling 1984), adapted for upstream controlled flow regimes by Syme (1991). Only the 2D component of TUFLOW was used in this study.

The resolution of the fine grid model (FGM) was 2m. This can be considered a very fine grid by current industry standards for a floodplain of this size. Given the flatness of the study site, it is also fine enough to resolve all topographic features of the floodplain, including buildings, which exert a significant influence on the flow. Buildings were not reinstated into the DTM. Instead, the blockage effect that they exert on the flow was taken into account through the use of a very high Manning’s n (roughness) value of 0.5 in computational cells found to coincide with buildings. This ensured that the flow velocities were reduced to near-zero values inside buildings, in such a manner that building blockage effects were effectively simulated. A value of Manning’s n of 0.035 was used uniformly in the rest of the domain. Computational times were between 35 and 89 hours for the 2m FGM depending on the inflow (with a 1s time step), and between 2.5 and 5.5 minutes for the 50m CGM (also with a 1s time step).

A resolution of 50m was used for the coarse grid model (CGM). This is larger than the typical sizes of buildings and of gaps between buildings and therefore did not allow an accurate representation of flow pathways through this urban area. However, the overall effects of buildings on the timing of inundation were represented by using higher values of Manning’s n, set in each computational cell depending on the percentage of surface area that was covered by buildings (Néelz & Pender, 2007). This was parameterised to minimize errors in the flood arrival times throughout the floodplain, but did not prevent the inherent “blurring” of flow pathways caused by the use of a coarse resolution, as shown in Figure 1.

Output data were in the form of time series of either water level or velocity at locations every 50m, i.e. coinciding with the cell centres in the 50m model. These also coincided with cell centres in the 2m model, so that no spatial interpolation was initially required in order to compare and process output data. The techniques were tested at a total of over 100 locations regularly spread across the floodplain. Results were assessed visually and a representative selection is reported on in this paper.

2.3 Combining FGM and CGM results

2.4 Overview of methodology

The method for improving CGM predictions is applicable to predictions of both water surface elevations and flow velocities. Water surface elevations are to be preferred to water depths, as these generally vary more smoothly even over complex flood topographies. For the given variable of interest, the residuals between the fine and coarse models were extracted at each time step.

As stated in Section 1, only 1 or 2 preliminary or “training” runs of the FGM may be available for a given breach location, because an uncertainty or risk analysis would typically involve many possible combinations of boundary conditions, thus severely restricting the amount of computational resources that can be allocated to running high resolution simulations. The study is therefore carried out with the assumption that no more than two FGM ‘training’ runs are available.

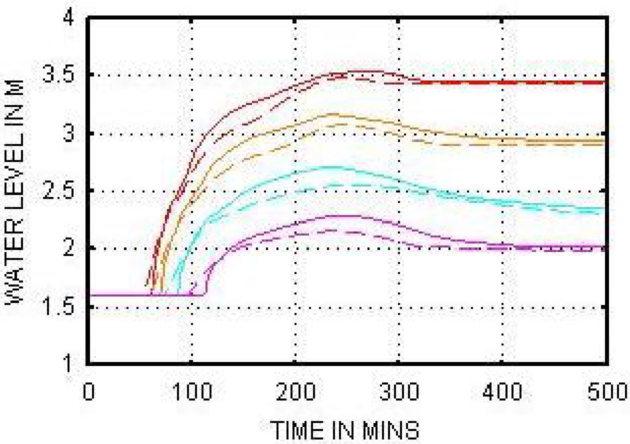

FIGURE 4. Time series output for water surface elevation from CGM (solid lines) and FGM (dashed lines) for a range of inflow discharges

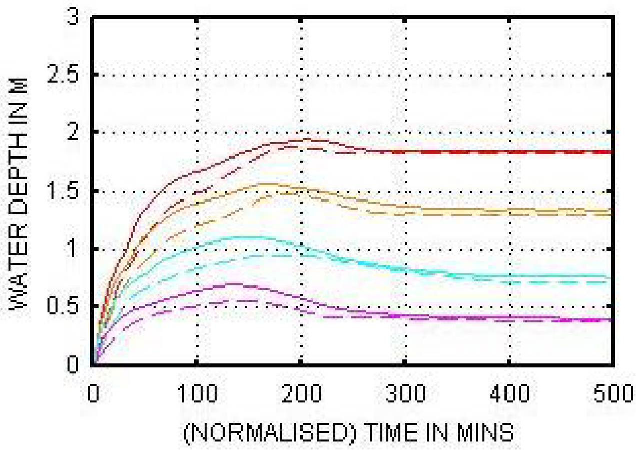

FIGURE 5. Time series output for water surface elevation from CGM (solid lines) and FGM (dashed lines), normalised to the arrival time of the flood front

Figure 4 shows time series outputs for water level, extracted from the FGM and the CGM at a selected point in the floodplain. In these time series, as with almost all output points, there is one single severe discontinuity at the time of arrival of the flood wave, making any direct extrapolation or interpolation of FGM/CGM residuals very inaccurate. The time series were therefore normalised to the arrival time of the flood wave, and the arrival times were treated as separate parameters (see Figure 5). There was therefore a need for appropriate approaches to predict 1) the arrival time, or 2) the time series of "normalised" depths or velocities. Such approaches can be summarised by the following equation:

$$\eta(q) = \eta_{CGM}(q) + \varepsilon(q) \tag{1}$$

where η is the quantity to be predicted (i.e. arrival time, or water depth or velocity normalised to the arrival time) at a given location, η_{CGM} is the value taken by η in the CGM run, q is the input flow discharge for which the prediction is sought and ε is the FGM-CGM residual (see for example Figure 6), estimated on the basis of a number of training runs at inflow discharges q_i, where it took values $\varepsilon(q_i)$, defined as follows:

$$\varepsilon(q_i) = \eta_{FGM}(q_i) - \eta_{CGM}(q_i) \tag{2}$$

With only 2 such training runs available, this estimation is done either by averaging:

$$\varepsilon = \frac{\varepsilon(q_1) + \varepsilon(q_2)}{2} \tag{3}$$

or by linear interpolation/extrapolation:

$$\varepsilon(q_i) = \varepsilon(q_1) + (\varepsilon(q_2) - \varepsilon(q_1))\frac{q_i - q_1}{q_2 - q_1} : \quad q_2 > q_1 \tag{4}$$

With only one training run available, $\varepsilon = \varepsilon(q_1)$ is the only option.

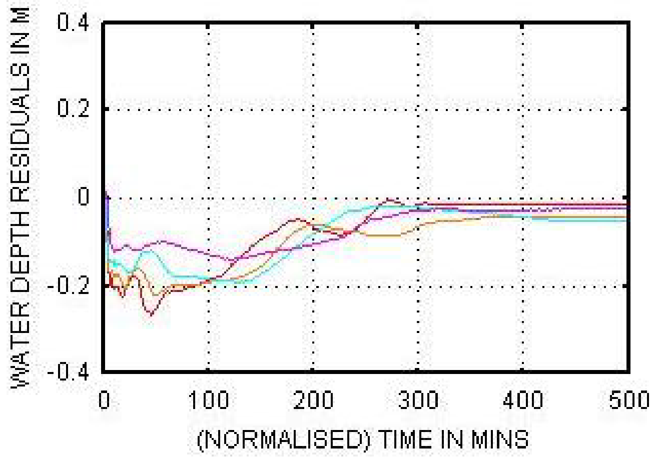

FIGURE 6. Time series output of water surface elevation residuals for a range of inflow discharges

An alternative to Equation 4 is also considered, where there are two training runs q_1 and q_2, in which interpolation/extrapolation is undertaken on the basis of the of the predictive variable $\eta(q)$ rather than using q as the independent variable:

$$\varepsilon(q_i) = \varepsilon(q_1) + \left(\varepsilon(q_2) - \varepsilon(q_1)\right)\frac{\eta_{CGM}(q_i) - \eta_{CGM}(q_1)}{\eta_{CGM}(q_2) - \eta_{CGM}(q_1)} : \quad q_2 > q_1 \qquad (5)$$

This approach relies on the residuals in the predicted variables (arrival time, depth, velocity) vary linearly with thes predictive variables from the CGM, $\eta_{CGM}(q)$, rather than with q.

depend on the inflow value q are independent of the grid resolution, instead of the assumption that FGM-CGM residuals vary linearly with q (as with the equation 2.4). It is tested in this study as a means to predict arrival times only.

It should also be mentioned that a much simpler alternative to using CGM results was also tested. This consisted in estimating predictions of η using spline interpolation, from a number of training runs of the FGM only. While this proved to return similar results in many test cases where up to 5 training runs were available, it is however clearly inappropriate in the present situation where only 2 training runs are available.

3. Results

In this study, the 2m FGM model was run with values of q equal to 10, 50, 100, 150, 200, 350 and 500 m^3/s, to provide runs for validation as well as 'training'. The range of values of q for which the techniques could be tested with any relevance proved to be limited by the dimensions and physical characteristics of the floodplain and the way predicted inundation patterns (and discrepancies between the CGM and the FGM results) depended on q and on the grid resolution.

In the range $50< q <350$ m^3/s, a significant proportion of the floodplain became inundated and flood depths exhibited an intermediate peak, which was significantly (up to 0.5 m) higher than final depths in many locations. There were also significant (up to 30mins) differences in terms of arrival times between FGM and CGM.

For $q < 50$ m^3/s, inundation remained confined to low lying pathways of small transverse dimensions. Also, grid resolution became insufficient in the 50m model to ensure a representation of the topography that was sufficiently accurate for shallow flows. Therefore (FGM-CGM) residuals were difficult to predict with any accuracy. The most obvious example is where the ground remained dry for low values of q (in both the FGM and the CGM), in which case extrapolating residuals always led to errors. There were also a number of locations that were predicted to be flooded in the FGM but not in the CGM, or the opposite. Most predictions for $q = 10$ m^3/s or using $q = 10$ m^3/s as one of the training runs were therefore of very low quality, and are not reported on any further in this paper.

For large values of q and particularly for $q > 350$ m^3/s, inundation spread faster and arrival time residuals became smaller (although it can be argued that there is a greater requirement of accuracy in cases of rapid spreading). Also, the intermediate peak was much reduced and the flood reached its final depth very quickly, with smaller differences in peak water levels between the FGM and the CGM.

3.1 Predicting arrival times

The methods consisting in extrapolating (Equation 4) or averaging (Equation 3) arrival time residuals performed badly, with errors as large as 30mins. On the other hand, the method consisting in predicting arrival time residuals by linear interpolation/extrapolation on the predictive variable (Equation 5) proved more or much more accurate in the majority of locations. For this reason, only results using this latter method to predict arrival times will be presented in the remainder of this paper (Figure 7). These conclusions were reached in tests using either i) q = 100 m^3/s and q = 200 m^3/s or ii) q = 100 m^3/s and q = 500 m^3/s as training cases.

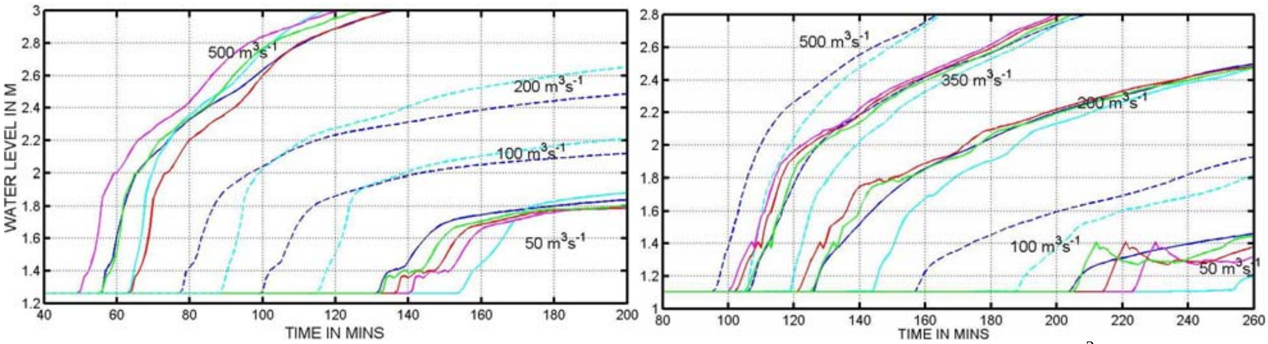

FIGURE 7. Close-up showing the beginning of flooding at two distinct locations, with a) q = 100 m^3/s and q = 200 m^3/s as training runs (left), and b) q = 100 m^3/s and q = 500 m^3/s as training runs (right). The methods used for arrival time prediction were: linear extrapolation/interpolation of residuals (Equation 4) (red), averaging of residuals (Equation 3) (purple), and the method based on the normalised arrival time (Equation 5) (green). (for clarity these methods to predict the arrival times are all combined with the same method to predict the water level time series, namely residual averaging). The predictions can be compared with FGM validation predictions (dark blue)

3.2 Predicting water depth and velocity

It was first attempted to predict water level time series using the training cases q = 100 m^3/s and q = 200 m^3/s, for inflow values ranging from q = 50 m^3/s to q = 500 m^3/s, *ie.* well outwith the range of training inflow values. The results (illustrated by 2 examples in Figure 8) suggested that linear extrapolation of residuals was often very inaccurate. This is attributed to the amplification of any local or time-dependent noise in the residuals, which was particularly apparent for the larger inflow values. Averaging of residuals was often significantly better but was also insufficiently accurate where residuals showed a high level of sensitivity to q.

In addition, the approach was applied to the prediction of velocity (also using the same training cases q = 100 m^3/s and q = 200 m^3/s). This was usually inaccurate, because of 1) the high level of noise (due to oscillations, either physical or numerical) in the velocity time series which made linear extrapolation of residuals very chaotic and 2) the high level of sensitivity to q which made residual averaging very inaccurate, particularly for predictions at the larger values of q.

Predictions based on only one training run were understandably worse than the predictions based on two training runs and are not reported on here.

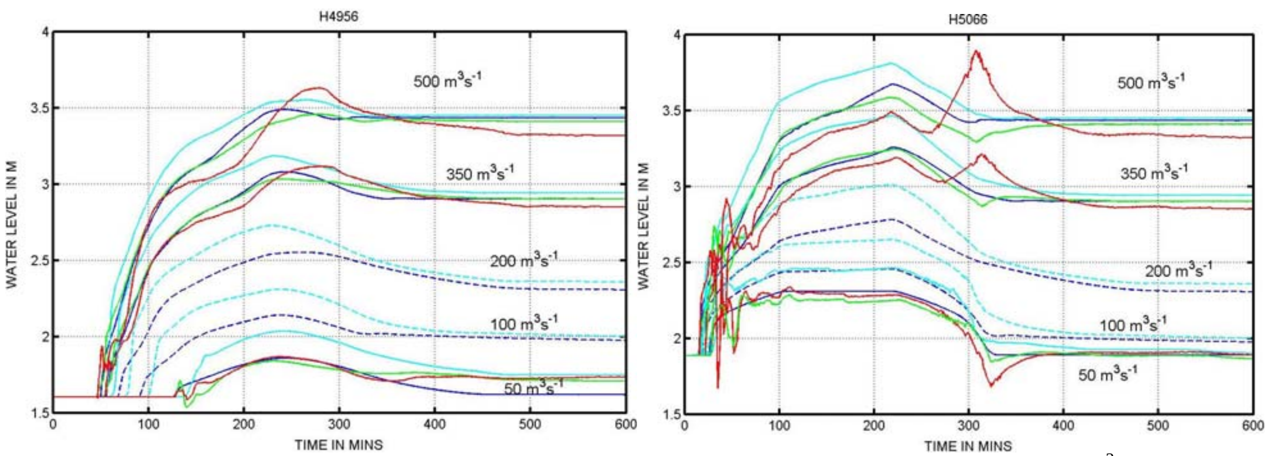

FIGURE 8. Prediction of water level at 2 locations on the floodplain. Training runs based on q = 100 m^3/s and q = 200 m^3/s. The methods used for water level time series prediction were: linear extrapolation of residuals (Equation 3) (red), and averaging of residuals (Equation 4) (green). Turquoise: CGM results. Dark blue: FGM results. Dashed lines: training runs.

The option consisting in running training cases with the inflow values q = 100 m^3/s and q = 500 m^3/s was also investigated. Inflow values much larger than 500 m^3/s are of limited relevance because they would result in a case of very quick floodplain inundation. Values much smaller than 100 m^3/s are likely to result in significantly different (limited in extent) inundation patterns and therefore should not be one of only two values used in the training runs.

The results, illustrated in Figure 9 (predictions for q = 50 m^3/s, q = 200 m^3/s and q = 350 m^3/s) were generally much more promising than when using the narrow training range, with many locations where variables such as the peak water levels or velocities were predicted very accurately. The most successful predictions were in general achieved using linear interpolation of residuals, with the averaging method being more accurate in only a small proportion of locations. These encouraging results were also achieved in the prediction of velocity, for which residuals can be very large (this is the case if an output point coincide with one of the many major flow routes, see Figure 2), and the use of the CGM alone is wholly inappropriate. However, prediction of velocity was often made more difficult by the presence of oscillations (either physical or numerical) in the results from the training runs or from the prediction CGM runs.

A number of limitations of the approach where apparent from the results, as many results suffered from a) over- or under-correction for at least a part of the time domain, or b) arrival time prediction errors. The most important limitation occurred at locations where flooding was shallow or inexistent for the training run q = 100 m^3/s. It is clear in Figure 9 (right) that choosing q = 50 m^3/s as one of the training inflow values would have resulted in very inaccurate predictions. Also, residuals during the first minutes of the inundation process were often excessively dependent on q, as differences in the way topography is represented in the FGM and the CGM often affect shallow flows only (and therefore have consequences for low values of q that last longer than for high values of q). This often caused noisy predictions during the rising limb of the flood, see Figure 10 (left), and also led to substantial errors in the arrival time predictions. In addition, the approach did not take into account how temporal characteristics of flood fronts varied depending on q. An example is shown in Figure 9 (right), where predictions inaccuracies for q = 200 m^3/s and q = 350 m^3/s between the times 150 and 300mins may be due to this. Predictions for q =50 m^3/s were very inaccurate at all locations where inundation was very shallow, see Figure 10 (left).

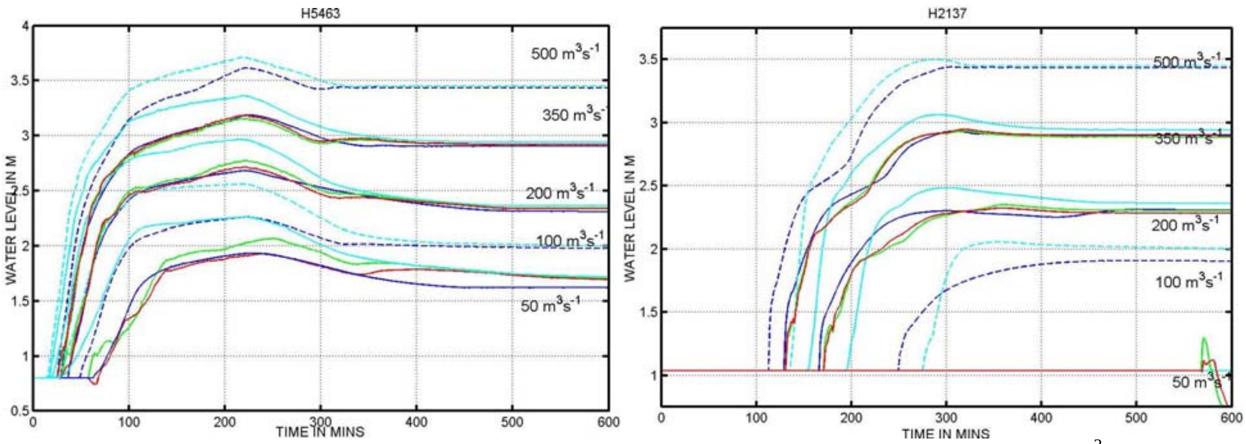

FIGURE 9. Prediction of water level at 2 locations on the floodplain. Training runs based on q = 100 m^3/s and q = 500 m^3/s. Methods used: linear interpolation/extrapolation of residuals (red); averaging of residuals (green). Light blue: CGM results. Dark blue: FGM results. Dashed lines: training runs

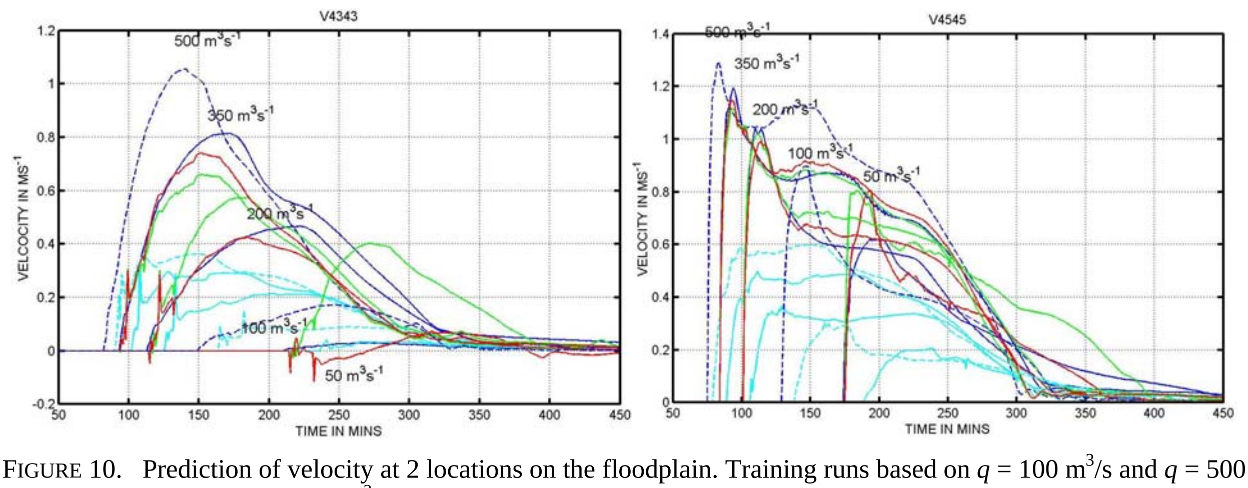

FIGURE 10. Prediction of velocity at 2 locations on the floodplain. Training runs based on q = 100 m^3/s and q = 500 m^3/s. See previous figure caption for colour codes.

3.3 Spatial interpolation

In addition to improving predictions of water levels and velocities at the computational grid points of the coarse grid model, the approach tested above can also be used to improve the spatial resolution of CGM model results. Predictions at any location on the floodplain can be obtained by spatial interpolation from grid node predictions, and can subsequently be improved using an estimate of residuals with the FGM. The inherent "blurring" effect obtained in CGM predictions is illustrated by Figure 8, which clearly shows that major flow routes through the urban environment are not adequately modeled using the 50m model and that very significant benefits may be gained from using the results from preliminary runs of the 2m models.

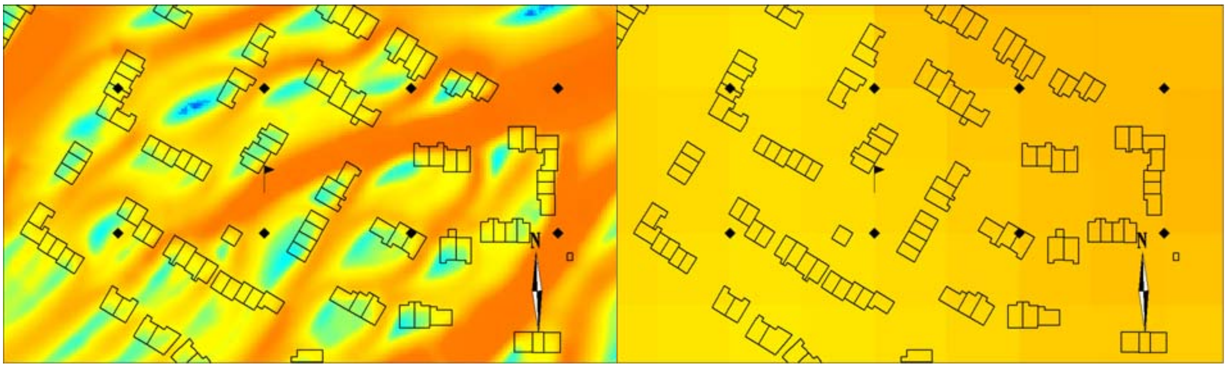

FIGURE 11. Peak velocity field predicted by a) the 2m FGM and b) the 50m CGM, for q = 200 m^3/s. the black dots are the computation nodes of the 50m model. The flag indicates the location where the approach is tested.

This was tested using three-dimensional cubic spline interpolation applied to the results from the 50m model at every output time step to obtain time series of velocity at a location coinciding with a computational node of the 2m model (for which FGM predictions were directly available) (see flag in Figure 11), which also was in the middle of a major flow routes induced by the presence of buildings. The same training and prediction inflow values as in Section 3.2 were used. Figure 12 shows that the CGM predictions for q = 200 m^3/s and q = 350 m^3/s were accurately corrected, while the prediction for the lower inflow value (50 m^3/s) was slightly overcorrected. These results are consistent with the results reported in Section 3.2.

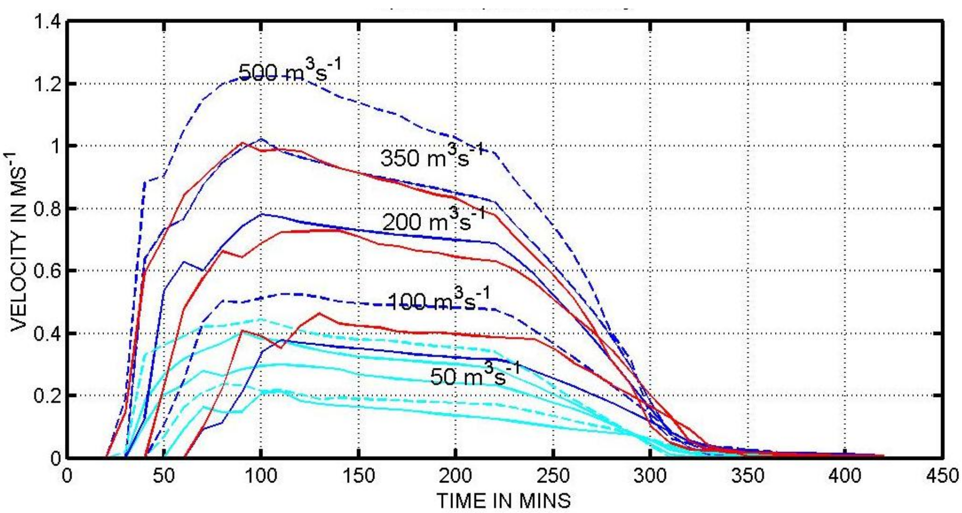

FIGURE 12. Prediction of velocity (red). Training runs based on q = 100 m^3/s and q = 500 m^3/s. Light blue: interpolated from CGM time series by spatial spline interpolation. Dark blue: output of FGM at exact location.

4. Conclusion

The study has demonstrated the potential for using a few fine grid model (FGM) runs to improve the accuracy of prediction of flood inundation models running on a coarser grid. The method has to be applied with some care. It is most effective where the (two) FGM runs span the range of flow conditions well. However, this is not straightforward as different areas in the floodplain will have different thresholds of flooding, so not all of the domain will be well characterized by the same FGM run. At low flood depths the method is not successful.

In flood risk analysis of systems of dikes it is usually necessary to test a very large number of inundation scenarios, which means that for any given scenario it will only be feasible to do a very small number FGM runs. Under these circumstances it is only possible to employ rather simple functions for fitting to the residuals between the FGM and CGM at any point, the most complex being a linear regression. The residuals between neighbouring CGM grid points are more or less uncorrelated, so no further benefit is obtainable by including neighbouring points in the regression analysis.

Correction of arrival time of the flood wave using the residuals between the FGM and CGM considerably increases the accuracy of the prediction method. Three methods of correcting the arrival time from the CGM were tested. The most effective proved to be a method based on scaling the residual in relation to the predicted arrival time.

Predictions of flood depth can be improved (following normalization of arrival time) by incorporation of linearly scaled or average residuals between the FGM and CGM. However, it is noteworthy that maximum flood depths (the most commonplace indicator used in flood risk analysis) is predicted reasonably accurately with the CGM, so the benefit of incorporating information from the FGM is marginal. Indeed for many applications it may be feasible to use an even faster inundation model than the CGM.

Prediction of velocity is problematic, given the oscillating nature of the velocity near the front of the flood wave. However, without incorporation of insights from the FGM, predictions of velocity can be poor. The improved spatial resolution of the FGM is a particular attraction, and it has been shown that insights from the FGM can be used to ‘infill’ spatial detail in the CGM.

Acknowledgements

The research reported in this paper was conducted as part of the Flood Risk Management Research Consortium. The FRMRC is supported by grant GR/S76304 from the Engineering and Physical Sciences Research Council, in partnership with the Natural Environment Research Council, the DEFRA/EA Joint Research Programme on Flood and Coastal Defence, UKWIR, the Scottish Executive and the Rivers Agency (Northern Ireland). This financial support is gratefully acknowledged. The authors are also grateful to the Environment Agency of England and Wales for providing LiDAR data, to WBM Australia for the use of their TUFLOW software, and Ordnance Survey for providing Mastermap® data.

REFERENCES

KENNEDY, M.C. and O'HAGAN, A. 2000.
Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, **87**(1), 1-13

KRUPKA, M., PENDER, G., WALLIS, S., SAYERS, P.B. and MULET-MARTI, J. 2007
A rapid flood inundation model, *32nd IAHR Congress, Venice, Italy, 1-6 July*.

NEELZ, S. and PENDER, G. 2007.
Parameterisation of square-grid hydrodynamic models of inundation in the urban area, *32nd IAHR Congress, Venice, Italy, 1-6 July*.

OAKLEY, J.E. and O’HAGAN, A. 2004.
Probabilistic Sensitivity Analysis of Complex Models: a Bayesian Approach. *J. of the Royal Statiscal Society: Series B (Statistical Methodology)*, **66**-3, 751-769.

STELLING G.S. 1984.
On the Construction of Computational Methods for Shallow Water Flow Problems. *Rijkswaterstaat Communications*, No 35/1984. The Hague, Netherlands.

SYME, W.J. 1991.
Dynamically Linked Two-Dimensional / One-Dimensional Hydrodynamic Modelling Program for Rivers, Estuaries & Coastal Waters. *M.Eng.Sc. (Research) Thesis*, Dept of Civil Engineering, The University of Queensland, Australia.