

# Optimisation of hardware setups for time-efficient HEC-RAS simulations

Ramtin Sabeti<sup>1</sup> , Thomas Rodding Kjeldsen<sup>1</sup>, Ioanna Stamatakis<sup>2</sup> and Babak Rahi<sup>3</sup>

<sup>1</sup>Department of Architecture and Civil Engineering, University of Bath, Bath, UK; <sup>2</sup>School of Engineering, University of Greenwich, Chatham Maritime, UK and <sup>3</sup>School of Computer Science, University of Nottingham, Nottingham, UK

## Research Article

**Cite this article:** Sabeti R, Kjeldsen TR, Stamatakis I and Rahi B (2024). Optimisation of hardware setups for time-efficient HEC-RAS simulations. *Cambridge Prisms: Water*, 2, e13, 1–10

<https://doi.org/10.1017/wat.2024.11>

Received: 29 May 2024

Revised: 18 September 2024

Accepted: 04 November 2024

### Keywords:

HEC-RAS 2D; simulation time; hardware optimisation; Google Cloud Platform

### Corresponding author:

Ramtin Sabeti;

Email: [rs3195@bath.ac.uk](mailto:rs3195@bath.ac.uk)

### Abstract

The necessity for high-resolution two-dimensional (2D) simulations in flood modelling often requires excessively long simulation times. This study evaluates the impact of various hardware configurations on Hydrologic Engineering Center-River Analysis System (HEC-RAS) 2D with particular emphasis on Central Processing Unit (CPU) speed, number of cores, Random Access Memory (RAM) capacity, addressing a critical gap in the optimisation of hardware setups for time-efficient simulations. These findings are invaluable for flood modellers and the HEC-RAS community, ultimately supporting more effective flood risk management and decision-making. Additionally, the study examines how different meshes, numerical solution methods, and solving equations perform within these hardware setups, aiming to examine the effects of computational techniques on overall simulation efficiency. Our investigations were carried out using both virtual machines on the Google Cloud Platform and a desktop PC. The findings indicate that optimal performance in HEC-RAS 2D simulations does not necessarily correlate with a higher number of cores or increased RAM. Instead, a well-adjusted configuration with 8 cores and 64 GB of RAM demonstrates superior efficiency. This result questions the usual assumptions about the need for greater computational power and emphasises the value of carefully optimising hardware for fast hydraulic modelling.

### Impact statement

A key challenge with 2D hydraulic models that analyse water flow in both length and width is their high computational demand, leading to long processing times. The relationship between hardware and model is crucial in reducing this time, and a more powerful setup does not always guarantee better results. Here, we focus on HEC-RAS 2D, a widely used hydraulic model in academia and industry, and provide insights into how optimising hardware is able to shorten the processing time of simulations. The impact of the results presented here is reducing the simulation time that facilitates running multiple simulations, essential for model validation, and enables the modellers to perform the uncertainty analysis, such as Monte Carlo requiring 1,000s of simulations.

### Introduction

In recent years, numerical modelling in the field of hydraulic engineering has rapidly improved, providing an invaluable platform for conducting risk assessments, forecasting, and supporting the design of hydraulic structures. Many models are available, however, HEC-RAS 2D from HEC models developed by the US Army Corps of Engineers Hydrologic Engineering Centre (Loucks, 2023) stands out for its widespread recognition and successful application across industry and academia supported by benchmarking of Brunner's (2018) study. However, the two-dimensional (2D) nature often raises concerns regarding simulation times.

Néelz and Pender (2010) conducted a comprehensive benchmarking study, which tested a range of 2D hydraulic modelling packages in the UK. The goal of the study was to establish the capability of these software packages to accurately model natural phenomena for flood risk assessment. Benchmarking cases were provided to developers of UK flood inundation software, and the report provided valuable insights into the accuracy and performance of different models. For all eight test cases, the study compared the performance of 14 modelling packages, documenting details, such as the numerical schemes used, hardware configurations, multi-processor capabilities, mesh size, time-stepping, and run times, which varied significantly across the models. However, while this study offered a general sense of run times, it did not provide specific recommendations on optimising hardware configurations for more efficient simulations. The necessity of optimising hardware configurations has been explored in other exiting studies due to the high computational times associated with 2D hydraulic modelling, such as Neal et al. (2009, 2010), Sanders et al. (2010, 2019), Lacasta et al. (2014), Dazzi et al. (2021), Morales-Hernández

© The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

 Cambridge Prisms

 CAMBRIDGE UNIVERSITY PRESS

(2021), and Buttinger-Kreuzhuber *et al.* (2022). However, specific evaluation along with recommendations on optimising hardware configurations in terms of speed of cores, number of cores, Boot disk, and random access memory (RAM) are still lacking for a considerable number of hydraulic models.

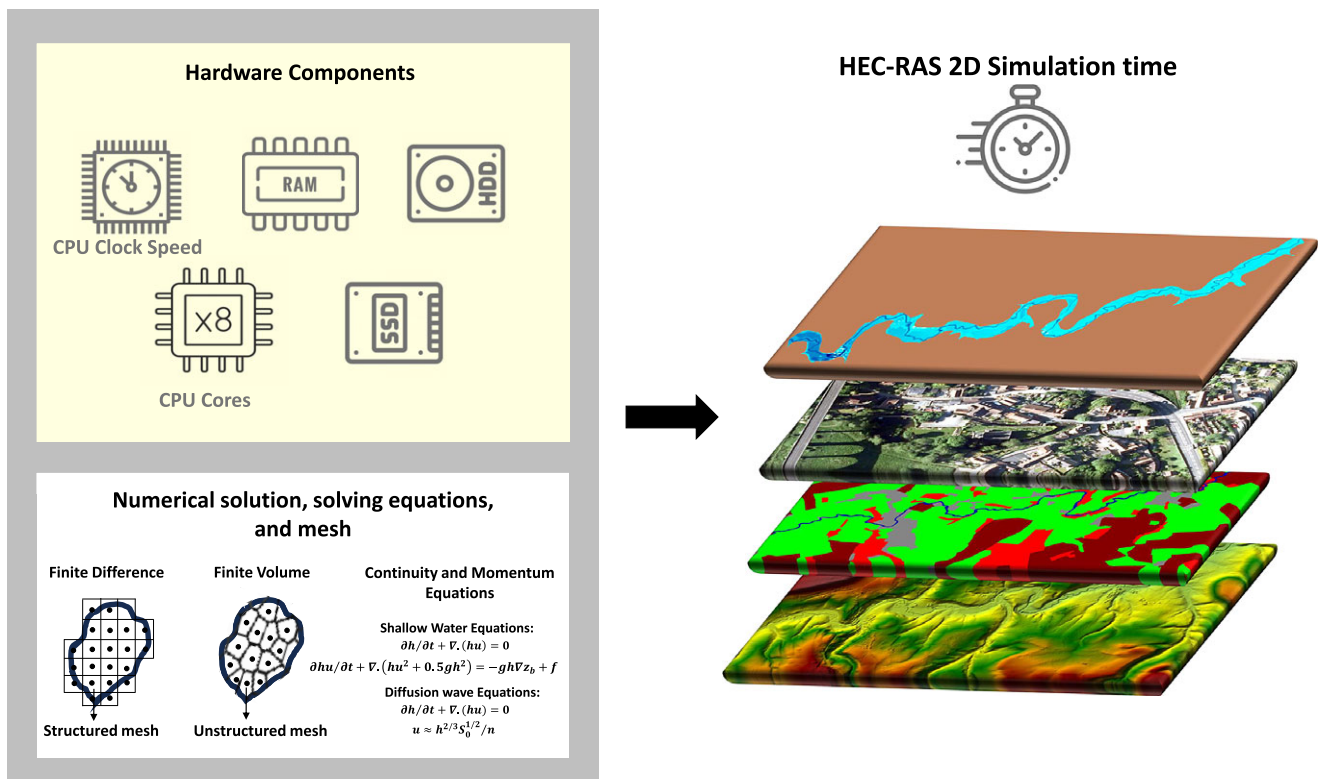
While HEC-RAS's manual (Hydrologic Engineering Center, 2021) has recommendations on hardware configurations to optimise simulation times, a noticeable gap remains in systematic research. Specifically, there is a lack of comprehensive studies on how varying hardware setups affect simulation durations. This gap is particularly evident in the limited guidance available for modellers regarding the optimal number of processor cores and the precise impact of this factor on reducing simulation time. Additionally, while the manual suggests a lower influence of RAM on simulation performance, it does not provide a detailed quantification of this effect or identify the specific threshold where additional RAM no longer contributes to performance improvements. Furthermore, the relationship between simulation time and factors, such as mesh type and resolution, solving equations (i.e., Diffusion Wave Equations [DWE] or Shallow Water Original Equations [SWOE]), and numerical solutions, that is, Finite Difference Classic (FDC) method or Finite Volume (FV) method across different hardware configurations has yet to be thoroughly explored. This novel research addresses these gaps through a comprehensive investigation of the effect of hardware configurations on the simulation time of the HEC-RAS 2D model (version 6.4.1) (Figure 1). Both Windows-based virtual machines on the Google Cloud Platform and a desktop PC were employed for conducting the tests in this study. A virtual machine emulates a physical computer in a digital format which can run operating systems and software, such as HEC-RAS. To evaluate the significance of optimised hardware, we analyse the impact of a balanced

configuration on time savings and cost efficiency using a hypothetical scenario of 1,000 simulations.

## Methods

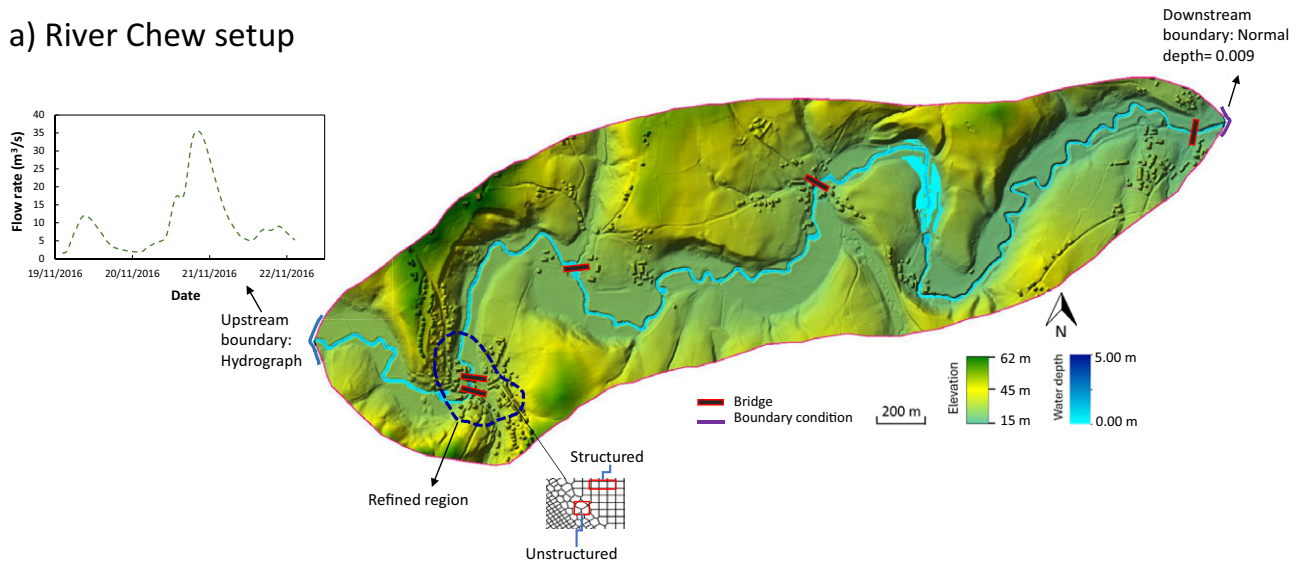
This section provides an overview of the setup for two Benchmark cases, which serve as baseline models for optimising hardware configurations in HEC-RAS 2D simulations. Details of setups and tests conducted on these two Benchmark cases, including different hardware configurations, as well as mesh, numerical solution, solving equations, and associated costs, are presented in the following.

**Benchmark Case 1:** The modelled area encompasses 2.83 km<sup>2</sup>, covering a 7.96-km stretch of the River Chew and the village of Pensford in Somerset, UK (Figure 2a). To create a flood inundation map, the Unsteady Flow module of HEC-RAS 2D was employed. The upstream boundary was defined using a flow hydrograph, whereas the downstream boundary was set with a normal depth (corresponding to a friction slope of 0.009). A hydrograph for the flood event from 20 to 23 November 2016 was generated using HEC-HMS (Figure 2a). While Sabeti *et al.* (2024) provides a detailed discussion of the model setup and calibration, this study only offers a brief overview of the HEC-HMS model configuration. The Soil Conservation Service (SCS)-Curve Number method was used to account for hydrological losses, and the SCS Unit Hydrograph method was selected for hydrological routing. For river channel routing, the Muskingum method was employed. The model accurately reproduced the event, achieving a 73% match with the measured flow rate. A Digital Elevation Model provided by SCALGO (<https://scalgo.com/>) was used in our HEC-RAS 2D setup with the following specifications: A Digital Terrain Model that incorporates

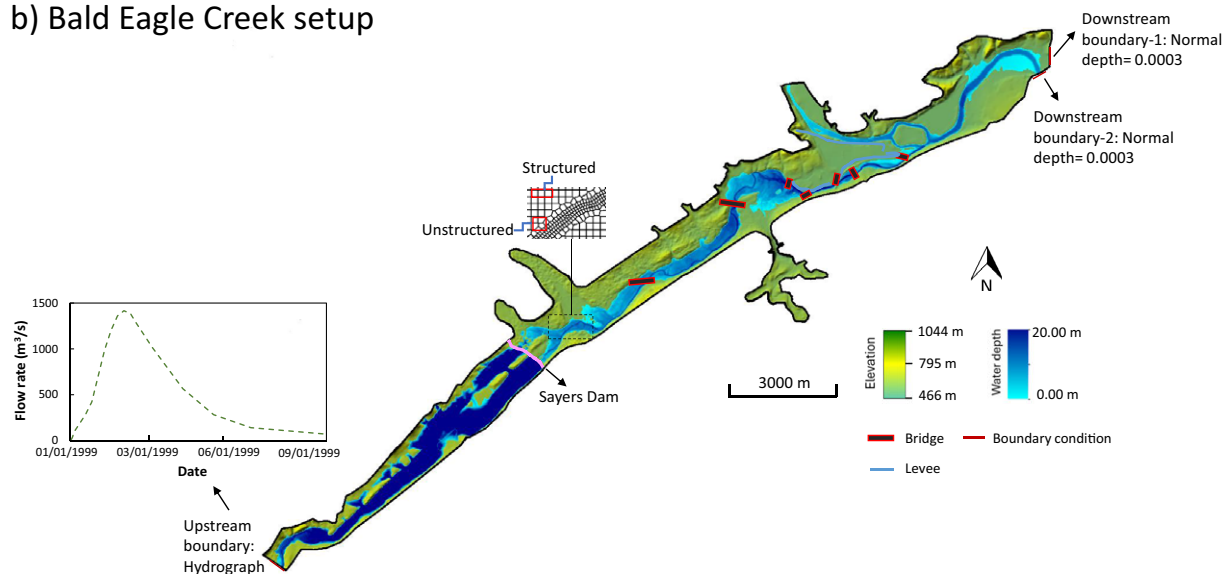


**Figure 1.** Overview of research methodology testing the effect of hardware configurations, numerical solution, solving equations, and mesh on HEC-RAS 2D simulation time.

## a) River Chew setup



## b) Bald Eagle Creek setup



**Figure 2.** HEC-RAS configuration: (a) River Chew, showing the water depth layer for the peak during the 20–23 November 2016 event. The Refined region in River Chew setup is applied only for Tests 15–18. (b) Bald Eagle Creek, depicting the water-depth layer for the peak during 01–09 January 1999.

buildings but not vegetation, resolution of 1 m, and Lidar originated (Figure 2a). The model includes five bridges, specified in the Geometry section of the HEC-RAS 2D.

In terms of the mesh, HEC-RAS 2D allows for the definition of varying mesh resolutions along with unstructured and structured (by default) meshes. The unstructured mesh consists of irregularly shaped cells with a maximum of eight sides per cell, offering greater flexibility in representing complex terrains and boundaries, although the irregularity in cell shapes and connectivity increases computational complexity. For the case of the River Chew, the initial setup used a structured mesh with a resolution of 10 m within the entire modelled area which was defined in the RASMapper. Accordingly, the first set of tests focused solely on the impact of hardware configurations on simulation time, using the River Chew setup with a structured mesh and a resolution of 10 m (Tests 1–9, Table 1). In order to determine whether increasing the number of cores improves the

efficiency of simulation time in HEC-RAS 2D as the number of cells increases, a series of tests is designed, in which we varied the mesh resolution from 10 to 5, 2.5, and 1 m in River Chew case (Tests 10–15, Table 1).

Additionally, we created a Refined region with a 0.5-m resolution covering the 0.071-km<sup>2</sup> residential area of Pensford, which has historically been affected by flooding (Figure 2a). The Perimeter spacing and Near repeats were set to 1 m and 4, meaning that for 4 m on either side of the Refined region, the mesh repeats the 1 m cells. The tests, including the Refined region, used an original mesh with resolution of 1 m covering the entire setup area, except for the Refined region. This modification was generated an unstructured mesh around the Refined region, allowing us to assess the combined impact of both very fine and unstructured meshes. Accordingly, a group of tests were designed to assess the impact of the speed of core and number of cores in simulation time within this setup (Tests 16–19, Table 1).

**Table 1.** Details of tests on the River Chew (RC) and Bald Eagle Creek (BEC) Example to assess the impact of CPU speed, number of CPUs, RAM capacity, Boot disk type (SSD vs. HDD), number of cells, numerical solutions (FV vs. FDC), solving equations (SWE vs. DWE), and mesh types on simulation time, along with the hourly cost of each virtual machine on GCP

No#	Platform type <sup>1</sup>	CPU (GHz)	Number of cores <sup>2</sup>	RAM (GB)	Boot disk <sup>3</sup>	Number of Cells <sup>4</sup>	Numerical solution	Solving equations	Benchmark Case	Mesh type	Hourly cost (\$) <sup>5</sup>
1	C2	3.10	15	120	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	2.00
2	C3	2.70	44	352	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	6.65
3	C2	3.10	30	240	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	3.94
4	C2	3.10	8	64	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	1.09
5	C2	3.10	4	64	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	0.53
6	E2	2.20	8	32	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	0.87
7	E2	2.20	8	64	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	0.96
8	E2	2.20	8	128	SSD (250 GB)	37,866	FDC	DWE	RC	Structured	1.15
9	C2	3.10	8	64	HDD (250 GB)	37,866	FDC	DWE	RC	Structured	1.05
10	C2	3.10	8	64	SSD (250 GB)	169,749	FDC	DWE	RC	Structured	1.09
11	C2	3.10	30	240	SSD (250 GB)	169,749	FDC	DWE	RC	Structured	3.94
12	C2	3.10	30	240	SSD (250 GB)	573,698	FDC	DWE	RC	Structured	3.94
13	C2	3.10	8	64	SSD (250 GB)	573,698	FDC	DWE	RC	Structured	1.09
14	C2	3.10	30	240	SSD (250 GB)	1,532,276	FDC	DWE	RC	Structured	3.94
15	C2	3.10	8	64	SSD (250 GB)	1,532,276	FDC	DWE	RC	Structured	1.09
16	E2	2.20	8	64	SSD (250 GB)	1,981,164	FDC	DWE	RC	Structured and Unstructured	3.94
17	C2	3.10	30	240	SSD (250 GB)	1,981,164	FDC	DWE	RC	Structured and Unstructured	3.94
18	C2	3.10	8	64	SSD (250 GB)	1,981,164	FDC	DWE	RC	Structured and Unstructured	1.09
19	C2	3.10	4	64	SSD (250 GB)	1,981,164	FDC	DWE	RC	Structured and Unstructured	0.53
20	E2	2.20	8	64	SSD (250 GB)	37,866	FV	DWE	RC	Structured	0.96
21	C2	3.10	8	64	SSD (250 GB)	37,866	FV	DWE	RC	Structured	1.09
22	C2	3.10	30	240	SSD (250 GB)	37,866	FV	DWE	RC	Structured	3.94
23	C2	3.10	4	64	SSD (250 GB)	37,866	FV	DWE	RC	Structured	0.53
24	E2	2.20	8	64	SSD (250 GB)	37,866	FDC	SWOE	RC	Structured	0.96
25	C2	3.10	8	64	SSD (250 GB)	37,866	FDC	SWOE	RC	Structured	1.09
26	C2	3.10	30	240	SSD (250 GB)	37,866	FDC	SWOE	RC	Structured	3.94
27	C2	3.10	4	64	SSD (250 GB)	37,866	FDC	SWOE	RC	Structured	0.53
28	E2	2.20	8	64	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	0.96
29	C2	3.10	8	64	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	1.09
30	C2	3.10	30	240	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	3.94
31	C2	3.10	4	64	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	0.53
32	Desktop PC	5.2	64	7	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	N/A
33	Desktop PC	5.2	64	8	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	N/A
34	Desktop PC	5.2	64	9	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	N/A
35	Desktop PC	5.2	64	12	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	N/A

(Continued)

Table 1. (Continued)

No#	Platform type <sup>1</sup>	CPU (GHz)	Number of cores <sup>2</sup>	RAM (GB)	Boot disk <sup>3</sup>	Number of Cells <sup>4</sup>	Numerical solution	Solving equations	Benchmark Case	Mesh type	Hourly cost (\$) <sup>5</sup>
36	Desktop PC	5.2	64	24	SSD (250 GB)	24,718	FDC	DWE	BEC	Structured and Unstructured	N/A
37	Desktop PC	5.2	64	24	SSD (250 GB)	24,718	FDC	DWE	RC	Structured and Unstructured	N/A

<sup>1</sup>C2, E2, and C3 are different series of virtual machines on GCP, each with similar specifications. The number, such as the “2” in C2, represents the generation of the series.

<sup>2</sup>In the creation stage of the virtual machines, the virtual CPU-to-core ratio is configured to 1:1 (for all virtual machines), meaning each core is assigned one virtual CPU.

<sup>3</sup>In GCP Boot disk setup, SSD Persistent Disk use SSD storage, while Standard Persistent Disk is based on HDD storage.

<sup>4</sup>For the RC case, the number of computational cells increases with mesh refinement: 10 m mesh results in 37,866 cells, 5 m mesh in 169,749 cells, 2.5 m mesh in 573,698 cells, 1 m mesh in 1,532,276 cells, and a combination of 1 m mesh with a 0.5-m refined region results in 1,981,164 cells. In the BEC case, the mesh size is 250 m.

<sup>5</sup>For all tests, the US-central1 (lowa) region was selected for costing.

The initial setup of the River Chew case utilised the FDC method along with DWE. To evaluate the impact of the numerical solution, and solver equations, a series of tests was specifically designed to assess the performance of the alternative approaches, the FV method and SWOE (Tests 20–23 and Test 24–27, respectively; Table 1). These variations allowed us to evaluate the impact of different solver equations and numerical approaches on simulation time across different hardware configurations.

In the simulations of the River Chew case, the flow velocity ( $v$ ) was 1.5 m/s. For the setups with mesh sizes ( $\Delta s$ ) of 10 m (Tests 1–9 and Tests 20–23), 5 m (Tests 10–11), 2.5 m (Tests 12–13), 1 m (Tests 14–15) (Table 1), and the setup with mesh size of 1 m and the Refined region of 0.5 m (Tests 16–19), the DWE was used. The time step ( $\Delta t$ ) for these tests was set to 10, 10, 5, 2, and 1 s, respectively, to satisfy the Courant number criterion  $C = v \frac{\Delta t}{\Delta s} \leq 1.0$  (with a maximum  $C = 5.0$ ), as recommended by the HEC-RAS manual (Hydrologic Engineering Center, 2021), ensuring model stability. For simulations using the Shallow Water Equations (Tests 24–27), which impose a stricter Courant number criterion of maximum  $C = 3.0$ , the mesh size was 10 m and we set the  $\Delta s$  to 10 s. Finally, the number of cores in the setup was matched with the number of the virtual machine or desktop PC cores, not “All Available” option. For instance, if the virtual machine’s core had 10 cores, we set the number of cores in the setup to 10.

**Benchmark Case 2:** For the second Benchmark Case, we utilised one of the official examples provided by HEC-RAS, specifically the Bald Eagle Creek Example (Figure 2b). Within this example, the “Single 2D Area with Bridges and Break-lines” scenario was selected. This simulation covers a region of 106.84 km<sup>2</sup> with a grid spacing ( $\Delta s$ ) of 250 m, resulting in 24,748 computational cells. The majority of the mesh was structured; however, certain regions, such as those with Breaklines (e.g., the area shown in Figure 2a within the main channel of this setup), were unstructured. The modelled area includes an urban region of 29.30 km<sup>2</sup> and features the Sayers Dam, seven bridges, and three levees. Regarding boundary conditions, we maintained the original setup, where the two downstream boundaries were both set to a normal depth, each with a friction slope of 0.0003, and the upstream boundary was defined by a flow hydrograph for the event occurring from 1 to 4 January 1999. It should be noted that the hydrograph in Figure 2b displays the full hydrograph from 1 to 9 January 1999. However, the original simulation period, from 1 to 4 January 1999, was retained as the default setting in our tests. In the Unsteady Flow Analysis settings, the time step was set to 5 s by default. In terms of the computational methods, the numerical solution was implemented using the FDC method. However, to ensure consistency with most of our tests on the River Chew (Benchmark Case 1), the solving equations in these setups (Benchmark Case 2) were set to

DWE. The Bald Eagle Creek example was hydraulically more complex than the River Chew case due to its coverage of a larger urban area, the presence of a dam, multiple bridges, and levees. This foundation allowed us to conduct an additional series of tests to determine the optimal hardware configurations for efficient HEC-RAS 2D simulation time (Tests 28–36, Table 1).

Leveraging the ability to adjust the number of cores in the Computation Options and Tolerances section of HEC-RAS 2D, alongside the virtual machines (Tests 28–31, Table 1) a desktop PC was employed in our assessment (Tests 32–36, Table 1) to identify the optimal core number for minimising simulation time. Bald Eagle Creek example was used in Tests 32–36. This assessment was considered to provide insights for modellers who do not have access to hardware-adjustable virtual machines.

The desktop PC is also used to conduct Test 37 which allows for a comparison with Test 4 in terms of CPU speed, as both tests share identical conditions except for the CPU speed.

The setups of Tests 1–9 were designed to be identical in terms of numerical solutions, solving equations, and the modelled region (River Chew). Consequently, any variation in simulation time was solely attributable to hardware configurations. This design was created a foundation for conducting a hypothetical analysis based on the assumption of 1,000 simulations, which was a common scenario in uncertainty analysis using tools, such as the Monte Carlo approach. By extrapolating the time for a single simulation to 1,000 runs and calculating the total cost using the hourly rate, we assessed how different hardware configurations impact overall time and cost. This analysis highlighted the importance of an optimised hardware setup in HEC-RAS 2D models in large-scale simulations. To quantify this impact, we applied an efficiency score ( $Es$ ) that incorporates both total time and total cost of simulations, with an equation derived from the approach presented in Fishburn’s (1979) study, as outlined below:

$$Es = 1 - (w_{\text{time}} \times \sqrt{n_{\text{time}}} + w_{\text{cost}} \times \sqrt{n_{\text{cost}}}) \quad (1)$$

where  $w_{\text{time}}$  and  $w_{\text{cost}}$  are the weights assigned to time and cost, respectively. Both weight values were set to 0.5 in this analysis, giving equal importance to each.  $n_{\text{time}}$  and  $n_{\text{cost}}$  are the normalised values of total time and total cost, derived for each test using min–max normalisation method as described by Amiri et al. (2014).

Normalised Sensitivity Coefficient (NSC) (Equation 2) was used as a sensitivity analysis index (Hamby, 1995), to quantify the sensitivity of an output variable (i.e., simulation time) to changes in input parameters (i.e., hardware configurations). NSC

Table 2. Simulation times for all 37 tests are provided

No#	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
Simulation time (min)	11.10	17.80	17.87	10.40	12.15	17.45	16.93	16.97	10.98	45.82	72.48	294.95	198.45	794.07	598.82	1,274.62	1,248.81	906.64	1,053.73
No#	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	
Simulation time (min)	17.12	10.47	17.92	12.44	25.83	16.46	27.04	20.68	18.58	14.08	20.81	15.81	09.17	08.72	08.98	09.73	11.02	08.55	

Note: Refer to Table 1 for detailed information on each test.

nondimensionalises parameter sensitivity, allowing for direct comparison and ranking of the most influential variables.

$$NSC = \frac{\Delta SimT / SimT_0}{\Delta p_i / p_{i0}} \times 100 \quad (2)$$

where  $SimT$  is the simulation time,  $p_i$  is the input parameter (i.e., RAM), and  $SimT_0$  and  $p_{i0}$  are the reference values of the simulation time and the input parameter, respectively. The tests used in this assessment were from Tests 1–9 and Test 37, with the setup of River Chew (Benchmark Case 1).

## Results

In this section, the effects of the following factors on simulation time were analysed: (1) Boot disk and RAM, (2) CPU speed, (3) number of cores, (4) NSC of hardware configurations, (5) Mesh, (6) numerical solution and solving equations, and (7) time and cost efficiency in large-scale simulations.

### Boot disk and RAM

Using Tests 4 and 9, we assessed the impact of the Boot disk on simulation time (Table 1). For these tests, all parameters, including hardware configurations, mesh, and computational methods, were kept constant. The results indicate Test 4 which used a SSD, completed the simulation 0.58 min faster than Test 9 with a HDD Boot disk (Table 2). This difference is attributed to the read and write speeds, with the SSD achieving 240 MB/s on average compared to 122 MB/s for the HDD in these setups.

Regarding the effect of RAM on simulation time, three tests were conducted (Tests 6, 7, and 8). All other parameters were kept constant, while the RAM set to 32, 64, and 128 GB, respectively, for these three tests. The results show an improvement in simulation time from 32 GB in Test 6 to 64 GB in Test 7, with a reduction of 0.52 min. However, no significant improvement was observed between 64 GB in Test 7 and 128 GB in Test 8. Increasing the RAM to 128 GB did not improve the simulation time instead; it caused a slight increase (0.04 min) in the overall simulation duration. However, this increase was minimal and almost negligible. This series of tests on RAM suggests that 64 GB of RAM could be considered optimal for HEC-RAS 2D setups in a well-balanced hardware configuration.

### CPU speed

The results indicate the simulation time was highly influenced by CPU speed as demonstrated by Test 4 and Test 7 (Table 2). Both tests had an identical number of cores, RAM and SSD but differed in CPU speeds (3.10 and 2.20 GHz, respectively), leading to Test 4 completing the simulation 6.57 min faster than Test 7. The significant influence of CPU speed is evident in Tests 28 and 29 (Table 2), where the Bald Eagle Creek setup, a more hydraulically complex scenario, resulted in the higher-speed CPU (Test 29 with 3.10 GHz), completing the simulation 4.5 min faster than Test 28 with 2.20 GHz CPU speed. Another pair of tests that further emphasise the impact of CPU speed is Test 33 and Test 29 (Table 1). In these tests, the desktop PC in Test 33 had a CPU speed of 5.20 GHz, compared to 3.10 GHz in Test 4, while all other parameters were kept constant. This resulted in a 5.36-min shorter simulation time for Test 33 (Table 2), representing a 38.07% improvement. Furthermore, comparing Test 5 (4 cores at

a)

Finished Processing Geometry

**Performing Unsteady Flow Simulation HEC-RAS 6.4.1 June 2023**

**Unsteady Input Summary:**  
2D Unsteady Diffusion Wave Equation Set (fastest)  
2D number of Solver Cores: 44

Overall Volume Accounting Error in 1000 m<sup>3</sup>: 0.6603  
Overall Volume Accounting Error as percentage: 0.03293  
Please review "Computational Log File" output for volume accounting details

**Writing Results to DSS**

**Finished Unsteady Flow Simulation**

1D Post Process Skipped (simulation is all 2D)

**Computations Summary**

Computation Task	Time(hh:mm:ss)
Completing Geometry, Flow and Plan	1
Preprocessing Geometry	<1
Unsteady Flow Computations	17:47
Complete Process	17:48

Computation Speed	Simulation/Runtime
Unsteady Flow Computations	243x
Complete Process	243x

b)

Finished Processing Geometry

**Performing Unsteady Flow Simulation HEC-RAS 6.4.1 June 2023**

**Unsteady Input Summary:**  
2D Unsteady Diffusion Wave Equation Set (fastest)  
2D number of Solver Cores: 15

Overall Volume Accounting Error in 1000 m<sup>3</sup>: 0.6605  
Overall Volume Accounting Error as percentage: 0.03294  
Please review "Computational Log File" output for volume accounting details

**Writing Results to DSS**

**Finished Unsteady Flow Simulation**

1D Post Process Skipped (simulation is all 2D)

**Computations Summary**

Computation Task	Time(hh:mm:ss)
Completing Geometry, Flow and Plan	1
Preprocessing Geometry	<1
Unsteady Flow Computations	14:37
Complete Process	14:39

Computation Speed	Simulation/Runtime
Unsteady Flow Computations	295x
Complete Process	295x

c)

Finished Processing Geometry

**Performing Unsteady Flow Simulation HEC-RAS 6.4.1 June 2023**

**Unsteady Input Summary:**  
2D Unsteady Diffusion Wave Equation Set (fastest)  
2D number of Solver Cores: 8

Overall Volume Accounting Error in 1000 m<sup>3</sup>: 0.6599  
Overall Volume Accounting Error as percentage: 0.03292  
Please review "Computational Log File" output for volume accounting details

**Writing Results to DSS**

**Finished Unsteady Flow Simulation**

1D Post Process Skipped (simulation is all 2D)

**Computations Summary**

Computation Task	Time(hh:mm:ss)
Completing Geometry, Flow and Plan	1
Preprocessing Geometry	<1
Unsteady Flow Computations	10:22
Complete Process	10:24

Computation Speed	Simulation/Runtime
Unsteady Flow Computations	416x
Complete Process	415x

d)

Finished Processing Geometry

**Performing Unsteady Flow Simulation HEC-RAS 6.4.1 June 2023**

**Unsteady Input Summary:**  
2D Unsteady Diffusion Wave Equation Set (fastest)  
2D number of Solver Cores: 4

Overall Volume Accounting Error in 1000 m<sup>3</sup>: 0.6591  
Overall Volume Accounting Error as percentage: 0.03287  
Please review "Computational Log File" output for volume accounting details

**Writing Results to DSS**

**Finished Unsteady Flow Simulation**

1D Post Process Skipped (simulation is all 2D)

**Computations Summary**

Computation Task	Time(hh:mm:ss)
Completing Geometry, Flow and Plan	1
Preprocessing Geometry	<1
Unsteady Flow Computations	12:09
Complete Process	12:10

Computation Speed	Simulation/Runtime
Unsteady Flow Computations	355x
Complete Process	355x

**Figure 3.** Summaries of HEC-RAS computation for (a) Test 2, (b) Test 3, (c) Test 4, and (d) Test 5 (see details of configuration of these tests in Table 1 and results in Table 2).

3.10 GHz, Figure 3d) and Test 2 (44 cores at 2.70 GHz, Figure 3a), Test 5 finished the simulation 5.67 min faster despite Test 2's higher core count (Table 2). This indicates that core speed has a more significant impact on simulation efficiency than the number of cores.

### Number of cores

The results of Tests 32–36, conducted using the Bald Eagle Creek case, show that the difference in simulation time between the longest (Test 36) and shortest (Test 33) was 2.3 min, driven by the varying number of cores (Table 2). Test 33, with 8 cores, showed

**Table 3.** Quantifying the impact of changing each tested hardware component on simulation time using Normalised Sensitivity Coefficients (NSC) index, calculated by Equation 2

Hardware component	Range of change	Test number	NSC value
CPU speed	2.20–3.10 GHz	7 and 4	+94.27%
CPU speed	3.10–5.20 GHz	4 and 37	+26.25%
Number of cores	4–8 Cores	5 and 4	+14.40%
Number of cores	8–15 Cores	4 and 1	–7.70%
Number of cores	15–30 Cores	1 and 3	–60.99%
RAM	32–64 GB	6 and 7	+2.98%
RAM	64–128 GB	7 and 8	–0.24%
Boot disk speed	122 MB/s (HDD) to 240 MB/s (SSD)	4 and 9	+5.46%

the shortest simulation time of 8.72 min. In all these Tests (32–36), the same desktop PC was used, and all other parameters, except the number of cores, were kept constant. This finding was consistent in Tests 29, 30, and 31, which used the Bald Eagle Creek setup and were conducted on GCP virtual machines (C2). The results indicated that Test 29, with 8 cores, completed the simulation 1.73 min faster than Test 31, which had fewer cores, and 6.73 min faster than Test 30, which had more cores (Table 2). Similarly, in another set of tests using the River Chew setup (Tests 1, 3, 4, and 5), Test 4, with 8 cores (Figure 3c), completed the simulation 0.7 min faster than Test 1, 7.47 min faster than Test 3 (Figure 3b), and 1.75 min faster than Test 5, due to the prior finding that RAM sizes greater than 64 GB do not significantly affect simulation time. Although the RAM sizes in these setups varied, they were all at least 64 GB, which facilitated this comparison. Based on these various comparisons, we conclude that using 8 cores results in shorter simulation times than either fewer or more cores.

### NSC of hardware configurations

The calculated NSC values (Table 3) using Equation 2 revealed that CPU speed has the most significant influence, followed by the number of cores, whereas RAM and Boot disk having a lesser impact. Notably, the NSC for CPU speed showed a 94.2% positive impact as the speed increased from 2.20 to 3.10 GHz (Table 3). Further increases from 3.10 to 5.20 GHz resulted in a more moderate impact of 26.2%, suggesting that further substantial reductions in simulation time could require disproportionately higher CPU speeds. Increasing the number of cores from 4 to 8 led to a clear 14.40% improvement, but additional cores showed diminishing returns, with performance decreases of –7.7% from 8 to 15 cores, and a –60.9% drop from 15 to 30 cores (Table 3). While upgrading the Boot disk from HDD to SSD and increasing RAM from 32 to 64 GB showed a positive impact, both improvements were within a modest range of less than 6%; further increasing the RAM from 64 to 128 GB resulted in a relatively minor negative impact.

### Mesh

The results from Tests 10 to 15 show that, although higher grid resolutions along with the smaller time steps increased simulation times, setups with 8-core CPUs consistently resulted in shorter

overall simulation durations, even as the number of cells rose to 1,532,276. The same outcome was observed in Tests 16 to 19, where the number of cells reached 1,981,164, and the mesh was partially unstructured with the Refined region of 0.5 m. Additionally, Tests 16–19 confirmed earlier findings, demonstrating that setups with higher core speeds could significantly reduce simulation times. In Test 18, the simulation using a 3.10-GHz CPU was completed 367.98 min (6.13 h) faster than in Test 16, which used a 2.20 GHz CPU. It is important to note that although HEC-RAS manual suggests that a larger number of cells necessitates a higher number of cores for optimal performance (Hydrologic Engineering Center, 2021), our research indicates that even with as many as 1,981,164 cells, a higher core number does not necessarily translate to improved efficiency. Note that HEC-RAS does not specify what constitutes a “high number” of cells, and their advice could still hold true if the number of cells exceeds those explored here.

### Numerical solution and solving equations

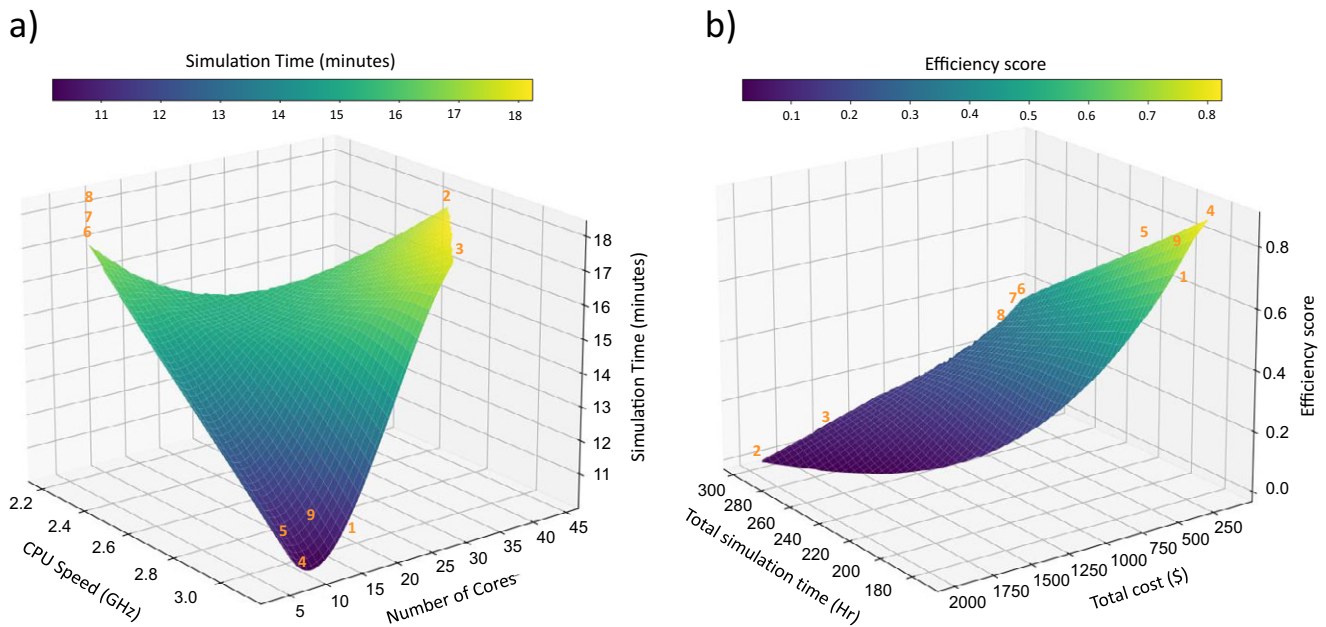
In terms of the numerical solution, Tests 20–23 highlight two key findings that align with our previous observations: (i) the higher core speeds, along with the use of (ii) 8 cores, generate the shortest simulation times. Although these tests employed the FV method instead of the FDC method, the outcomes remained consistent. Specifically, Tests 20 and 21 demonstrate that the configuration with a core speed of 3.10 GHz (Test 21) completed the simulation 6.65 min faster than the configuration with a core speed of 2.20 GHz (Test 20). A review of Tests 20, 22, and 23, which focused on the impact of core count on simulation time, reveals that the setup in Test 21 with 8 cores completed the simulation faster than both the 4-core configuration in Test 23 and the 30-core configuration in Test 22. Comparing Tests 20, 21, 22, and 23, which used the FV method, with their corresponding Tests 7, 4, 3, and 5 that utilised the Finite Difference method, reveals a longer simulation time for all tests using the FV method. This rise in simulation time due to use of FV method is relatively small, with the maximum difference being less than 0.29 min across all four Tests 20–23 and their counterparts. The additional computational complexity of the FV method, particularly in relation to flux calculation and the enforcement of conservation laws, could result in longer simulation time than setups using the Finite Difference method.

Regarding the solving equations, results of Tests 24–27, where the DWE was replaced by the Shallow Water Equations, indicate that our previous findings still hold. Specifically, a higher core speed (3.10 GHz in Test 25 vs. 2.20 GHz in Test 24) and the use of 8 cores (Test 25) instead of 30 cores (Test 26) or 4 cores (Test 27) results in the shortest simulation time. Comparing Tests 24, 25, 26, and 27 with their corresponding Tests 7, 4, 3, and 5 shows the expected outcome of longer simulation times due to the use of the Shallow Water Equations instead of the DWE. This increase is more significant with changes in the numerical solution method. On average, the simulation time increased by 8.15 min for Tests 24–27 compared to their counterparts. This increase is due to the fact that the Shallow Water Equations account for full dynamic flow behaviour, including inertia and momentum, which require more complex calculations. In contrast, the DWE simplifies flow dynamics by neglecting these terms, leading to faster simulations.

### Time and cost efficiency in large-scale simulations

The Efficiency scores calculated using Equations 1 for the hypothetical scenario of 1,000 simulations across Tests 1–9 reveal significant differences in performance. Test 4 with higher CPU speed,





**Figure 4.** 3D surface plots illustrating (a) the relationship between CPU speed, number of cores, and simulation time for Tests 1–9. (b) Efficiency scores as a function of total time and total cost for 1,000 simulations across various configurations (Tests 1–9). In both (a) and (b), the test numbers are shown in orange bold colour on the 3D surface.

8 cores, and 64 GB of RAM demonstrated the highest efficiency, achieving a score of 0.86 and completing 1,000 simulations in 173 h at a total cost of 188 USD, outperforming all other tests (Figure 4b). This result highlights the importance of an optimised system configuration. In contrast, Tests 2 and 3, with unbalanced hardware configurations (15 cores with core speed of 3.10 GHz; and 44 cores with 2.7 GHz, respectively), led to much longer simulation times and significantly higher costs (Figure 4b). Specifically, Test 2 required 296 h and a total cost of 1,972 USD, while Test 3 took 297 h and cost 1,173 USD. These inefficiencies resulted in hundreds of additional hours and dollars. Figure 4b presents a detailed comparison of the total simulation time, cost, and Efficiency scores for Tests 1–9.

## Conclusion

This study highlights that for achieving the shortest simulation time, understanding the optimal balance of available hardware resources is important. Specifically, while faster core speeds reduce simulation times, merely having the highest number of cores or the largest amount of RAM does not guarantee improved outcomes. Notably, a configuration with 8 cores and 64 GB of RAM delivers superior performance compared to setups with either higher or lower core counts or varying RAM sizes. This suggests that optimal hardware configurations for HEC-RAS simulations involve more than just maximising individual components.

**Open peer review.** To view the open peer review materials for this article, please visit <http://doi.org/10.1017/wat.2024.11>.

**Data availability statement.** All data used in this study are given in the body of the article.

**Acknowledgements.** This study acknowledges the funding provided by the Leverhulme Trust and the Google Cloud Research Credits program. Support from SCALGO (<https://scalgo.com/>) in providing access to their platform is appreciated. Gratitude is extended to the University of Bath Institutional Open Access Fund for covering the Article Processing Charge of this paper, allowing it

to be published as Open Access. We would like to express our gratitude to Mr. Luis Partida for his valuable technical advice on our HEC-RAS setup.

**Author contribution.** Conceptualisation: R.S. Investigation: R.S., T.K., and B.R. Visualisation: R.S., B.R., and I.S. Writing and Reviewing: R.S., T.K., and I.S.

**Financial support.** This study was financially supported by Leverhulme Trust project grant: RPG-2022-306 and Google Cloud Research Credits program: GCP19980904.

**Competing interest.** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

- Amiri IS, Akanbi OA and Fazeldehkordi E (2014) *A machine-learning approach to phishing detection and defense*. Syngress. [https://www.researchgate.net/profile/Iraj\\_Sadegh\\_Amiri/publication/320702619\\_A\\_Machine\\_Learning\\_Approach\\_to\\_Phishing\\_Detection\\_and\\_Defense/data/5b3054e8aca2720785e3ae37/A-Machine-Learning-Approach-to-Phishing-Detection-and-Defense.pdf](https://www.researchgate.net/profile/Iraj_Sadegh_Amiri/publication/320702619_A_Machine_Learning_Approach_to_Phishing_Detection_and_Defense/data/5b3054e8aca2720785e3ae37/A-Machine-Learning-Approach-to-Phishing-Detection-and-Defense.pdf)
- Brunner GW (2018) *Benchmarking of the HEC-RAS Two-Dimensional Hydraulic Modelling Capabilities*. Davis CA: US Army Corps of Engineers, 1–137.
- Buttinger-Kreuzhuber A, Konev A, Horváth Z, Cornel D, Schwerdorf I, Blöschl G and Waser J (2022) An integrated GPU-accelerated modeling framework for high-resolution simulations of rural and urban flash floods. *Environmental Modelling and Software* **156**. 105480. <https://doi.org/10.1016/j.envsoft.2022.105480>.
- Fishburn PC (1979) *Utility Theory for Decision Making*. New York, NY: John Wiley and Sons. <https://apps.dtic.mil/sti/citations/AD0708563>.
- Hamby DM (1995) A comparison of sensitivity analysis techniques. *Health Physics* **68**, 195–204.
- Hydrologic Engineering Center (2021) *HEC-RAS 2D Modeling User's Manual*. Davis, CA: U.S. Army Corps of Engineers,
- Lacasta A, Morales-Hernández M, Murillo J and García-Navarro P (2014) An optimized GPU implementation of a 2D free surface simulation model on unstructured meshes. *Advances in Engineering Software* **78**, 1–15. <https://doi.org/10.1016/j.advengsoft.2014.08.007>.
- Morales-Hernández M, Sharif MB, Kalyanapu A, Ghafoor SK, Dullo TT, Gangrade S, Kao S-C, Norman MR and Evans KJ (2021) TRITON: a multi-

- GPU open source 2D hydrodynamic flood model. *Environmental Modelling and Software* **141**, 105034. <https://doi.org/10.1016/j.envsoft.2021.105034>.
- Loucks DP** (2023) *Hydroinformatics: A Review and Future Outlook*. Cambridge: Cambridge University Press, 1–26.
- Neal J, Fewtrell T and Trigg M** (2009) Parallelisation of storage cell flood models using OpenMP. *Environmental Modelling and Software* **24** (7), 872–877. <https://doi.org/10.1016/j.envsoft.2008.12.004>.
- Neal JC, Fewtrell TJ, Bates PD and Wright NG** (2010) A comparison of three parallelisation methods for 2D flood inundation models. *Environmental Modelling and Software* **25** (4), 398–411. <https://doi.org/10.1016/j.envsoft.2009.11.007>.
- Néelz S and Pender G** 2010. Benchmarking of 2D hydraulic modelling packages. SC080035/R2 Environmental Agency, Bristol, 169.
- Sabeti R, Stamataki I and Kjeldsen TR** (2024) Reconstructing the 1968 river chew flash flood: merging a HEC-RAS 2D hydraulic modelling approach with historical evidence. *Geomatics, Natural Hazards and Risk* **15**(1), 2377655. <https://doi.org/10.1080/19475705.2024.2377655>.
- Sanders BF and Schubert JE** (2019) PRIMo: parallel raster inundation model. *Advances in Water Resources* **126**, 79–95. <https://doi.org/10.1016/j.advwatres.2019.02.007>.
- Sanders BF, Schubert JE and Detwiler RL** (2010) ParBreZo: a parallel, unstructured grid, Godunov-type, shallow-water code for high-resolution flood inundation modeling at the regional scale. *Advances in Water Resources* **33** (12), 1456–1467. <https://doi.org/10.1016/j.advwatres.2010.07.007>.
- Dazzi S, Shustikova I, Domeneghetti A, Castellarin A and Vacondio R** (2021) Comparison of two modelling strategies for 2D large-scale flood simulations. *Environmental Modelling and Software* **146**, 105225. <https://doi.org/10.1016/j.envsoft.2021.105225>.