Hydrology & Water
Resources Symposium
12-15 November | Sydney

National Committee on Water Engineering

ISBN number for HWRS 2023 is 978-1-925627-81-7

Achieving accuracy, stability, and parallelism in a new 1D hydraulic scheme for TUFLOW HPC

Greg Collecutt, Duncan Kitts, Russell Jones

TUFLOW

greg.collecutt@tuflow.bmt.org

ABSTRACT

Recent trends in 2D hydraulic modelling have seen an increased ability to resolve surface channels within the 2D domain. However, for urban environments representing the storm water network remains an essentially 1D component, elsewhere hydraulic structures are often best represented in 1D, and 1D can still be the pragmatic choice where large scale and/or fast run-times are needed.

The 2D finite-volume scheme used in TUFLOW HPC has been adapted to represent 1D network elements. The new scheme has been benchmarked against a suite of tests developed by the English Environment Agency, which include subcritical flow, supercritical flow, hydraulic jumps, branched networks, irregular channels, tidal channel transient response, and flood wave propagation. The results are extremely encouraging with close to perfect agreement with analytical solutions.

In addition to its accuracy, the new scheme is unconditionally stable provided typical Courant and celerity timestep considerations are met. Further still, being an explicit formulation, the scheme is parallelised for multicore/GPU execution. This enables it to be integrated with the existing TUFLOW HPC code base and executed efficiently on GPU devices concurrently with the 2D solution.

The relevant details of the scheme are presented along with benchmarking results.

INTRODUCTION

Advancements in computing power over the last decade have come from both an increase in single thread performance as well as an increased number of threads available on a single device. Single thread performance of Central Processinc Units (CPUs) has increased by approximately a factor of 2.5 over the last 10 years, while multi-thread CPU performance has increased by approximately a factor of 5 over the same period [cpubenchmark]. However, in recent years a substantially more significant gain in performance has come with the availability of general compute on Graphics Processing Unit (GPU) hardware, where current devices have a processor core count of over 10^4. For example, TUFLOW HPC running a given medium-sized 2D benchmark model on a 32 core Intel i9-13900KS CPU takes 108 minutes, and the same model (using the same code but compiled for GPU) on an Nvidia RTX4090 GPU takes just 2.5 minutes [TUFLOW Hardware Benchmarks] – i.e. 43x faster on GPU compared to multicore CPU (as of 2022 hardware). The speed difference becomes even more significant for larger models.

Interestingly, TUFLOW Classic, which uses an Alternating Direction Implicit (ADI) 2D scheme, is able to run the same benchmark model in 33 minutes on a single thread. This is 3x faster than the

HPC solution on multi-core CPU, but 13x slower than the HPC solution on GPU. So the question must be asked: how fast could TUFLOW Classic be made to run on GPU? This leads us to the crux of the matter: the key to fully utilising the compute power of current GPUs is using an algorithm that can be parallelised. TUFLOW HPC utilises an explicit finite volume formulation of the 2D Shallow Water Equations (SWE), which is ideally suited to concurrent computation [Collecutt and Syme 2017]. Whereas in the case of TUFLOW Classic, the ADI scheme used was considered to be practically impossible to efficiently paralllelise.

The use of 1D elements in a hydraulic surface water model remains common place for both open channel flow paths and for urban storm water pipe networks. Recent advances in 2D modelling schemes, in particular sub-grid-sampling (SGS) [Huxley et.al 2022] for more accurate cell storage and face conveyance representation combined with 2nd order spatial interpolation scheme, enable open surface channels to be more accurately represented in the 2D domain. The need to represent rivers, tributaries, and urban surface drain channels with 1D elements has significantly reduced [Gao et. al. 2022]. However, cities and their storm water networks will, for the foreseeable future, require complex 1D networks to represent the pipes and pipe junctions (manholes). Also, hydraulic structures are often still best represented in 1D, and 1D can still be the pragmatic choice where large scale and/or fast run-times are needed (eg. Monte Carlo analyses; real-time flood forecasting). The need for coupled 1D/2D modelling, increasingly with larger and more complex 1D networks, will continue for some time.

To statisfy the input/output needs of thousands and tens of thousands of cores working in parallel, GPU designers utilise fast on-card RAM with a dedicated memory bus and cache controller. Currently, GPU enabled software utilise both the CPU and the GPU, with binary execution code being run on each asynchronously, and with regular synchronisation points usually involving transfer of data from one to the other. Transfer of data between the GPU RAM and the GPU cores is extremely fast, but transfer of data between the main motherboard RAM and the GPU (via the PCI lanes) is much slower. The key to achieving optimal overall computation speed is to either (1) keep the number of synchronisation points and volume of data transfers (between CPU and GPU) to a minimum, or (2) ensure that the GPU has work to do while data transfers are effected. The latter approach is feasible for once-only processing tasks such as manipulating a video stream, but where the task involves iteratively evolving a data set the former approach must be used. As much of the model as possible needs to reside in GPU memory keeping data exchanges with the CPU to a minimum.

The TUFLOW HPC engine currently only processes the 2D solution of a hydraulic model on GPU. The 1D component is calculated with the existing (legacy) ESTRY engine on CPU, requiring some synchronisation and data transfer between CPU and GPU compute processes. There are two types of connection depending on whether the 1D engine receives levels and computes flows or receives flows and computes levels. Either way, the computation steps are mutally exclusive and CPU and GPU processes must wait for the other to complete before proceeding.

It is therefore apparent that for models with substantial 1D networks, there is a performance gain to be realised by moving the 1D computation step onto the GPU. And further, to ensure the concurrent compute capability of the GPU is optimal, it is preferable to adopt an explicit finite volume formulation. Essentially, the objective is to create a 1D version of the existing 2D scheme used by TUFLOW HPC, adapting it as necessary for specific 1D element types.

METHODS

An explicit 1D finite volume scheme has been prototyped, first in Python, and then in C for CPU execution. The 1D network uses an offset grid for the conserved variables where volume is evolved at nodes (in-line points or junctions of 1D paths) and momentum is evolved at path mid-points (a path connects two nodes). This is illustrated in Figure 1.

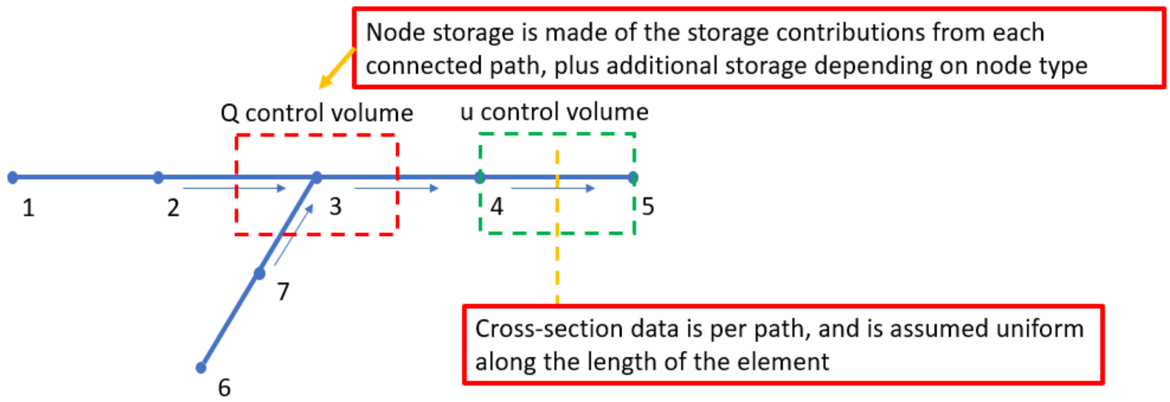

Figure 1. Offset volume and momentum control volumes.

The equations of motion are shown in Equation (1) for volume and Equation (2) for momentum:

$$\frac{\partial Q}{\partial t} = \sum \phi_i + S \tag{1}$$

where Q is the volume stored at a node, ϕ_i are the path fluxes into the node, and S is a local source/sink at the node.

$$\frac{\partial u}{\partial t} = \frac{\phi_1 u_1 - \phi_2 u_2 - u\frac{\partial Q}{\partial t}}{Q} - g\frac{\partial e}{\partial x} - g\frac{n^2 u^2}{R_h^{4/3}} - \frac{k}{L}\frac{u^2}{2} \tag{2}$$

where u is the area averaged flow velocity in the path, $\phi_1 u_1$ and $\phi_2 u_2$ are the momentum fluxes entering and leaving the u control volume at its upstream and downstream boundaries respectively (upstream/downstream are with respect to the digitisation direction, i.e. positive u), Q in this case is stored volume within the u control volume, $\partial e/\partial x$ the free surface slope (again with respect to digitisation direction), n the Mannings bed friction number, R_h the hydraulic radius of the path as defined by flow area divided by wetted perimeter, and finally k is an entry or internal energy loss coefficient and L the length of the path element. Exit losses are accounted for in the treatment of momentum transfer at nodes. Nodes that are simply an in-line point along a channel (and therefore have no storage in addition to the that of half of the upstream and downstream path elements) should transfer all of the exit momentum ($\phi_2 u_2$) from the upstream path as inlet momentum ($\phi_1 u_1$) for the downstream path, corresponding to a zero exit loss for the upstream path. However, the situation for path junctions is more complex – the exit loss for a particular path will depend on the relative angle between the upstream and downstream path, which can be dynamic if the there are two downstream paths. In the extreme case of manholes with large storage areas, the momentum transfer is near zero. Finally, path elements that represent a hydraulic structure utilise equations for flow based on the structure type, rather than evolving momentum.

The explicit nature of the equations of motion requires a compressible solution – i.e. the relationship between nodal depth and volume must be smooth and monotonically increasing. When path elements are closed, this requirement is no longer met once the paths run full. In this case, weak compressibility is achieved by allowing the water level at nodes to rise above the obverts of the connected paths. Commonly known as the ‘Preissmann Slot Method’, the nodal area for the storage extension is based on a small fraction of the path width times the path length.

The equations of motion are evolved using the standard 4th order Runge-Kutta integrator as is also used by TUFLOW HPC’s 2D solution. At each iteration:

1) The water levels at each node are computed as a diagnostic variable from the nodal volumes. As the relationship between level and volume at a node may not be easily inverted, a numerical procedure is used to solve for the node water level that yields the current node volume. Note that the nodal volume includes half of each connected path, and any additional nodal area relationship or data tables.
2) Once the nodal water levels are estabilished, path flow areas and volumes fluxes can be computed and subsequently the momentum fluxes at the ends of each path can be computed, along with the water surface slope for each path.
3) The bed friction and loss coefficient terms are solved for using a locally implicit calculation based on the full time step *dt*. The change in *u* within the step is then back computed to a time local derivative.
4) The nodal volume and path velocity derivatives are computed and passed to the Runge-Kutta integrator.

Concurrency was implemented using OpenMP compiler directives to parallelise all `for` loops that run over either nodes or path elements. Computed results were found to be identical between single and multi-threaded execution.

RESULTS AND DISCUSSION

Stage 1 testing, as presented in this paper, was performed using a selection of the English Environment Agency's (EA) 1D benchmark tests [EA 2004] as listed in Table 1. Tests A-D are mostly steady state cases, where the TUFLOW HPC 1D engine was run for a sufficient model time for the solution to reach equilibrium. Tests E and F are specifically transient flow test cases. As these tests are for open channels, Stage 2 testing (underway at the time of writing) will extend to other benchmarking specifically for pipe flow, manholes and hydraulic structures.

Table 1 Environment Agency Benchmark Cases

Test Case	Description
Test A	Straight uniform rectangular channel flow, subcritical, supercritical, transitions
Test B	Looped system, non-uniform flow split
Test C	Straight uniform triangular channel
Test D	Weirs
Test E	Ippen Wave (tidal response)
Test F	Monoclinic rising wave (propagating solitary wave flood front)

Test A: Uniform Rectangular Channel

Each of the Test A cases utilise a straight uniform rectangular channel of fixed Manning's friction with constant flow and a prescribed downstream water level, save for Part 6, which uses a time-varying downstream water level. Refer [EA 2004 Test A] for full details. Analytical solutions exist for the steady state cases: setting the left sides of Equations (1) and (2) to zero, it is possible to derive Equation (3) which expresses the spatial gradient of water depth h. For a given end boundary condition, this equation may be integrated forward or backwad in x to derive an exact solution, and therefore these benchmark tests form a comprehensive and exacting study. The HPC 1D results are compared against the analytical solutions for the steady state cases 1-5 in Figure 2. The depth percentage error is also shown on these figures. Mostly the errors are less than about 0.5% except for near hydraulic jumps where some amount of numerical dispersion smooths the jump slightly compared to the analytical answer. The locations of the hydraulic jumps are well predicted.

$$\frac{\partial h}{\partial x} = \frac{-\frac{\partial z}{\partial x} - \frac{n^2|u|u}{R_h^{4/3}}}{1 - \frac{u^2}{gh}} \tag{3}$$

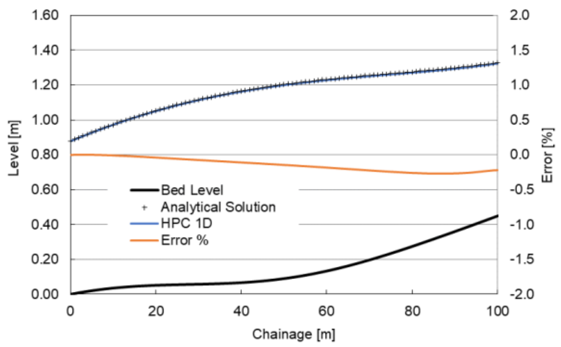

1 Subcritical

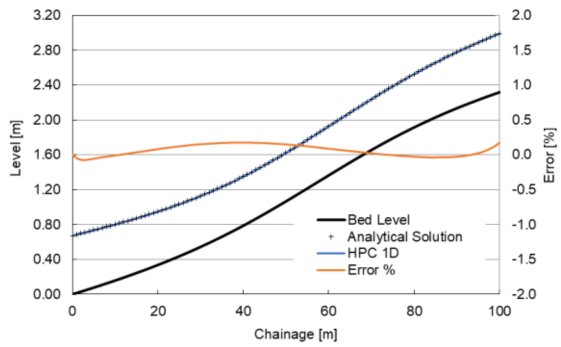

2 Supercritical

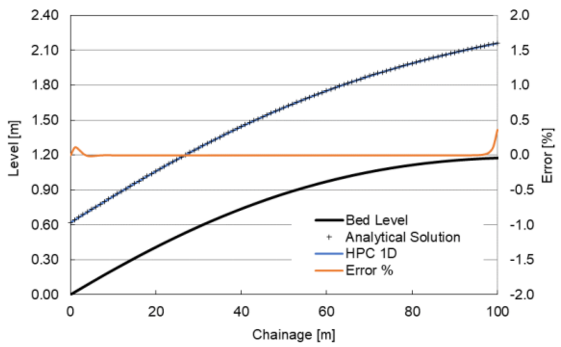

3 Subcritical to supercritical

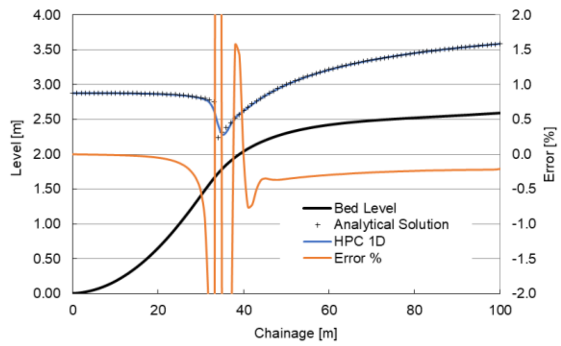

4 Supercritical to subcritical

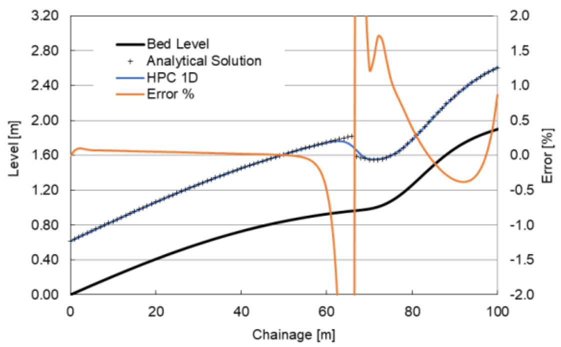

5 Subcritical to supercritical to subcritical

Figure 2. HPC 1D Results, EA Test A Parts 1-5

Test B: Looped System

The looped system test is comprised of a single flow path that branches into two and then recombines as shown in Figure 3, which as been reproduced from [EA 2004 Test B]. The details of the channel widths, gradients, Manning's numbers, and flow rates are also listed in this reference. There is no analytical solution for this test.

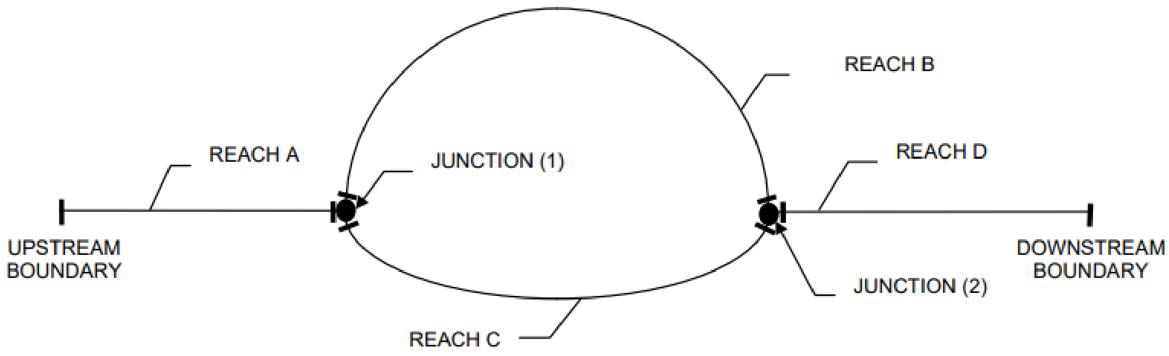

Figure 3. Looped System Test, reproduced from EA 2004 Test B

The results for the first steady state case are shown in Figure 4 a long with the published results for other software. Interestingly, the water levels (Figure 4, plot a) in Reach D (the most upstream) are predicted to be slightly lower than the levels at the upstream end of Reaches B and C. Initially this may seem odd, but when energy elevations are considered (Figure 4, plot b) the results appear entirely sensible. The velocity head in Reach A is significant and its energy is mostly recovered due to the conservation of momentum being applied across the junctions. Unlike some software, the HPC 1D scheme makes no attempt to force either water surface elevations or total energy elevations in connected paths to match at junctions – it returns to first principles and applies conservation of volume at each node and conserves momentum.

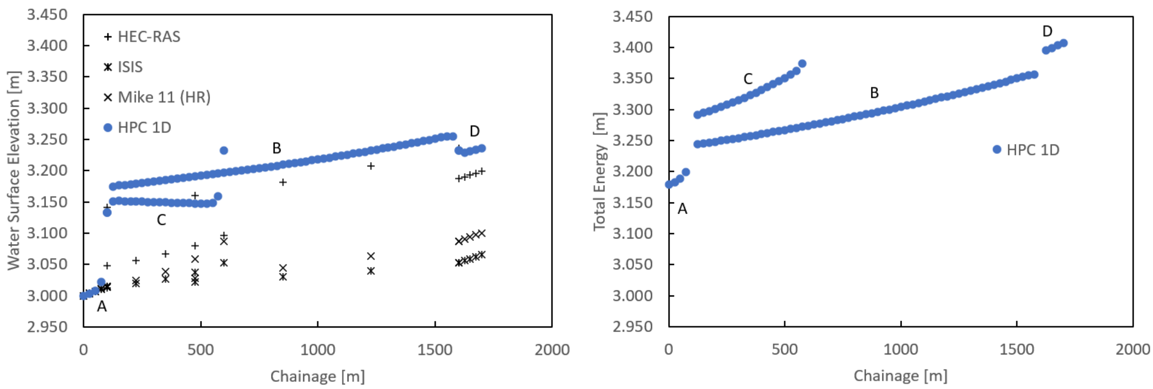

(a) Nodal water surface elevations (b) Nodal total energy elevations

Figure 4. Looped System Test, steady state case 1

Table 2 UK EA Test B results

Test Case	Downstream level (m)	Predicted flow split (Path B %)	U/S Water Surface level [m AD]
SS1	3.000	49.983	3.236
SS2	1.600	51.381	2.605

The predicted flow fraction for Path B and the upstream water surface elevations for both cases are listed in Table 2.There is some variation in the reported test results in [EA 2004 Test B] for the various software trialled. The results presented for HPC 1D agree well with those reported from the HEC-RAS steady state solutions.

Test C: Uniform Triangular Channel

An exact flow depth for triangular channels can be computed using Equation 4 [EA 2004 Test C]:

$$h = \left(\frac{4(B^2+1)}{B^2}\right)^{\frac{1}{8}} \left[\frac{Qn}{B\sqrt{S}}\right]^{\frac{3}{8}} \tag{4}$$

Where B is the reciprocal of the side slope (i.e. H/V), and S is the longitudinal slope. In this benchmark, B = 2, Mannings bed friction is 0.035, and the flow rate is 20 m^3/s. Two cases are used, one subcritical, the other supercritical. The tests were run in transient mode with the HPC 1D engine for sufficient time to establish the steady state depth at the upstream cross-section to 6 significant figures, with the results listed in Table 3. The discrepancies between the HPC 1D results and the theoretical depths are effectively zero.

Table 3 EA Test C cases and results

Test Case	Slope	Downstream level (m AD)	Normal Depth (m)	HPC 1D depth (m)	Error (%)
Part 1 Subcritical	0.001	3.000	3.01229	3.01223	0.0019
Part 2 Supercritical	0.020	1.700	1.71773	1.71772	0.0004

Test D Weirs

The HPC 1D engine allows for weir elements which provide a prescribed flow as a function of upstream and downstream water levels. The benchmark tests in [EA 2004 Test D] are for a rectangular broad-crested weir, and a triangular Crump weir, both of which are tested in the free flow and drowned flow states. Currently, HPC 1D only implements rectangular weirs using the same advanced weir equation as TUFLOW:

$$Q = \frac{2}{3} C_d C_{sf} W \sqrt{2g}\, H_u^{ex} \tag{5}$$

where W is the width, H_u is the upstream elevation above the crest of the weir, ex the exponent, C_d the discharge coefficient. C_{sf} is the submergence factor, which is implemented with the Villemont equation:

$$C_{sf} = \left[1 - \left(\frac{H_d}{H_u}\right)^a\right]^b \tag{6}$$

where H_u and H_d are the upstream and downstream elevations above the crest of the weir – note that if the downstream elevation is below the crest of the weir then the weir is operating in free flow mode and the submergence factor defaults to 1. Also note Equations 5 and 6 use total energy elevation for the upstream H_u, and water surface elevation for the downstream H_d.

The benchmark cases for the rectangular broad-crested weir were implemented using a width 0.9 m, discharge coefficient 0.577, exponent 1.5, and a = 8.55, b = 0.556 for the Villemont submergence factor relation. The upstream inflow was set at 0.150 m^3/s. The resulting water levels either side of the weir are listed in Table 4, along with the theoretical weir flow based on these levels. The comparison shows effectively zero error.

Table 4 EA Test D cases and results

Test Case	Downstream level (m)	Weir U/S Elev. (m)	Weir D/S Elev. (m)	Weir Equation Flow (m^3/s)	Error (%)
Part 1 Free Flow	0.300	0.708479	0.346688	0.150007	0.0044
Part 1 Drowned Flow	0.800	0.816145	0.804118	0.150011	0.0073

Test E Ippen Wave

Benchmark tests E and F assess the transient response of the solution scheme. The Ippen Wave [EA 2004 Test E] uses a uniform rectangular channel with one end closed and the other forced with a time-varying water surface elevation. If the nonlinear bed friction is linearised, and the advective momentum term is neglected, then an analytical solution to the transient forcing can be found (the solution is detailed in the test documentation).

The case was constructed as per the test documentation and run with the HPC 1D engine. The time-series results for water level at x = -75 km are shown in Figure 5 along with the analytical solution and the published results for HEC-RAS and Mike 11. All of the numerical engine results deviate from the analytical solution, however HPC 1D and Mike 11 appear to capture a shorter wavelength wave in the solution. One possible explanation is that HPC 1D and Mike 11 are correctly implementing nonlinear bed friction and the advective momentum term that the analytical solution has neglected.

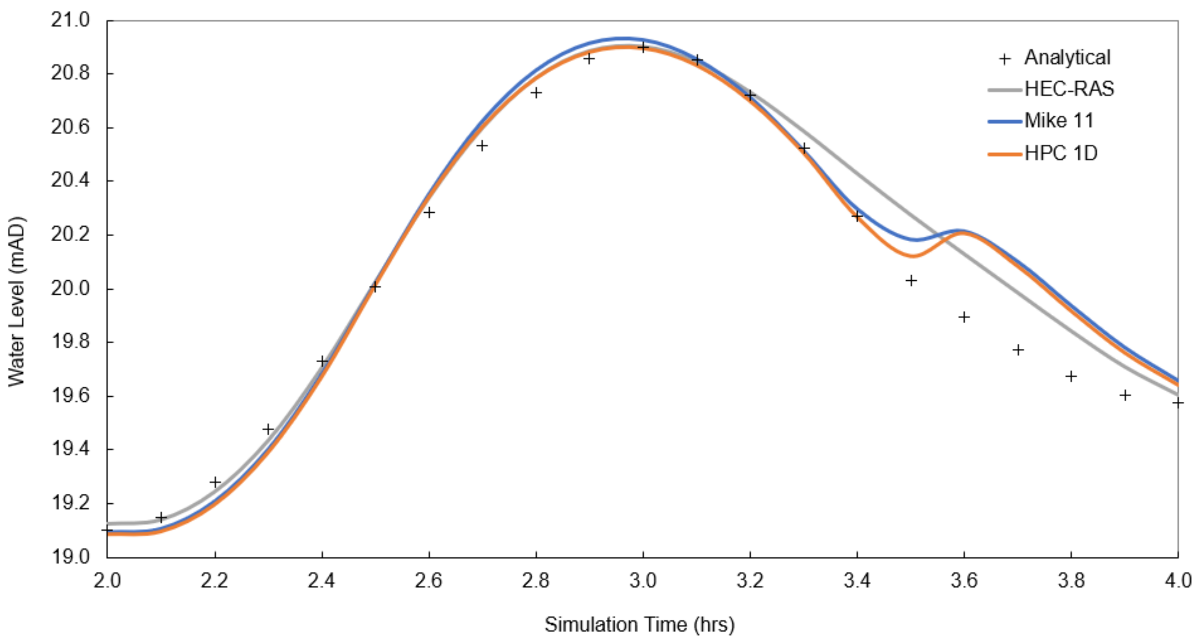

Figure 5. Ippen Wave Water Levels at -75 km

Test F Monclinic Rising Wave

The final benchmark test case presented is that of the monoclinic rising wave. When bed friction is assumed to follow the Chezy friction formulation, solitary wave solutions exist for a rising flood wave that preserves its shape as it propagates. The benchmark model was built as detailed in [EA 2004 Test F], with a depth-dependent Manning's bed friction curve created to represent a Chézy friction coefficient of 55. The results for the flood wave are shown in Figure 6 at 5 hr intervals (propagating from right to left). The analytical solutions are shown with crosses and the HPC 1D solutions are shown with lines. The HPC 1D solution propagates at the correct velocity and its shape is well preserved.

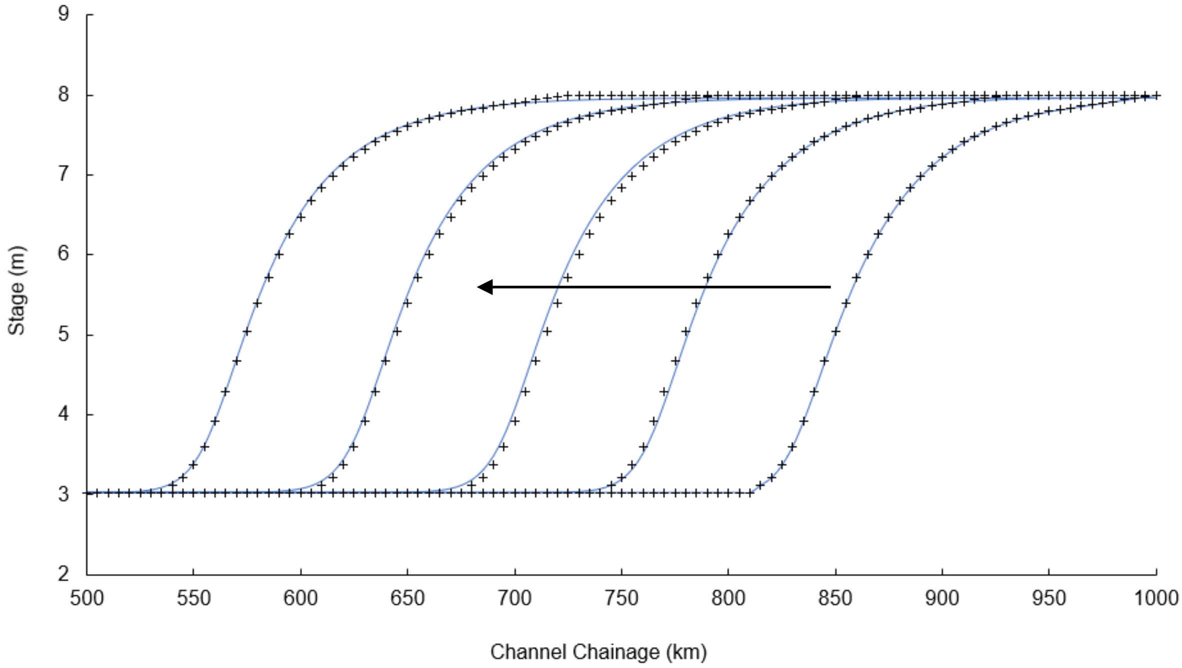

Figure 6. Monoclinic Rising Wave (analytical solution as markers, HPC 1D solution as lines)

Stability

Where equations of motion admit a wave solution and have advective fluxes, the normal stability criterion for a transient numerical solution is shown in Equation 7:

$$\frac{(C + |u|)dt}{dx} \leq CFL_{lim} \tag{7}$$

where C is the wave speed, or celerity, and dx the length of the path element. For path elements of arbitrary cross-section, the wave celerity is given by:

$$C = \sqrt{\frac{gA_f}{W_s}} \tag{8}$$

where A_f is the cross-sectional area of the water flowing within the element and W_s the width of the free surface. Note that when pipe sections begin to run more than half full (and particularly when the flow is in the slot extension), the wave speed can become more significant. HPC 1D uses an adaptive timestep such that the Courant Friedrichs Lewy (CFL) control number is maintained at or below an appropriate limit. The test cases presented all ran stably with a CFL limit of 1.6.

CONCLUSION

An explicit finite volume solution scheme for unsteady 1D flow in open or closed channels of arbitrary cross section has been developed and protyped. The method, called TUFLOW HPC 1D, is based on the existing 2D scheme employed by TUFLOW HPC.

The method is unconditionally stable subject to the standard Courant and celerity criteria typical for such schemes, and is parallelised for multi-threaded CPU and GPU hardware. It is easily offloaded to GPU where it can access TUFLOW HPC's 2D domain data without device synchronisation or memory transfers, further increasing the computational efficiency of 1D/2D coupled models.

A number of the EA 1D hydraulic scheme benchmarks have been constructed and run with HPC 1D. Where analytical solutions are available the scheme has performed flawlessly. For the remaining tests where analytical solutions are not available, the results appear sensible and agree well, or are an improvement, on the other 1D solutions benchmarked.

Reference List

cpubenchmark "https://www.cpubenchmark.net/year-on-year.html"

TUFLOW Hardware Benchmarks "https://wiki.tuflow.com/Hardware_Benchmarking_-_Results#CPU_Results"

Collecutt and Syme 2017 "Experimental Benchmarking of Mesh Size and Time-Step Convergence for a 1st and 2nd Order SWE Finite Volume Scheme" *Proceedings of the 37th IAHR World Congress* pp 5141-5150

Gao et. al. 2022 The End of the 1D Open Channel Cross-Section? *HWRS 2022*

Huxley et. al. 2022 Hydraulic Modelling 2D Cell Size Result Convergence – Comparing the Performance of Different Shallow Water Equation Solution Schemes *HWRS 2022*

EA 2004 "R&D Technical Report: W5-105/TR0" https://assets.publishing.service.gov.uk/media/602ac638e90e07055499e80a/_Benchmarking_of_Hydraulic_River_Modelling_Software_Packages_-_Overview.pdf

EA 2004 Test A "R&D Technical Report: W5-105/TR2A" https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/290889/scho0305bixp-e-e.pdf

EA 2004 Test B "R&D Technical Report: W5-105/TR2B" https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/290890/scho0305bixq-e-e.pdf

EA 2004 Test C "R&D Technical Report: W5-105/TR2C" https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/290891/scho0305bixr-e-e.pdf

EA 2004 Test D "R&D Technical Report: W5-105/TR2D" https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/290892/scho0305bixs-e-e.pdf

EA 2004 Test E "R&D Technical Report: W5-105/TR2E" https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/290893/scho0305bixt-e-e.pdf

EA 2004 Test F "R&D Technical Report: W5-105/TR2F" https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/290894/scho0305bixu-e-e.pdf

BIOGRAPHY

Greg Collecutt is the principal GPU software developer at TUFLOW. He has degrees in mechanical engineering and a PhD in theoretical physics, and has spent most of the last twenty years working in computational fluid dynamics and flood modelling. In addition to working with the TUFLOW team, Greg has also worked for the Australian Bureau of Meteorology, in the solar thermal power industry and in the aerospace industry. Greg's core passion is for scientific coding and pushing the boundaries of computing technology, and has found a niche within the team at TUFLOW. In this role he is heavily involved with the implementation and benchmarking of new modelling features.