

## Performance Enhancement of Flexpart Atmospheric Transport Model for GPU Environment

Don Morton<sup>1</sup>, Wolfgang Sommerer<sup>2</sup>, Jolanta Kusmierczyk-Michulec<sup>2</sup>, Robin Schoemaker<sup>2</sup>, Anne Tipka<sup>2</sup>

<sup>1</sup>Boreal Scientific Computing, Fairbanks, Alaska, USA

<sup>2</sup>CTBTO, Vienna, Austria



### INTRODUCTION

This is an overview of an ongoing EU-funded project to enhance the performance of the Flexpart Atmospheric Transport Model by using Graphical Processing Units (GPUs)

### METHODS/DATA

This section outlines the full workflow from GPU software installation, to setting up baseline Flexpart case, to compilation, execution, and first steps of OpenACC usage

START

### RESULTS

Description of various problems and solutions, creation of test mockup, and ongoing results

### CONCLUSION

This continues to be a work in progress. Much has been accomplished, and there is still much to do for verification of correctness and performance enhancement. Additional work includes preliminary test on Flexpart 11

The CTBTO Atmospheric Transport Modeling (ATM) group runs a large number of compute-intensive operational simulations on a daily basis in support of its mission.

The foundation of the simulations is the legacy Flexpart Lagrangian transport and dispersion model, which has undergone extensive modifications and enhancements in its 25 years.

With the recent CTBTO acquisition of GPU hardware, it was decided to explore the usefulness of modifying Flexpart for the new hardware in anticipation of shorter simulation times.

Previously, groups have explored Flexpart parallelisation with MPI (distributed memory) and OpenMP (shared memory) methods, finding that the most computationally-expensive components of the code are the processing of large GRIB datasets for the input of meteorology, and the computation of new particle positions at each timestep.

In this project, we aim to explore the application of GPU programming to Flexpart for the first time.



From *International Monitoring System (CTBTO)*



From *NVIDIA Tesla V100 GPU Architecture*

```

-----
| NVIDIA-SMI 418.126.02   Driver Version: 418.126.02   CUDA Version: 10.1   |
|-----+-----+-----+-----+-----+-----+-----+-----|
| GPU   Name               Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  0    Tesla V100-SXM2...    Off          | 00000000:06:00:0 Off  |    0%      Default  |
| N/A   28C    P0     44W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  1    Tesla V100-SXM2...    Off          | 00000000:07:00:0 Off  |    0%      Default  |
| N/A   29C    P0     44W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  2    Tesla V100-SXM2...    Off          | 00000000:0A:00:0 Off  |    0%      Default  |
| N/A   29C    P0     42W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  3    Tesla V100-SXM2...    Off          | 00000000:0B:00:0 Off  |    0%      Default  |
| N/A   26C    P0     42W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  4    Tesla V100-SXM2...    Off          | 00000000:85:00:0 Off  |    0%      Default  |
| N/A   28C    P0     44W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  5    Tesla V100-SXM2...    Off          | 00000000:86:00:0 Off  |    0%      Default  |
| N/A   29C    P0     44W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  6    Tesla V100-SXM2...    Off          | 00000000:89:00:0 Off  |    0%      Default  |
| N/A   29C    P0     43W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|
|  7    Tesla V100-SXM2...    Off          | 00000000:8A:00:0 Off  |    0%      Default  |
| N/A   27C    P0     43W / 300W |  0MiB / 32480MiB |             |
|-----+-----+-----+-----+-----+-----+-----+-----|

```

🏠

INTRODUCTION

OBJECTIVES

METHODS/DATA

RESULTS

CONCLUSION

⏪ ⏩

Please do not use this space, a QR code will be automatically overlaid

P4.3-486

## Objectives

This is an exploratory venture intended to utilise newly-installed GPU hardware for the performance enhancement of the complex atmospheric transport model, Flexpart.

Several groups have previously implemented parallelism in Flexpart using MPI for distributed processing and OpenMP for a shared-memory paradigm. Through this work it has been recognized that the most logical components for Flexpart parallelisation are the ingest and processing of meteorological data files, and the computation of new particle positions through the timestepping process.

Before jumping into the GPU paradigm, it needs to be understood that this is simply another parallel paradigm, and as such, performance enhancement is limited – more than we often think – by the code that is inherently sequential.

Still, in a saturated operational environment, even “some” performance improvement will lead to the ability to run more workload.

The primary objectives for this work are to

- establish a functional GPU-based development and testing environment on the new CTBTO system
- use the current CTBTO Flexpart code to establish a baseline for model output and performance, serving as a control for subsequent development and testing
- conduct a performance analysis of the current CTBTO Flexpart code to find the computational “hot spots,” and determine which part of the code will be targeted for GPU optimization
- create a development / test mockup environment of the targeted code to facilitate rapid and iterative experimentation of promising methods
- implement, test and assess the GPU-optimised code



INTRODUCTION

OBJECTIVES

METHODS/DATA

RESULTS

CONCLUSION



Please do not use this space, a QR code will be automatically overlaid

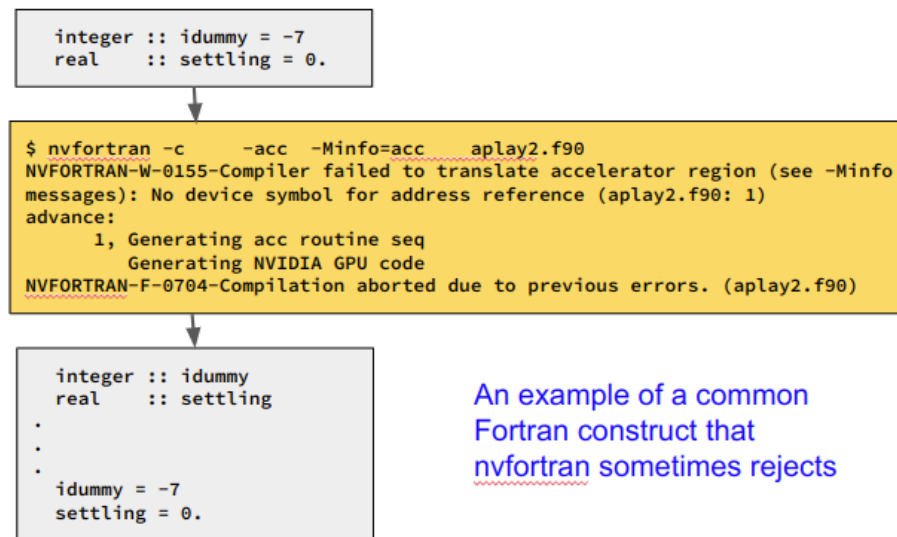
P4.3-486

Gfortran compilation and execution of the original Flexpart-CTBTO was performed, creating a baseline by which to compare results and performance of the future GPU-instrumented code.

Installation of full NVIDIA HPC SDK in user space was surprisingly straightforward, but in order for it to work correctly, compatible low-level NVIDIA libs must be installed at root-level. This was all subsequently tested on a basic linear algebra code in serial, CUDA and OpenACC to ensure a correct development environment.

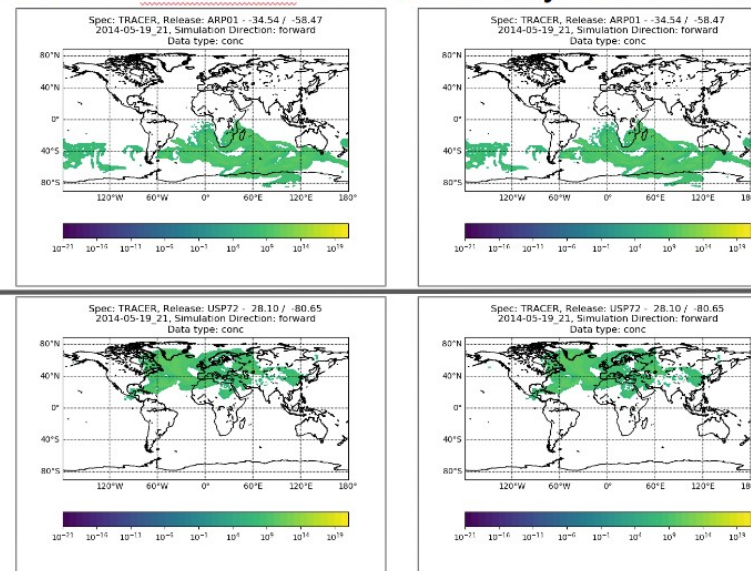
Initial attempts at compilation with the NVIDIA HPC nvfortran were problematic with both minor and major issues needing to be resolved. One issue – limiting the size of the problems we could handle - required a bug report to NVIDIA with a subsequent fix in the next release of the compiler. An additional problem is currently being considered for a bug report.

Numerical and visual comparison of the gfortran vs nvfortran codes was performed with satisfactory results.



An example of a common Fortran construct that nvfortran sometimes rejects

gfortran vs nvfortran, end of 10-day fwd simulation



- INTRODUCTION
- OBJECTIVES
- METHODS/DATA
- RESULTS
- CONCLUSION



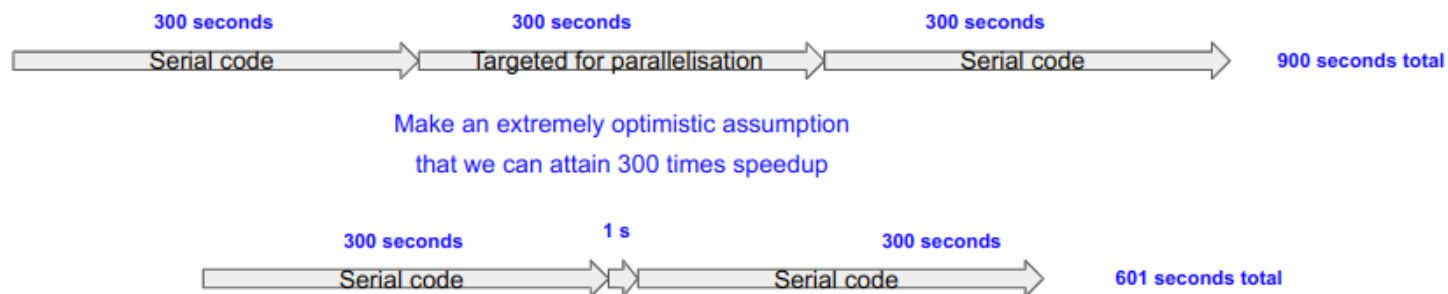
Please do not use this space, a QR code will be automatically overlaid





We knew this was going to be difficult for many reasons, and we knew that even upon completion, Amdahl's Law would significantly limit overall speedups that could be obtained.

Example of Amdahl's Law as applied to a typical, complex legacy code



However, given the popularity of Flexpart and its long legacy for 25 years, and its complexity, we see it as an opportunity to explore and report on – for better and worse – issues related to parallelisation in general, and GPU-isation specifically

This continues to be a work in progress. Upcoming activities include

- Evaluate – and revise current test code where necessary – the correctness of outputs and the performance of the new code
- Time permitting, experiment with cleaner allparticles loops, where iterations are truly independent (requires immense code refactoring)
- Apply what we have learned to the evolving Flexpart 11 in order to assess the value of continued efforts in this area