MIKE 21 Flow Model FM

Parallelisation using multi-GPU

Benchmarking report

DHI headquarters
Agern Allé 5
DK-2970 Hørsholm
Denmark

+45 4516 9200 Telephone
+45 4516 9333 Support
+45 4516 9292 Telefax

mikebydhi@dhigroup.com
www.mikebydhi.com

PLEASE NOTE

PRINTING HISTORY

July 2014 .... Release 2014

CONTENTS

MIKE 21 Flow Model FM
Parallelisation using multi-GPU
Benchmarking report

1 Vision and Scope

A set of well-defined test cases for the GPU version of the MIKE 21 Flow Model FM, have been established. These test cases also cover MIKE FLOOD using the GPU version of the MIKE 21 Flow Model FM for the 2D surface flow calculation. The test-suite is used to test the performance across platforms with different graphics cards. It is essential that it is possible to run the simulation with different spatial resolutions to be able to evaluate the scalability of the parallelisation. The main focus is to benchmark the GPU parallelisation of the flexible mesh modelling system.

2 Methodology

2.1 GPU Parallelisation

The GPU computing approach uses the computers graphics card to perform the computational intensive calculations. This approach is based on CUDA™ by NVIDIA and can be executed on NVIDIA graphics cards with Compute Capability 2.0 or higher.

Depending on the available hardware it is possible to launch a simulation using a single or multiple GPUs. The multiple GPU approach is based on the domain decomposition concept, where the communication between the processors is done using MPI (Message Passing Interface).

Currently, only the computational intensive hydrodynamic calculations are performed on the GPU. The additional calculations are for each sub-domain in the domain decomposition performed locally on the CPU and these calculations are further parallelised based on the shared memory approach, OpenMP.

As default, the program uses one MPI process per GPU, but it is possible to assign more processes to the same GPU. In this way simulations, where the hydrodynamic calculations are less time consuming than the calculations performed in the other modules, will benefit from the MPI parallelisation.

2.2 Hardware

The benchmarks have been performed using the following hardware platforms and GPUs:

Table 2.1 Hardware platforms used for benchmarking

	Computer	Processor	Memory	Operating system	GPUs
1	DELL Precision T7610 (workstation)	2 x Intel®Xeon® E5-2687W v2 (8 cores, 3.40 GHz)	32 GB	Windows 7 Professional SP1, 64-bit	2 x GeForce GTX Titan
2	Anselm Cluster (209 nodes and 23 of these withh GPU, infiniband)	2 x Intel Sandy Bridge E5-2470, 8-core, 2.3 GHz processors per node	96 GB of physical memory per node	Linux	1 x Tesla K20m per node

The Anselm Cluster Is located at IT4Innnovations National Supercomputing Center, Ostrava.

Table 2.2 GPU specifications

GPU	Compute Capability	Number of CUDA cores	Memory (GB)	Single/Double precision floating point performance
GeForce GTX Titan	3.5	2688	6	4.5 Tflops / 1.27 Tflops
Tesla K20m	3.5	2496	5	3.52 Tflops / 1.17 Tflops

2.3 Software

All benchmarks have been performed using the Release 2014, Service Pack 2, version of the MIKE by DHI software.

2.4 Performance of the GPU Parallelisation

The parallel performance is illustrated by measuring the speedup factor. The speedup factor is defined as the elapsed time using the existing CPU version of MIKE 21 Flow Model FM (one core/thread) divided by the elapsed time using the new GPU version (one core/thread for the CPU part of the calculation). The elapsed time is the total elapsed time (excluding pre- and post-processing). The performance metric is highly dependent on not only the GPU hardware but also the CPU hardware.

For the GPU simulations the number of threads per block on the GPU is 128.

3 Description of Test Cases

3.1 Mediterranean Sea

This test case has been established for benchmarking of the MIKE 21 Flow model FM.

3.1.1 Description

In the Western parts of the Mediterranean Sea tides are dominated by the Atlantic tides entering through the Strait of Gibraltar, while the tides in the Eastern parts are dominated by astronomical tides, forced directly by the Earth-Moon-Sun interaction.

3.1.2 Setup

The bathymetry is shown in Figure 3.1. Simulations are performed using five meshes with different resolution (see Table 3.1). The meshes are generated specifying the value for the maximum area of 0.04, 0.005, 0.00125, 0.0003125 and 0.000078125 degree^2, respectively. The simulation period for the benchmarks covers 2 days starting 1 January 2004 for the simulations using mesh A, B and C. The simulation period is reduced to 6 hours for the simulations using mesh D and E.

At the Atlantic boundary a time varying level boundary is applied. The tidal elevation data is based on global tidal analysis (Andersen, 1995).

For the bed resistance the Manning formulation is used with a Manning number of 32. For the eddy viscosity the Smagorinsky formulation is used with a Smagorinsky factor of 1.5. Tidal potential is applied with 11 components (default values).

The shallow water equations are solved using both the first-order scheme and the higher-order scheme in time and space.

Table 3.1 Computational mesh for the Mediterranean Sea case

Mesh	Element shape	Elements	Nodes	Max. area Degree^2
Mesh A	Triangular	11287	6283	0.04
Mesh B	Triangular	80968	41825	0.005
Mesh C	Triangular	323029	164161	0.00125
Mesh D	Triangular	1292116	651375	0.0003125
Mesh E	Triangular	5156238	2588665	0.000078125

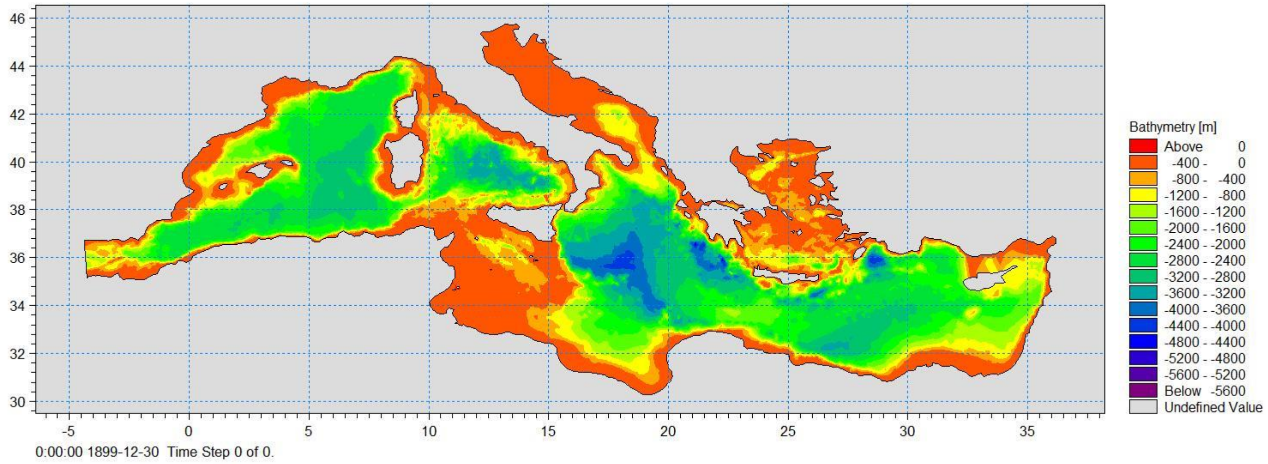

Figure 3.1 Bathymetry for the Mediterranean Sea case

The averaged time step for the simulations using Mesh A, B, C, D and E is 17.65s, 5.61s, 2.86s, 1.43s and 0.69s, respectively, for both the first-order scheme and the higher-order scheme in time and space.

3.1.3 Data and specification files

The data and specification files used for the present case are located in the directory:

Benchmarking\Mediterranean_Sea

The tests are performed using the following specification files:

2004_tide_0.04_1st.m21fm

2004_tide_0.005_1st.m21fm

2004_tide_0.00125_1st.m21fm

2004_tide_0.0003125_1st.m21fm

2004_tide_0.00078125_1st.m21fm

2004_tide_0.04_2nd.m21fm

2004_tide_0.005_2nd.m21fm

2004_tide_0.00125_2nd.m21fm

2004_tide_0.0003125_2nd.m21fm

2004_tide_0.00078125_2nd.m21fm

3.2 Ribe Polder

This test case has been established for benchmarking of the MIKE 21 Flow Model FM.

3.2.1 Description

The model area is located, on the southern part of Jutland, Denmark, around the city of Ribe. The area is protected from storm floods in the Wadden Sea to the west by a dike. The water course Ribe Å runs through the area and crosses the dike through a lock.

The flood condition where the dike is breached during a storm flood is illustrated by numerical modelling. The concept applied to model the breach failure in the hydrodynamic model is based on prescribing the breach by a dynamic bathymetry that change in accordance with the relation applied for the temporal development of the breach. Use of this method requires that the location of the breach is defined and known at an early stage, so that it can be resolved properly and built into the bathymetry. The shape and temporal development of the breach is defined with a time-varying distribution along the dike crest. It is further defined how far normal to the crest line the breach can be felt. Within this distance the bathymetry is following the level of the breach, if the local level is lower than the breach level no changes are introduced. The area of influence of the breach will therefore increase with time.

The breach and flood modelling has been carried out based on a historical high water event (24 November, 1981), shown in Figure 3.2. Characteristic for this event is that high tide occurs at the same time as the extreme water level. Højer sluice is located about 40 km south of the breach, while Esbjerg is located about 20 km to the north. Based on the high water statistics for Ribe the extreme high water level has been estimated for an event having a return period of 10,000 years. The observed water level at Højer is hereafter adjusted gradually over two tidal cycles to the extreme high water level estimated for the given return periods at Ribe, as indicated in Figure 3.2. The water level time series established in this way are shown in Figure 3.2.

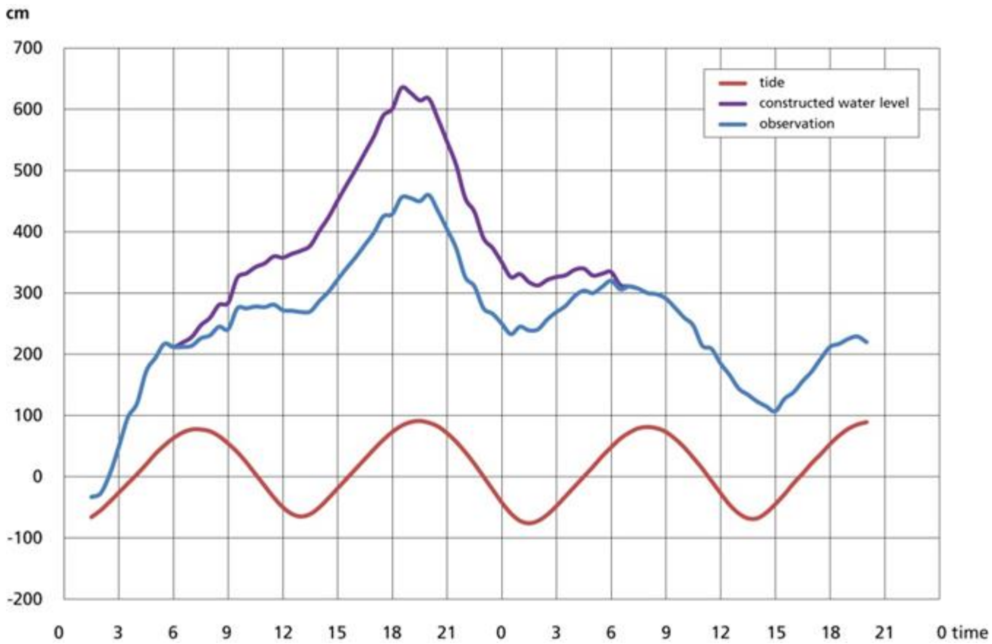

Figure 3.2 Runoff from the catchment is included as specified discharges given for the two streams Ribe Å and Kongeåen

The crossing between the dike and Ribe Å is shown in Figure 3.4. The crossing is in the form of a navigational chamber lock. It is represented in the model bathymetry as a culvert that can be closed by a gate. The points defining the dike next to the creek are modified to have increased levels in order to ensure a well-defined bathymetry where flow only occurs through the cells defining the creek proper. The sluice is defined as a check valve allowing only flow towards the sea.

A constant discharge of 9.384 m^3/s and 14.604 m^3/s, respectively, are applied for the two streams Ribe Å and Kongeåen. For the bed resistance the Manning formulation is used with a Manning number of 18. For the eddy viscosity the Smagorinsky formulation is used with a Smagorinsky factor of 0.28.

3.2.2 Setup

The bathymetry is shown in Figure 3.3. The computational mesh contains 173101 elements. A satisfactory resolution of the breach is obtained by a fine mesh of structured triangles and rectangles as shown in Figure 3.4. The areas in- and offshore of the dike is defined by a relatively fine mesh to avoid instabilities due to humps or holes caused by large elements with centroids just outside the area of influence from the breach. The simulation period is 42 hours.

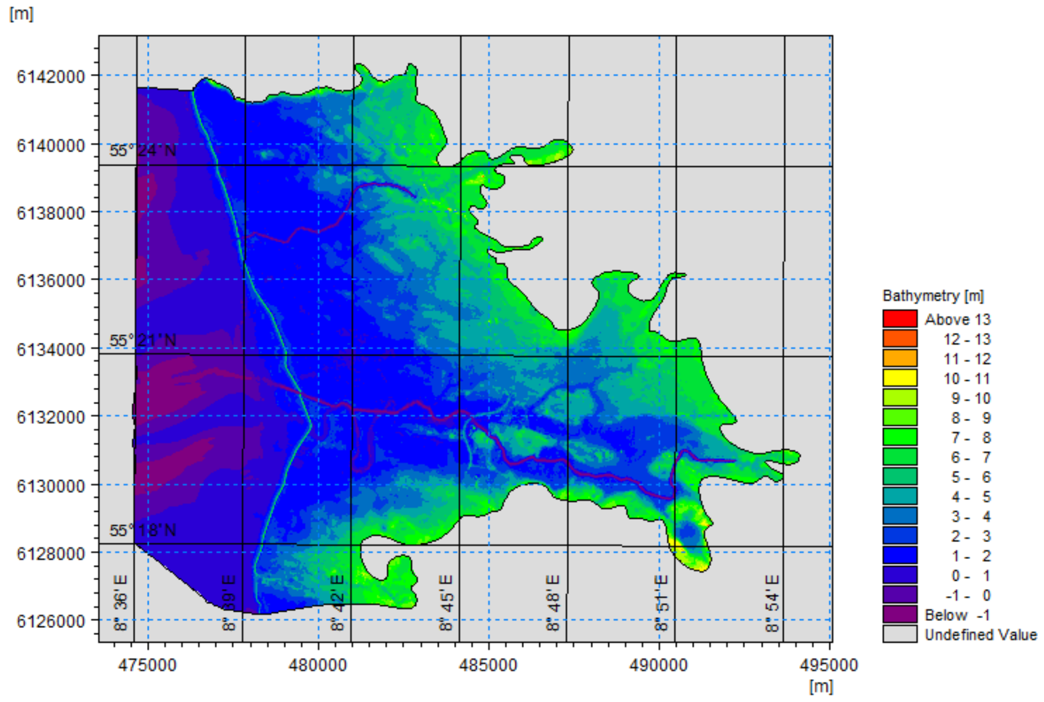

Figure 3.3 Bathymetry for the Ribe Polder case

The shallow water equations are solved using both the first-order scheme and higher-order scheme in space and time.

The averaged time step is 0.21s for both the first-order scheme and higher-order scheme in space and time.

3.2.3 Data and specification files

The data and specification files used for the present case are located in the directory Benchmarking\Ribe_Polder

The tests are performed using the following specification files:

Event_10000_1st.m21fm

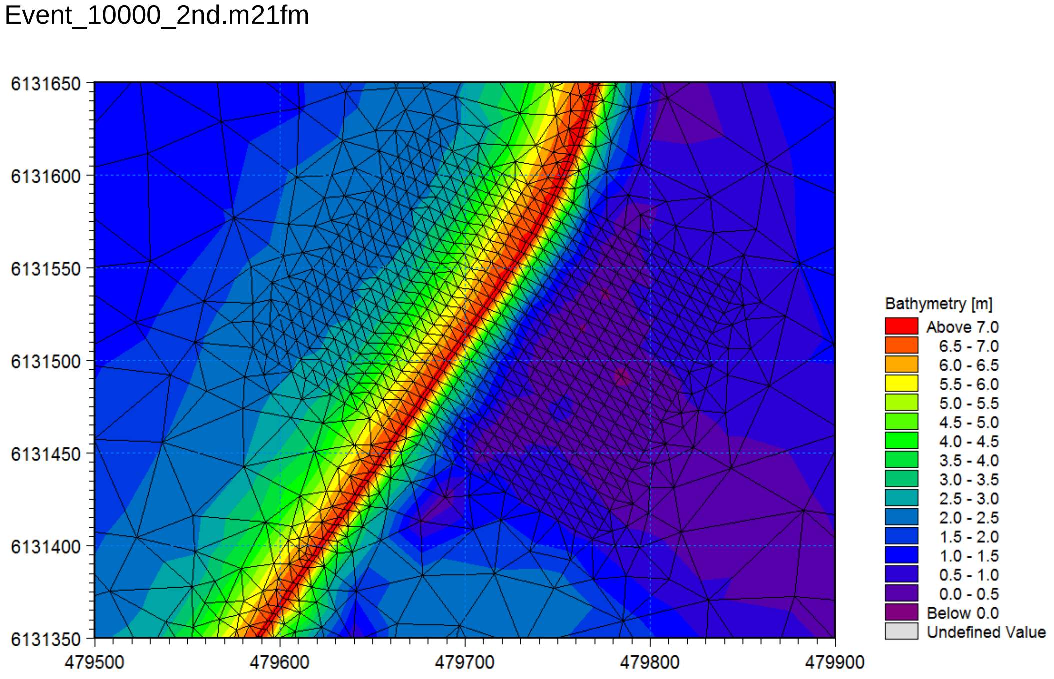

Figure 3.4 Close-up of the bathymetry

At the offshore boundary a time series of level variations is applied.

3.3 EA2D Test 8A

This test is Test 8A in the benchmarks test developed during the Joint Defra/Environment Agency research programme. This tests the package's capability to simulate shallow inundation originating from a point source and from rainfall applied directly to the model grid, at a relatively high resolution. This test case has been established for benchmarking of the MIKE 21 Flow model FM.

3.3.1 Description

The modelled area is approximately 0.4 km by 0.96 km and covers entirely the DEM provided and shown in Figure 3.5. Ground elevations span a range of ~21m to ~37m.

The flood is assumed to arise from two sources:

- a uniformly distributed rainfall event illustrated by the hyetograph in Figure 3.6. This is applied to the modelled area only (the rest of the catchment is ignored).
- a point source at the location (264896, 664747) (Map projection: British national grid), and illustrated by the inflow time series in Figure 3.7. (This may for example be assumed to arise from a surcharging culvert.)

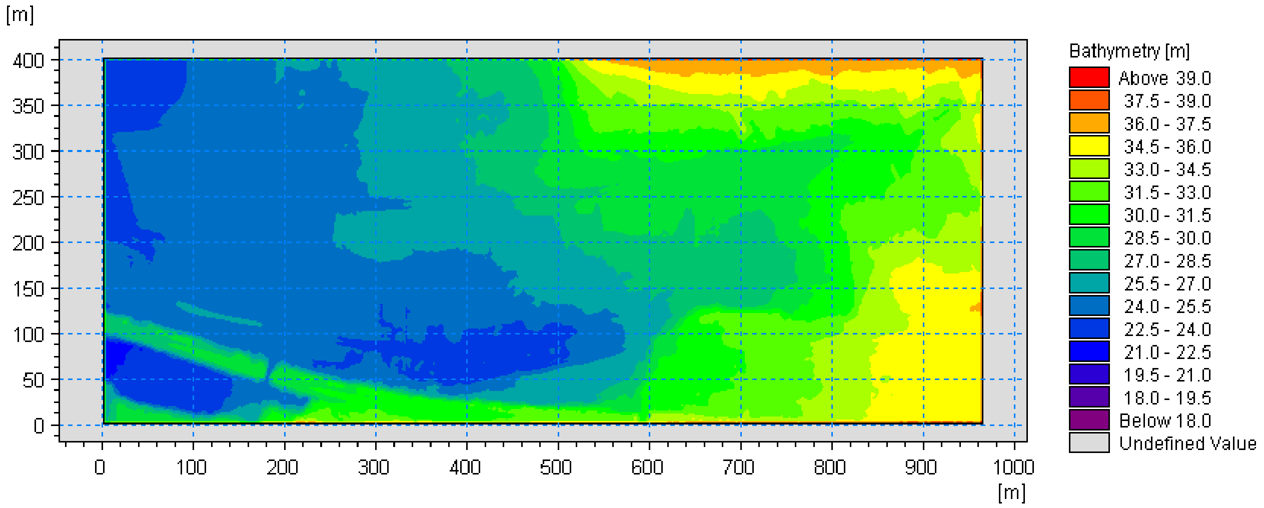

Figure 3.5 Bathymetry for the EA2D Test8A case

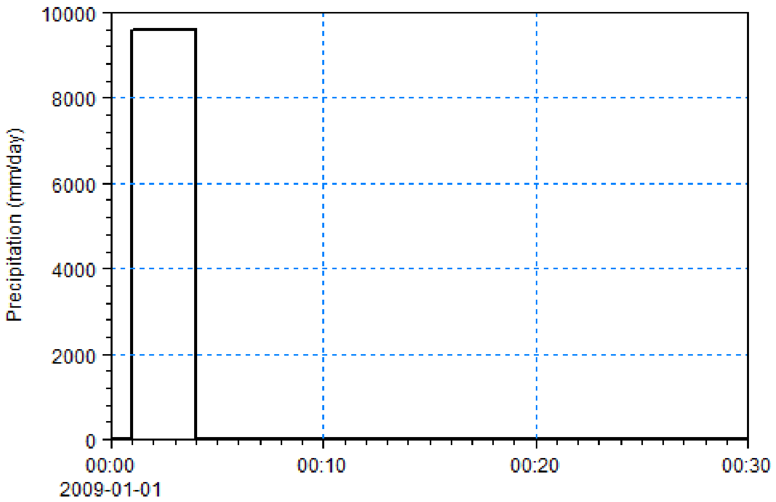

Figure 3.6 Rainfall

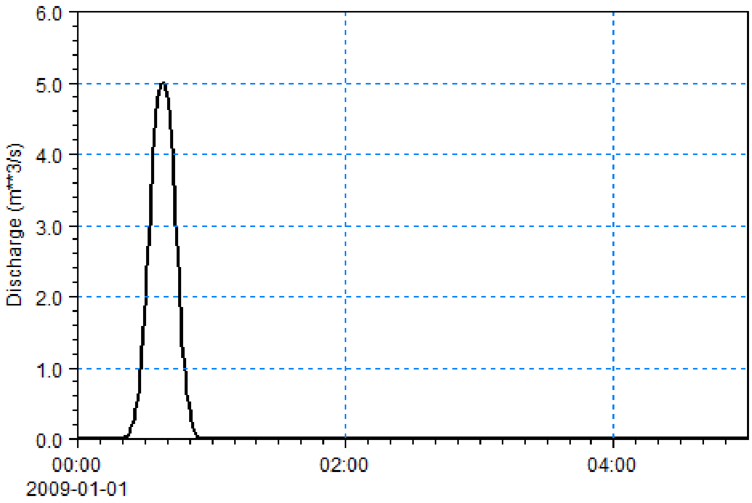

Figure 3.7 Discharge from source

DEM is a 0.5m resolution Digital Terrain Model (no vegetation or buildings) created from LiDAR data collected on 13th August 2009 and provided by the Environment Agency (http://www.geomatics-group.co.uk). Model grid resolution should be 2m (or ~97000 nodes in the 0.388 km^2 area modelled).

All buildings at the real location (Cockenzie Street and surrounding streets in Glasgow, UK) are ignored and the modelling is carried out using the “bare-earth” DEM provided.

A land-cover dependent roughness value is applied, with 2 categories: 1) Roads and pavements; 2) Any other land cover type. Manning’s n = 0.02 is applied for roads and pavements n = 0.05 everywhere else.

All boundaries in the model area are closed (no flow) and the initial condition is dry bed. The model is run until time T = 5 hours to allow the flood to settle in the lower parts of the modelled domain.

3.3.2 Setup

Simulations are performed using four meshes with different resolution (see Table 3.2). The four meshes uses regular quadrilateral elements with grid spacing 2m, 1m, 0.5m and 0.25m, respectively. Mesh A corresponds to the original mesh used in the EA2D test, and the additional meshes are obtained by refining this mesh.

Table 3.2 Computational mesh for the EA2D Test 8A case

Mesh	Element shape	Elements	Nodes	Grid spacing metres
Mesh A	Quadrilateral	95719	96400	2
Mesh B	Quadrilateral	384237	385600	1
Mesh C	Quadrilateral	1539673	1542400	0.5
Mesh D	Quadrilateral	6164145	6169600	0.25

The shallow water equations are solved using the first-order scheme in time and space.

The averaged time step for the simulation using Mesh A, B, C and D is 0.22s, 0.10s, 0.5s and 0.027s, respectively.

3.3.3 Data and specification files

The data and specification files used for the present case are located in the directory:

Benchmarking\EA2D_Test_8A

The tests are performed using the following specification files:

Test8A_quadratic_2m.m21fm

Test8A_quadratic_1m.m21fm

Test8A_quadratic_0.5m.m21fm

Test8A_quadratic_0.25m.m21fm

3.4 EA2D Test 8B

This test is Test 8B in the benchmarks test developed during the Joint Defra/Environment Agency research programme. This tests the package’s capability to simulate shallow inundation originating from a surcharging underground pipe, at relatively high resolution. This test case has been established for benchmarking of the MIKE Flood using MIKE 21 Flow model FM for the 2d surface flow calculation and MIKE Urban for calculation of the pipe flow.

3.4.1 Description

The modelled area is approximately 0.4 km by 0.96 km and covers entirely the DEM provided and shown in Figure 3.8. Ground elevations span a range of ~21m to ~37m.

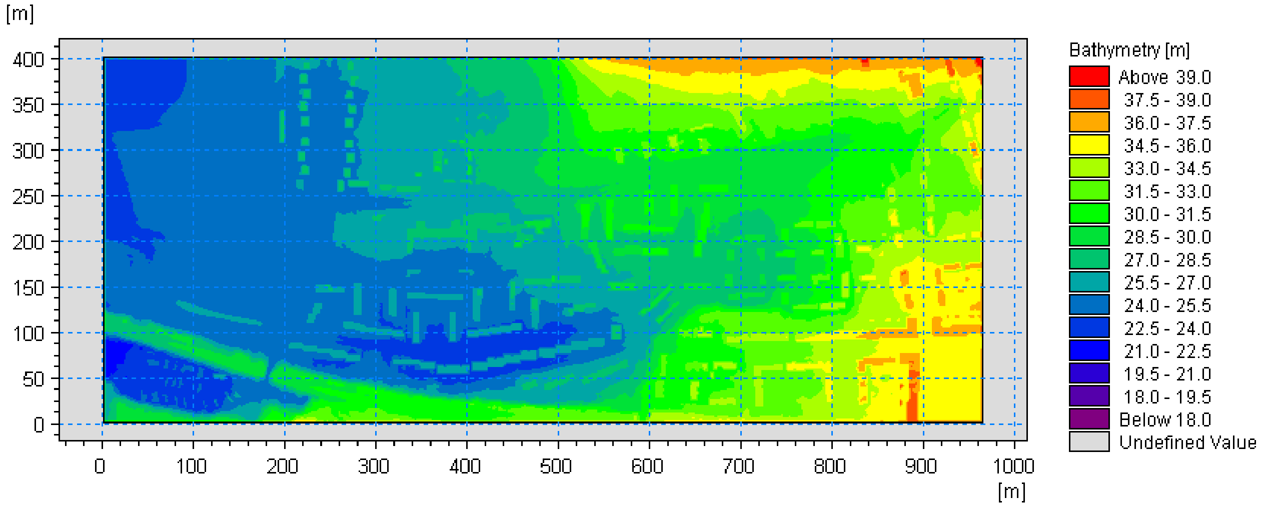

Figure 3.8 Bathymetry for the EA2D Test8B case

A culverted watercourse of circular section, 1400mm in diameter, ~1070m in length, and with invert level uniformly 2m below ground is assumed to run through the modelled area. An inflow boundary condition is applied at the upstream end of the pipe, illustrated in Figure 3.9. A surcharge is expected to occur at a vertical manhole of 1m^2 cross-section located 467m from the top end of the culvert, and at the location (264896, 664747). For the downstream boundary condition free out fall (critical flow is assumed). The base flow (uniform initial condition) is 1.6 m/s. The manhole is connected to the grid in one point and the surface flow is assumed not to affect the manhole outflow.

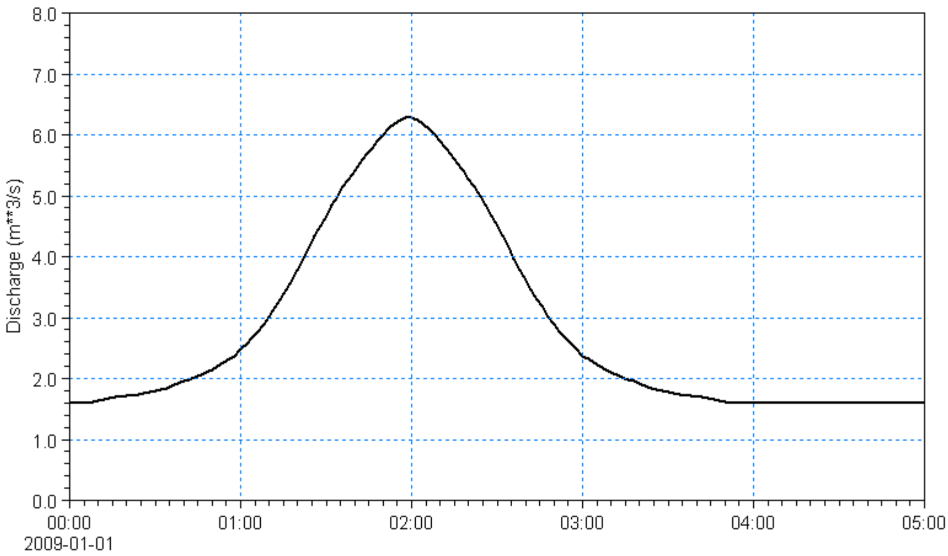

Figure 3.9 Inflow hydrograph applied for the EA2D Test8B at upstream end of culvert

DEM is a 0.5m resolution Digital Terrain Model (no vegetation or buildings) created from LiDAR data collected on 13th August 2009 and provided by the Environment Agency (http://www.geomatics-group.co.uk). Model grid resolution should be 2m (or ~97000 nodes in the 0.388 km^2 area modelled).

The presence of a large number of buildings in the modelled area is taken into account. Building outlines are provided with the dataset. Roof elevations are not provided.

A land-cover dependent roughness value is applied, with 2 categories: 1) Roads and pavements; 2) Any other land cover type. Manning's n = 0.02 is applied for roads and pavements n = 0.05 everywhere else.

All boundaries in the model area are closed (no flow) and the initial condition is dry bed. The model is run until time T = 5 hours to allow the flood to settle in the lower parts of the modelled domain.

3.4.2 Setup

Simulations are performed using four meshes with different resolution (see Table 3.3). The four meshes uses regular quadrilateral elements with grid spacing 2m, 1m, 0.5m and 0.25m, respectively. Mesh A corresponds to the original mesh used in the EA2D test, and the additional meshes are obtained by refining this mesh.

Table 3.3 Computational mesh for the EA2D Test 8B case

Mesh	Element shape	Elements	Nodes	Grid spacing metres
Mesh A	Quadrilateral	95719	96400	2
Mesh B	Quadrilateral	384237	385600	1
Mesh C	Quadrilateral	1539673	1542400	0.5
Mesh D	Quadrilateral	6164145	6169600	0.25

The shallow water equations are solved using the first-order scheme in time and space.

The averaged time step for the simulation using Mesh A, B, C and D is 0.27s, 0.15s, 0.76s and 0.025s, respectively.

3.4.3 Data and specification files

The data and specification files used for the present case are located in the directory:

Benchmarking\EA2D_Test_8B

The tests are performed using the following specification files:

Test8B_quadratic_2m.couple
Test8B_quadratic_1m.couple
Test8B_quadratic_0.5m.couple
Test8B_quadratic_0.25m.couple

4 Benchmarking using a GeForce GTX Titan Card

These tests have been performed using both a single GeForce GTX Titan graphics card and two GeForce GTX Titan graphics card on hardware platform 1 specified in Table 2.1.

4.1 Mediterranean Sea

Table 4.1 Simulations are carried out using single precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	64.19	11.24	5.71	13.51	4.75 (0.83)
Mesh B	1454.61	62.55	23.25	53.94	26.97 (1.15)
Mesh C	16814.01	269.68	62.34	193.53	86.88 (1.39)
Mesh D	13235.29	177.43	74.59	103.25	126.18 (1.71)
Mesh E		1256.99		676.05	(1.85)

Table 4.2 Simulations are carried out using single precision and using the higher-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	167.00	21.28	7.84	27.33	6.11 (0.77)
Mesh B	4325.12	157.33	27.49	133.18	32.47 (1.18)
Mesh C	46896.29	834.25	56.21	539.65	86.90 (1.54)
Mesh D	39764.76	648.09	61.35	363.18	109.49 (1.78)
Mesh E		5062.73		2680.93	(1.88)

Table 4.3 Simulations are carried out using double precision and using the first-order scheme in time and space.

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	64.19	12.07	5.31	13.96	4.59 (0.86)
Mesh B	1454.61	82.22	17.69	60.93	23.87 (1.34)
Mesh C	16814.01	337.21	49.86	227.56	73.88 (1.48)
Mesh D	13235.29	241.49	54.80	135.63	97.58 (1.78)
Mesh E		1775.31		949.72	(1.86)

Table 4.4 Simulations are carried out using double precision and using the higher-order scheme in time and space.

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	167.00	23.48	7.11	28.25	5.91 (0.83)
Mesh B	4325.12	202.00	21.41	151.25	28.59 (1.19)
Mesh C	46896.29	1017.99	46.06	648.47	72.31 (1.56)
Mesh D	39764.76	831.94	47.79	456.76	87.05 (1.82)
Mesh E		6930.60		3553.71	(1.95)

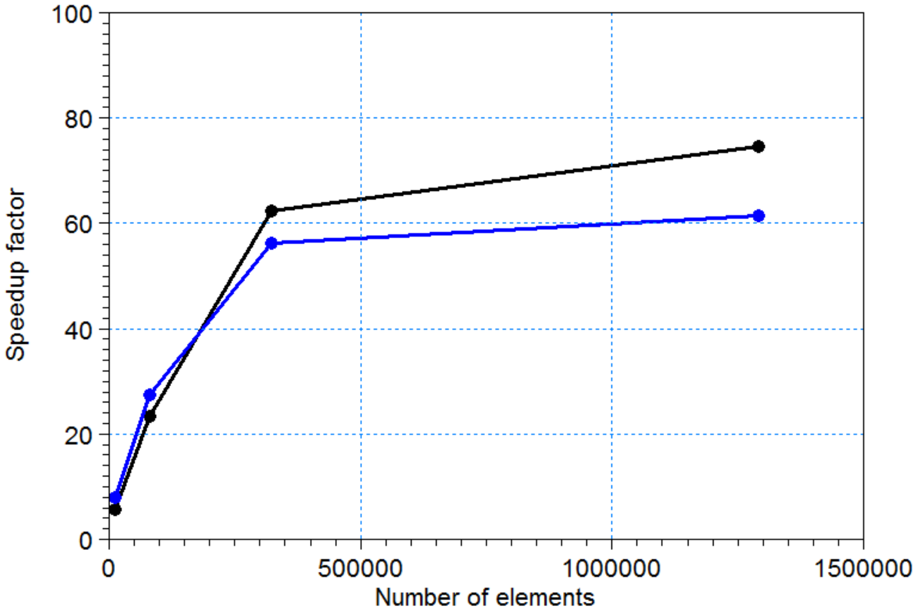

Figure 4.1 Speedup factor for one GeForce GTX Titan card using single precision. Black line: first-order scheme; Blue line: higher-order scheme

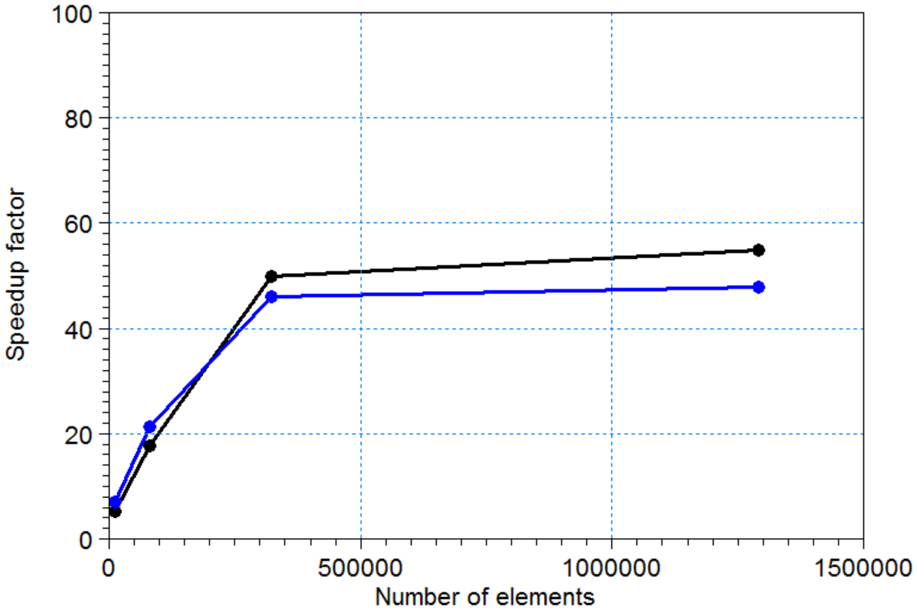

Figure 4.2 Speedup factor for one GeForce GTX Titan card using double precision. Black line: first-order scheme; Blue line: higher-order scheme

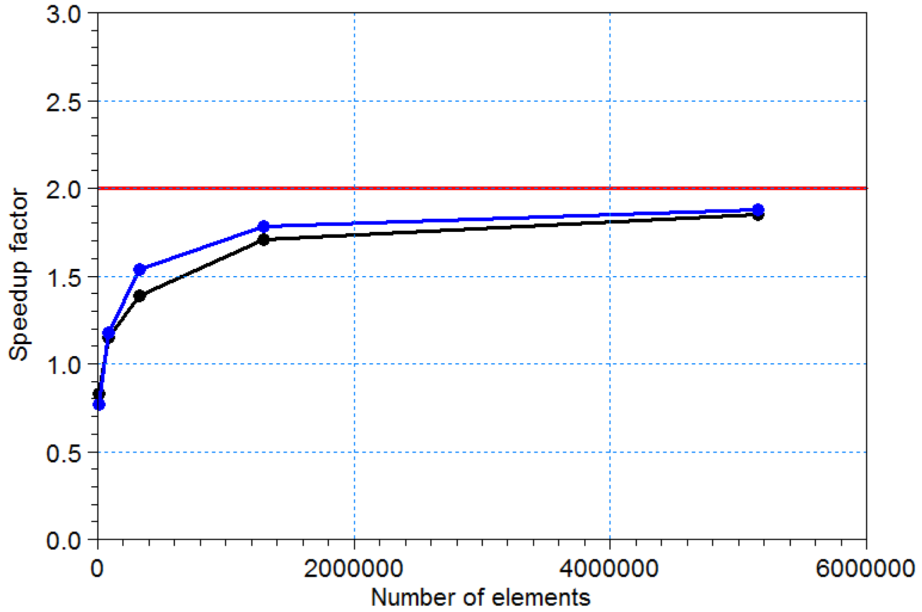

Figure 4.3 Speedup factor for two GeForce GTX Titan cards relative to a single GeForce GTX Titan card using single precision. Black line: first-order scheme; Blue line: higher-order scheme; Red line: Ideal speedup factor

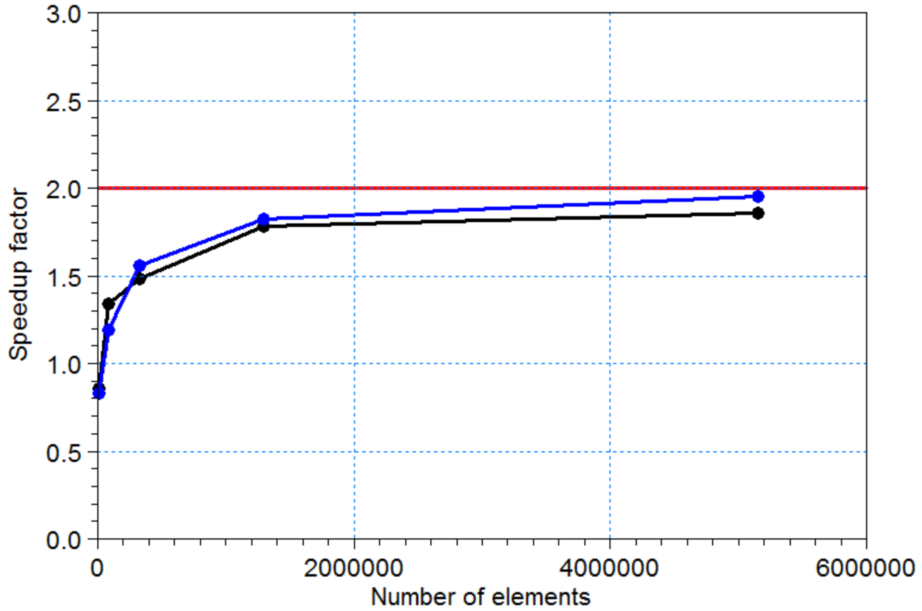

Figure 4.4 Speedup factor for two GeForce GTX Titan cards relative to a single GeForce GTX Titan card using double precision. Black line: first-order scheme; Blue line: higher-order scheme; Red line: Ideal speedup factor

4.2 Ribe Polder

Table 4.5 Simulations are carried out using single precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	26555.23	3751.94	7.07	2802.01	9.47 (1.33)

Table 4.6 Simulations are carried out using single precision and using the higher-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	97717.09	7350.22	13.29	5786.44	16.88 (1.27)

Table 4.7 Simulations are carried out using double precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	26555.23	4142.66	6.41	3238.33	8.20 (1.27)

Table 4.8 Simulations are carried out using double precision and using the higher-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	97717.09	8108.07	12.05	6762.41	14.45 (1.19)

4.3 EA2D Test 8A

Table 4.9 Simulations are carried out using single precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	1638.88	221.04	7.41	220.32	7.43 (1.00)
Mesh B	15725.61	1015.62	15.48	830.92	18.92 (1.22)
Mesh C	128194.12	5230.58	24.50	3333.22	38.45 (1.56)
Mesh D		35892.25		19201.26	(1.86)

Table 4.10 Simulations are carried out using double precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	1638.88	266.36	6.15	260.35	6.29 (1.02)
Mesh B	15725.61	1202.33	13.07	929.36	16.92 (1.29)
Mesh C	128194.12	6678.67	19.19	4003.44	32.02 (1.66)
Mesh D		42745.21		23521.65	(1.81)

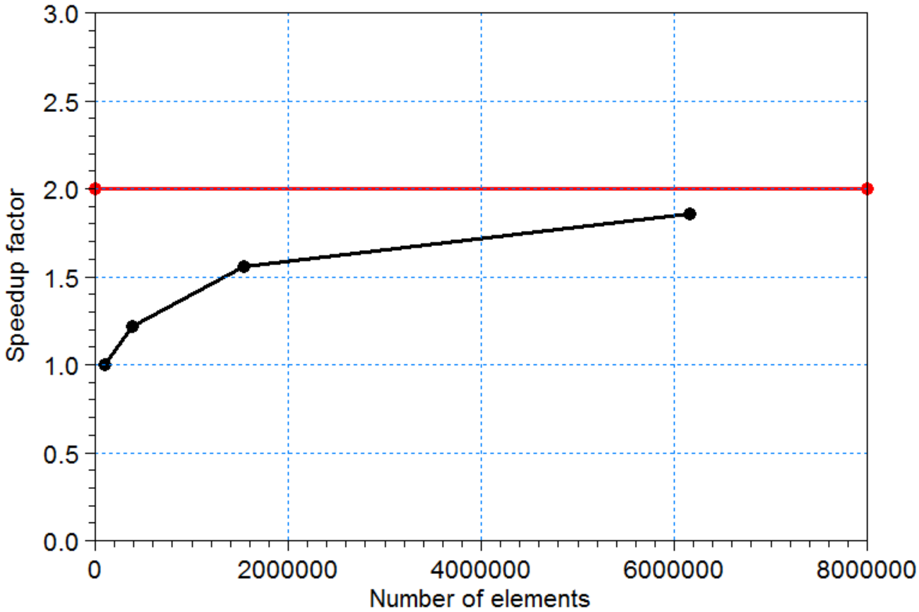

Figure 4.5 Speedup factor for two GeForce GTX Titan cards relative to a single GeForce GTX Titan card using single precision. Black line: first-order scheme; Red line: Ideal speedup factor

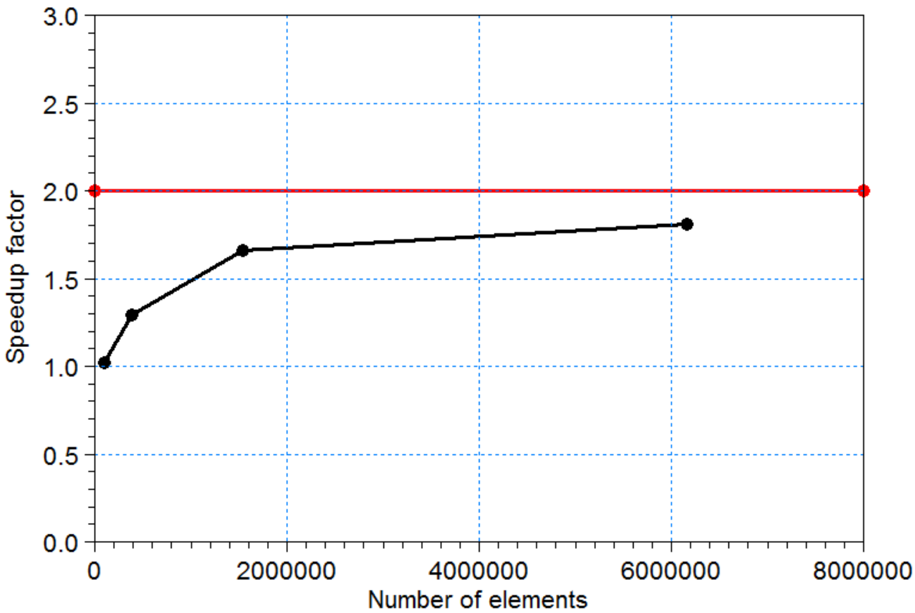

Figure 4.6 Speedup factor for two GeForce GTX Titan cards relative to a single GeForce GTX Titan card using double precision. Black line: first-order scheme; Red line: Ideal speedup factor

4.4 EA2D Test 8B

Table 4.11 Simulations are carried out using single precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	810.08	207.49	3.90	201.22	4.02 (0.96)
Mesh B	6568.61	710.77	9.24	558.70	11.75 (1.27)
Mesh C	52093.89	2792.37	18.65	1789.02	29.11 (1.56)
Mesh D		14465.55		8210.86	(1.76)

Table 4.12 Simulations are carried out using double precision and using the first-order scheme in time and space

Mesh	Time (s) CPU (1 core)	Time (s) GPU (1)	Speedup factor	Time (s) GPU (2)	Speedup factor
Mesh A	810.08	238.14	2.90	199.59	4.05 (1.19)
Mesh B	6568.61	756.56	8.68	541.93	12.12 (1.39)
Mesh C	52093.89	3235.74	16.09	1945.25	26.78 (1.66)
Mesh D		18084.01			

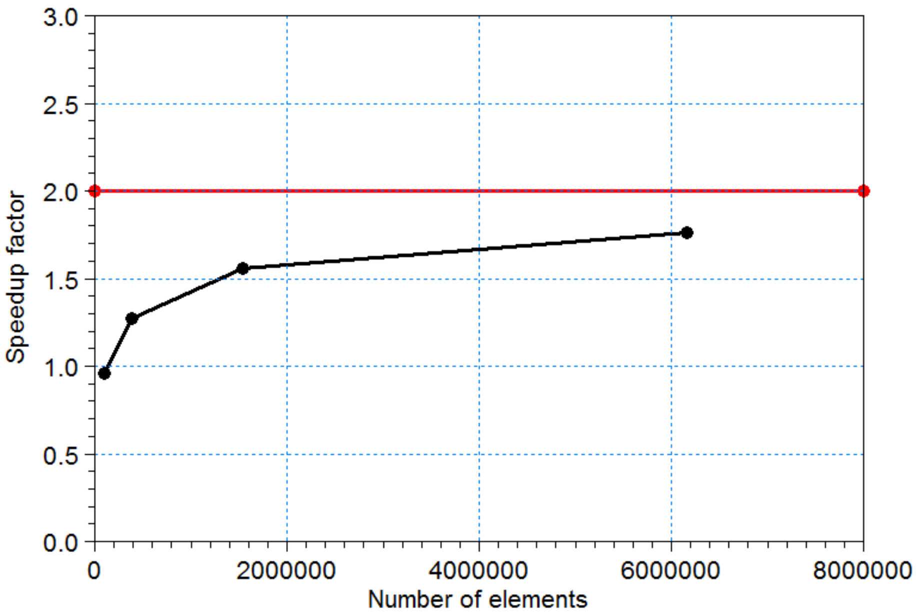

Figure 4.7 Speedup factor for two GeForce GTX Titan cards relative to a single GeForce GTX Titan card using single precision. Black line: first-order scheme; Red line: Ideal speedup factor

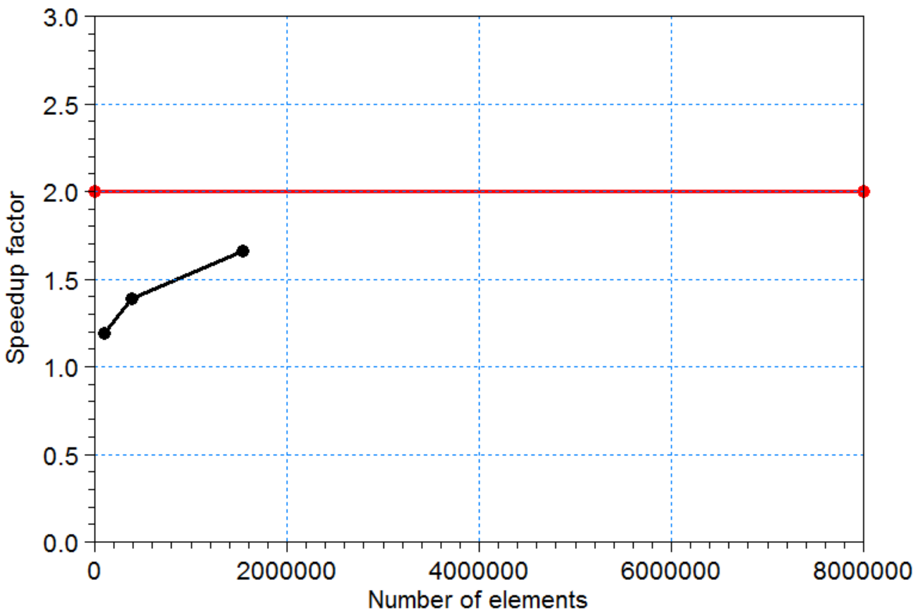

Figure 4.8 Speedup factor for two GeForce GTX Titan cards relative to a single GeForce GTX Titan card using double precision. Black line: first-order scheme; Red line: Ideal speedup factor

5 Benchmarking using a GPU cluster

These tests have been performed using the GPU cluster specified as hardware platform 2 in Table 2.1.

5.1 Mediterranean Sea

Table 5.1 Simulations are carried out using MIKE 21 Flow Model FM with multi-GPU parallelisation and using single precision and the higher-order scheme in time and space

Mesh	No. of GPUs	CPU time	Speedup factor
Mesh C	1	873.27	1.0
	2	498.72	1.75
	4	297.51	2.93
	8	197.63	4.41
	16	147.28	5.92

Mesh	No. of GPUs	CPU time	Speedup factor
Mesh D	1	794.67	1.0
	2	421.01	1.88
	4	227.28	3.49
	8	122.96	6.46
	16	75.84	10.47

Mesh	No. of GPUs	CPU time	Speedup factor
Mesh E	1	6555.68	1.0
	2	3355.75	1.95
	4	1713.48	3.82
	8	909.73	7.20
	16	497.35	13.18

Table 5.12 Simulations are carried out using MIKE 21 Flow Model FM with multi-GPU parallelisation and using single precision and the higher-order scheme in time and space

Mesh	No. of GPUs	CPU time	Speedup factor
Mesh C	1	1188.54	1.0
	2	683.36	1.73
	4	382.04	3.11
	8	251.04	4.73
	16	173.87	6.83

Mesh	No. of GPUs	CPU time	Speedup factor
Mesh D	1	1100.86	1.0
	2	567.36	1.94
	4	309.99	3.55
	8	177.97	6.18
	16	111.81	9.84

Mesh	No. of GPUs	CPU time	Speedup factor
Mesh E	1	9086.42	1.0
	2	4695.33	1.93
	4	2401.80	3.78
	8	1254.00	7.24
	16	700.99	12.96

Using the CPU version and 1 core, the computational times for the simulation with mesh C and D are 68065.85 s and 65100.9 s, respectively. Using mesh C and D the speedup factors for the single precision simulation with 16 GPUs compared to the simulation with the CPU version and 1 core are 462.15 and 858.39. For the double precision simulation the speedup factors are 391.47 and 582.24.

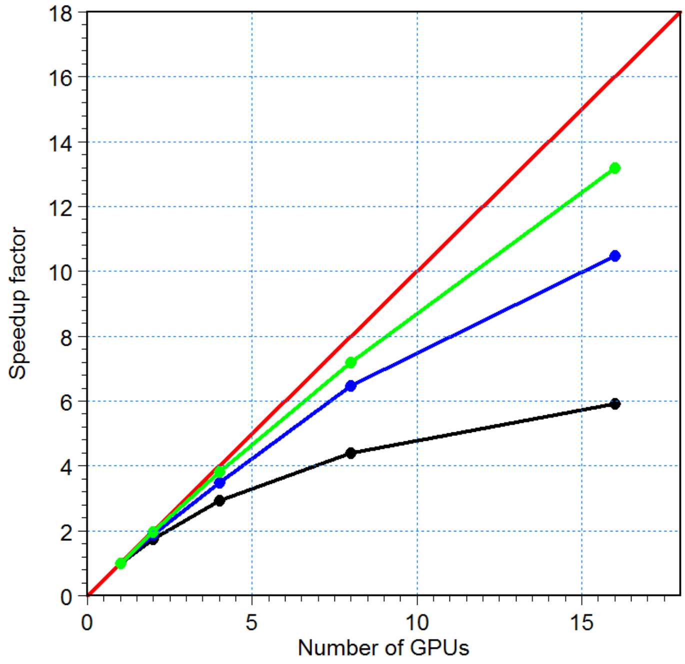

Figure 5.1 Speedup factor using single precision and higher-order scheme in time and space. Black line: GPU parallelisation with mesh C; Blue line: GPU parallelisation with mesh D; Green line: GPU parallelisation with mesh E; Red line: Ideal speedup factor.

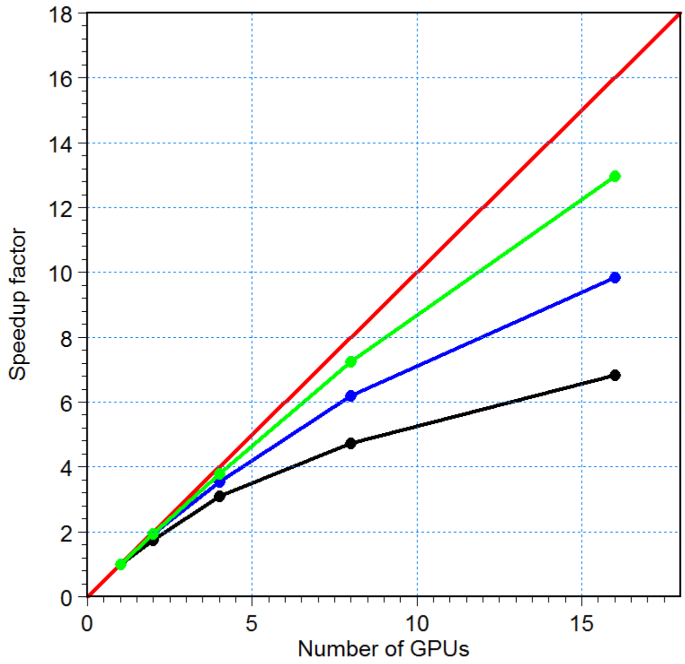

Figure 5.2 Speedup factor using double precision and higher-order scheme in time and space. Black line: GPU parallelisation with mesh C; Blue line: GPU parallelisation with mesh D; Green line: GPU parallelisation with mesh E; Red line: Ideal speedup factor.

6 Conclusions

The overall conclusions of the benchmarks are

- The numerical scheme and the implementation of the GPU version of the MIKE 21 Flow Model FM are identical to the CPU version of MIKE 21 Flow Model FM. Simulations without flooding and drying produces identical results using the two versions. Simulations with extensive flooding and drying produce results that may contain small differences.
- The performance of the new GPU version of MIKE 21 Flow Model FM depends highly on the graphics card and the model setup. When evaluating the performance by comparing with a single core (no parallelisation) CPU simulation the performance also depend highly on the specifications for the CPU.
- The speedup factor of simulations with no flooding and drying increases with increasing number of elements in the computational mesh. When the number of elements becomes larger than approximately 400.000 then there is only a very limited increase in the speedup factor for increasing number of elements.
- The use of multi-GPU shows excellent performance. To get the optimal speedup factor a large number of elements is required for each sub-domain.

7 References

/1/ Andersen, O.B., 1995, Global ocean tides from ERS-1 and TOPEX/POSEIDON altimetry, J. Geophys Res. 100 (C12), 25,249-25,259.

/2/ Bo Brahtz Christensen, Nils Drønen, Peter Klagenberg, John Jensen, Rolf Deigaard and Per Sørensen, 2013, Multiscale modelling of coastal flooding, Coastal Dynamics 2013, paper no. 053.

/3/ Néelz S., Pender G., 2010, Benchmarking of 2D Hydraulic Modelling Packages, Report published by Environment Agency, www.environment-agency.gov.uk, Copies of this report are available from the publications catalogue: http://publications.environment-agency.gov.uk.