

GPU Performance Tests of the RiverFlow2D Model

March 2024

RiverFlow2D™ model and documentation produced by Hydronia, LLC. Pembroke Pines, FL. USA.

Information in this document is subject to change without notice and does not represent a commitment on part of Hydronia, LLC.

RiverFlow2D™ is a trademark of Hydronia, LLC. 2011-2024.

All other products or service names mentioned herein are trademarks of their respective owners.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of Hydronia, LLC.

Technical documentation prepared by Asier Lacasta, Nicholas Calero, and Reinaldo Garcia.

Last document modification date: January, 2024.

Technical Support: support@hydronia.com

Web site: www.hydronia.com

Contents

CONTENTS..... III

LIST OF FIGURES IV

LIST OF TABLES V

1 INTRODUCTION 1

2 TEST CASES 2

2.1 TEST 1.....2

2.1.1 *Test 1 Results* 3

2.2 TEST 2.....4

2.2.1 *Test 2 Results* 6

2.3 TEST 3.....8

2.3.1 *Test 3 Results* 9

2.4 TEST 4.....11

2.4.1 *Test 4 Results* 12

3 COMMENTS..... 13

List of Figures

Figure 1 RiverFlow2D Plus triangular-cell mesh.....	1
Figure 2 Main window at the end of the simulation for test 1 mesh 3.	2
Figure 3 Test 1: Speed up of the GPU solution compared against the non-parallelized CPU version.	3
Figure 4 Test 2: Speed up of the GPU solution compared against the parallelized CPU version.	4
Figure 5 Main window at the end of the simulation for test 2 mesh 3.	5
Figure 6 Test 2: Speed up of the GPU solution compared against the non-parallelized CPU version.	6
Figure 7 Test 2: Speed up of the GPU solution compared against the parallelized CPU version.	7
Figure 8 Main window at the end of the simulation for test 3 using Tesla K40 (top) and using Tesla K80 (bottom).	8
Figure 9 Test 3: Speed up of the GPU solution compared against the non-parallelized CPU version.	9
Figure 10 Main window at the end of the simulation for test 4 using Tesla K40 (top) and using Tesla K80 (bottom).	11
Figure 11 Test 4: Computational cost (in seconds) for the Tesla K80, Tesla P100, Tesla V100 and RTX 2080 Ti devices.	12

List of Tables

Table 1 Technical specification summary of NVIDIA GPU hardware. 1

Table 2 Test 1: Run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model. 3

Table 3 Test 2: Run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model. 6

Table 4 Test 3: Run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model. 9

Table 5 Test 4: Run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model. 12

1 Introduction

RiverFlow2D, is suite of two-dimensional finite-volume models for rivers, floodplains and estuaries that include flow hydrodynamics, and add-on modules for erosion and deposition simulations, mud and debris flows, and pollutant dispersion. RiverFlow2D can route floods in rivers and simulate inundation over complex terrain at high resolution and with remarkable stability, accuracy and speed. The use of adaptive triangular-cell meshes enables the flow field to be resolved around key features in complex river environments. The GPU version allows performing hydrodynamic computations up to than 680 times faster than non-parallelized models. RiverFlow2D hydraulic simulation core has been developed in collaboration with the Computational Hydraulics Group of the University of Zaragoza in Spain.

This document presents several tests to demonstrate the performance of the RiverFlow2D GPU model on a variety of real project applications using several meshes with different resolutions and utilizing various NVIDIA GPU hardware cards (see Table 1).

Table 1 Technical specification summary of NVIDIA GPU hardware.

	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100
CUDA cores	2 x 2,496	3,584	3,584	5,120	4,352	10,496	16384	6912
Memory	24 Gb	11 Gb	16 Gb	16 Gb	11 Gb	24 GB	24 GB	40-80 GB

Note: The sequential version of the code was run on a computer with an Intel Core i7-3820 @ 3.60 GHz CPU.

In the tests described in this document we report runtimes for each application and calculate model speed ups with respect to the non-parallelized CPU model (using one core), which is the standard procedure to compute speedups. For instance, if the speedup is reported to be 100, it means that the model performs 100 times faster than the non-parallelized version.




Figure 1 RiverFlow2D Plus triangular-cell mesh.

2 Test Cases

We present different tests to illustrate the performance of the RiverFlow2D GPU model in five real applications using various GPU cards.

2.1 Test 1

The first test case involves the model application to a short reach of the Green River (USA) using three mesh resolutions: 19,079 cells (Mesh 1), 154,880 cells (Mesh 2), and 1,878,607 (Mesh 3).




Figure 2 Main window at the end of the simulation for test 1 mesh 3.

2.1.1 Test 1 Results: GPU vs. Single CPU Core

Table 2: Test 1 run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model.

Mesh	No. Cells	Intel CPU	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100	Max Speedup
Mesh1	19,079	00:00:08:14	00:00:00:18	00:00:00:38	00:00:00:13	00:00:00:11	00:00:00:46	00:00:00:31	00:00:00:21	00:00:00:12	45x
Mesh2	154,880	00:03:23:47	00:00:02:38	00:00:02:44	00:00:01:24	00:00:00:51	00:00:03:07	00:00:01:56	00:00:01:03	00:00:00:49	250x
Mesh3	1,878,607	08:23:17:47	00:01:28:04	00:01:08:28	00:00:33:40	00:00:18:49	00:01:00:39	00:00:41:06	00:00:24:47	00:00:13:15	975x




Figure 3 Test 1: Speed up of the GPU solution compared against the non-parallelized CPU version.

2.1.2 Test 1 Results: GPU vs. 16 CPU Cores

Table 3: Test 1 run times for RiverFlow2D in different GPU hardware. This chart compares with parallelized CPU model.

Mesh	No. Cells	16 Core Intel CPU	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100	GPUx16 Speedup
Mesh1	19,079	00:00:00:40	00:00:00:18	00:00:00:38	00:00:00:13	00:00:00:11	00:00:00:46	00:00:00:31	00:00:00:21	00:00:00:12	4x
Mesh2	154,880	00:00:29:09	00:00:02:38	00:00:02:44	00:00:01:24	00:00:00:51	00:00:03:07	00:00:01:56	00:00:01:03	00:00:00:49	37x
Mesh3	1,878,607	01:03:16:04	00:01:28:04	00:01:08:28	00:00:33:40	00:00:18:49	00:01:00:39	00:00:41:06	00:00:24:47	00:00:13:15	123x




Figure 4 Test 2: Speed up of the GPU solution compared against the parallelized CPU version.

2.2 Test 2

The second test is about an application of a hydraulic structure in New Orleans at high resolution. We present results for three meshes: 21,001 Cells in Mesh 1, 539,177 cells in Mesh 2 and 1,640,606 Cells in Mesh 3. The project was provided by Stantec.



Figure 5 Main window at the end of the simulation for test 2 mesh 3.

2.2.1 Test 2 Results: GPU vs. Single CPU Core

Table 3: Test 2 run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model.

Mesh	No. Cells	Intel CPU	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100	Max Speedup
Mesh 1	21,001	00:00:37:07	00:00:01:24	00:00:02:42	00:00:00:53	00:00:00:44	00:00:03:15	00:00:01:49	00:00:01:00	00:00:00:49	51x
Mesh 2	539,177	02:22:39:24	00:00:38:12	00:00:35:30	00:00:16:36	00:00:09:49	00:00:32:08	00:00:23:52	00:00:12:18	00:00:08:19	510x
Mesh 3	1,640,606	16:05:34:31	00:03:18:32	00:02:37:37	00:01:15:36	00:00:40:55	00:02:17:59	00:01:44:38	00:00:56:49	00:00:32:09	727x




Figure 6: Test 2 speed up of the GPU solution compared against the non-parallelized CPU version.

2.2.2 Test 2 Results: GPU vs. 16 CPU Cores

Table 3: Test 2 run times for RiverFlow2D in different GPU hardware. This chart compares with parallelized CPU model.

Mesh	No. Cells	16 Core Intel CPU	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100	CPUx16 Speedup
Mesh 1	21,001	00:00:03:28	00:00:01:24	00:00:02:42	00:00:00:53	00:00:00:44	00:00:03:15	00:00:01:49	00:00:01:00	00:00:00:49	5x
Mesh 2	539,177	00:12:23:49	00:00:38:12	00:00:35:30	00:00:16:36	00:00:09:49	00:00:32:08	00:00:23:52	00:00:12:18	00:00:08:19	89x
Mesh 3	1,640,606	02:06:08:49	00:03:18:32	00:02:37:37	00:01:15:36	00:00:40:55	00:02:17:59	00:01:44:38	00:00:56:49	00:00:32:09	101x




Figure 7: Test 2 speed up of the GPU solution compared against the parallelized CPU version.

2.3 Test 3

This test case represents an event for the simulation of a river in California (USA) including 357,611 cells. The event covers a period of 6 days and 23 hours.



Figure 8: Main window at the end of the simulation for test 3 using Tesla K40 (top) and using Tesla K80 (bottom).

2.3.1 Test 3 Results: GPU vs Single CPU Core

Table 4: Test 3 run times for RiverFlow2D in different GPU hardware. Intel CPU corresponds to the non-parallelized model.

No. Cells	Intel CPU	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100	Max Speedup
357,611	06:00:30:01	00:01:51:47	00:01:47:10	00:00:49:49	00:00:34:38	00:01:59:56	00:01:06:10	00:00:33:24	00:00:28:19	250x




Figure 9: Test 3 speed up of the GPU solution compared against the non-parallelized CPU version.

2.3.2 Test 3 Results: GPU vs 16 Core CPU

Table 5: Test 3 run times for RiverFlow2D in different GPU hardware. This chart compares with parallelized CPU model.

No. Cells	16 Core Intel CPU	Tesla K80	GTX 1080 Ti	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090	RTX 4090	Tesla A100	CPUx16
										Speedup
357,611	01:06:41:21	00:01:51:47	00:01:47:10	00:00:49:49	00:00:34:38	00:01:59:56	00:01:06:10	00:00:33:24	00:00:28:19	63x




Figure 10 Test 3 speed up of the GPU solution compared against the parallelized CPU version.

2.4 Test 4

This test reports on results of an ongoing collaboration with the National Oceanic and Atmospheric Administration (NOAA) of the USA. It shows a simulation of 420-mile reach of the Red River of the North located in Minnesota (USA). The event involves the routing of 3-month hydrographs.




Figure 11: Main window at the end of the simulation for test 4 using Tesla K40 (top) and using Tesla K80 (bottom).

2.4.1 Test 4 Results

The computer times of the non-parallelized CPU model is impractical for this test. Therefore, only the RiverFlow2D GPU model was used.

Table 6 Test 4: Run times for RiverFlow2D in different GPU hardware.

Mesh	No. of cells	Tesla K80	Tesla P100	Tesla V100	RTX 2080 Ti	RTX 3090
Mesh 1	4,616,546	00:20:50:46	00:09:23:12	00:06:22:28	00:12:22:54	00:10:02:24




Figure 12 Test 4: Computational cost (in seconds) for the devices capable of running this mesh.

2.5 Test 5

The fifth test is our latest test, added to see the impact of adding the mud/tailings flow equations into the simulations. This test case is a lower resolution study of the tailings dam failure that occurred in Brumadinho, Brazil on January 25, 2019. This simulation contains a 9 kilometer channel for the spill to reach the Paraopeba River in the southwestern portion of the map. The largest mesh size used was 1,096,772 cells.




Figure 13: Main window at the end of the simulation for test 5 using GTX 1080ti (top) and using RTX 4090 (bottom).

2.5.1 Test 5 Results: GPU vs. Single CPU Core

2.5.2

Mesh	No. Cells	Intel CPU	GTX 1080 Ti	RTX 3090	RTX 4090	Max Speedup
Mesh 1	19,792	00:00:12:08	00:00:00:43	00:00:00:27	00:00:00:19	16x
Mesh 2	150,055	00:04:16:31	00:00:08:45	00:00:04:03	00:00:02:05	125x
Mesh 3	1,096,772	04:04:03:32	00:02:38:27	00:01:03:15	00:00:33:12	181x




Figure 14: Test 5 speed up of the GPU solution compared against the non-parallelized CPU version.

2.5.3 Test 5 Results: GPU vs. 16 CPU Cores

Mesh	No. Cells	16 Core Intel CPU	GTX 1080 Ti	RTX 3090	RTX 4090	CPUx16 Speedup
Mesh 1	19,792	00:00:03:53	00:00:00:43	00:00:00:27	00:00:00:19	11x
Mesh 2	150,055	00:01:21:56	00:00:08:45	00:00:04:03	00:00:02:05	27x
Mesh 3	1,096,772	01:04:43:56	00:02:38:27	00:01:03:15	00:00:33:12	52x




Figure 15: Test 5 speed up of the GPU solution compared against the parallelized CPU version.

3 Comments

This report presents performance results of the Riverflow2D GPU model in several NVIDIA GPUs including the latest generation RTX 4090 and Tesla A100 cards. The Tesla A100 has taken the lead from all the tested devices, but the NVIDIA RTX 4090 card is a great option as it can be purchased at lower cost and boasts 24Gb of memory for larger meshes, and its acceleration capabilities are a substantial upgrade from the previous performance / value card, the GTX 1080 Ti and even the RTX 3090.

As demonstrated in the tests presented in this document, the remarkable performance of the RiverFlow2D GPU has several implications including:

- Computer run times are reduced from days to a few hours, or from hours to minutes, and from minutes to seconds in some cases.
- The RiverFlow2D GPU allows evaluating river flooding simulations of large river reaches that were impractical until recently due to excessive runtimes.
- The use of GPU technology developed in the RiverFlow2D code also allows using models with large resolution meshes involving millions of cells.
- The emergence of Pay-per-Use Cloud Services such as the Google Cloud where the P100, V100, and the new A100 cards are available at very attractive costs, facilitates the use of the RiverFlow2D GPU model for a wide range of applications.