Environmental Modeling by Using the EuroHPC Petascale Supercomputer DISCOVERER

Tzvetan Ostromsky

Institute of Information and Communication Technologies,
Bulgarian Academy of Sciences,
Acad. G. Bonchev st., bl. 25-A, 1113 Sofia, Bulgaria

EuroCC conference: "High Performance Computing
for the Benefit of Scientists and Society"

November 21-22, 2022, Expo Hotel, Sofia, BULGARIA

Outline of the talk

- Introduction
- Description of the Danish Eulerian Model (DEM)
- High-performance parallel implementation of DEM (UNI-DEM)
- Numerical experiments with UNI-DEM on the Bulgarian EuroHPC supercomputer DISCOVERER
- Conclusions

Introduction

The Danish Eulerian Model in its unified MPI-parallel version (UNI-DEM) has been implemented on the new petascale supercomputer DISCOVERER, installed last year by Atos company in Bulgaria (at Sofia Tech Park). The machine is part of the European High Performance Computing Joint Undertaking (EuroHPC), which is building a network of 8 powerful supercomputers across the European Union – 3 pre-exascale (in Finland, Italy and Spain), the rest 5 – petascale.

Results of scalability experiments with different spatial grid versions (coarser-grain and fine-grain) of UNI-DEM on DISCOVERER supercomputer will be presented and discussed.

Description of the Danish Eulerian Model

DEM is a powerful and sophisticated large scale air pollution model, with some 30-year development history*. Over the years it was successfully applied in different long-term environmental studies in various areas, for instance:

- environment protection,
- human healthcare,
- agriculture (yield & losses of crops estimation),
- forestry, wildlife protection, etc.
- culture heritage protection, etc.

By processing a huge amount of data the model is able to calculate the variable concentrations of a number of pollutants and other chemical species over a long time period.

*[Hov et al., 1988; Zlatev, 1995; Alexandrov et al., 1997; Ostromsky & Zlatev, 2001; Zlatev & Dimov, 2006; Dimov et al., 2010]

Mathematical representation of DEM

The model can be described mathematically by a system of PDE for calculating the concentrations of a number of pollutants (and other chemical species that interact with the pollutants) in the atmosphere above the studied geographical region.

The main physical and chemical processes (advection, convection, diffusion, chemical reactions, emissions and deposition) are represented with separate terms in the right-hand-side of the system, as follows:

$$\frac{\partial c_s}{\partial t} = -\frac{\partial (uc_s)}{\partial x} - \frac{\partial (vc_s)}{\partial y} - \frac{\partial (wc_s)}{\partial z} +$$
$$+\frac{\partial}{\partial x}\left(K_x \frac{\partial c_s}{\partial x}\right) + \frac{\partial}{\partial y}\left(K_y \frac{\partial c_s}{\partial y}\right) + \frac{\partial}{\partial z}\left(K_z \frac{\partial c_s}{\partial z}\right) +$$
$$+E_s + Q_s(c_1, c_2, \ldots, c_q) - (k_{1s} + k_{2s})c_s, \quad s = 1, 2, \ldots, q.$$

Notation:

q - number of equations = number of chemical species,

c_s - concentrations of the chemical species considered,

u, v, w - components of the wind along the coordinate axes,

K_x, K_y, K_z - diffusion coefficients,

E_s - emissions in the space domain,

k_{1s}, k_{2s} - coefficients of dry and wet deposition respectively ($s = 1, \ldots, q$),

$Q_s(c_1, c_2, \ldots, c_q)$ - non-linear functions that describe the chemical reactions

Computational difficulties in the numerical solution of the model

- The above PDE system is too complicated for direct numerical solution – needs splitting with respect to the main physical / chemical processes;
- The numerical methods to solve it require spatial and time discretization with sufficient resolution (to ensure stable calculations and reliable results);
- Spatial discretization results in forming a large computational domain (spatial grid, covering the whole Europe);
- Time discretization for a long time period (for example - one year), leeds to many time-steps (suffisiently small, esp. on the chemistry stage);
- Huge amount of I/O data – requires special treatment in order to reduce the slow-down.

That's why the numerical solution of DEM requires powerful supercomputers and smart parallel algorithms with options for tunning on various hardware platforms.

Splitting into submodels

$$\frac{\partial c_s^{(1)}}{\partial t} = -\frac{\partial (u c_s^{(1)})}{\partial x} - \frac{\partial (v c_s^{(1)})}{\partial y} + \frac{\partial}{\partial x}\left(K_x \frac{\partial c_s^{(1)}}{\partial x}\right) + \frac{\partial}{\partial y}\left(K_y \frac{\partial c_s^{(1)}}{\partial y}\right) \quad (1)$$

$$= A_1 c_s^{(1)}(t)$$ horizontal advection & diffusion

$$\frac{\partial c_s^{(2)}}{\partial t} = E_s + Q_s(c_1^{(2)}, c_2^{(2)}, \ldots c_q^{(2)}) - (k_{1s} + k_{2s}) c_s^{(4)} = A_2 c_s^{(2)}(t) \quad (2)$$

chemistry, emissions & deposition

$$\frac{\partial c_s^{(3)}}{\partial t} = -\frac{\partial (w c_s^{(3)})}{\partial z} + \frac{\partial}{\partial z}\left(K_z \frac{\partial c_s^{(3)}}{\partial z}\right) = A_3 c_s^{(3)}(t) \quad (3)$$

vertical transport

The discretization of the spatial and time derivatives in each of the sub-models (1) – (3) results in forming respectively three large systems of ODE.

UNI-DEM package: versions and control parameters

The development and improvements of DEM throughout the years resulted in a variety of different versions with respect to the grid-size/resolution, the number of layers (2D / 3D) and the number of species in the chemical scheme. The most useful of them has been united under a common driver routine, called UNI-DEM. It provides an user-friendly access to the available up-to-date versions of the model with an easy way of selecting the appropriate parameters.

The versions, incorporated in UNI-DEM, and the parameters to select one of them are shown below.

Choosable parameters for selecting an optional UNI-DEM version				
Parameter	Description	Optional values		
NX = NY	Grid size (Grid step)	96×96 (50 km)	288×288 (16.7 km)	480×480 (10 km)
NZ	# layers (2D/3D)	1 or 10		
NEQUAT	# chem. species	35, 56 or 168		

Parallelization strategy

The main parallelization level (MPI) in UNI-DEM is based on the space domain partitioning.

- Requires communication on each time step (communication overhead)
- Domain overlapping of the advection-diffusion subproblems – computational overhead, grows up with increasing the number of MPI tasks.
- Improving the data locality for more efficient **cache utilization** by using chunks to group properly the small tasks in the chemistry-deposition and vertical exchange stages. An input parameter CHUNK-SIZE is provided for the purpose, which controls the amount of short-term reusable data in order to reduce the transfer between the cache and the main memory.
- Uses **MPI standard** routines for maximum portability.
- Additional pre-processing and post-processing stages are needed for scattering the input data and gathering the results.

Numerical experiments on the Bulgarian EuroHPC petascale supercomputer DISCOVERER

Characteristics of the DISCOVERER supercomputer

- Based on the Atos' platform BullSequana XH2000
- CPU type: AMD EPYC 7H12, 2.6GHz, 280W (code name Rome)
- 12 racks, 376 blades, 1128 nodes, 144384 cores in total
- 128 cores per node, RAM: 128 GB per node
- Total RAM 300 TB; Disk storage 2 PB
- Interconnection: Dragonfly+ with 200Gbps (IB HDR) bandwidth per link
- Sustained performance: 4.518 petaflops
- Theoretical peak performance: 6 petaflops
- TOP500 ranking: # 91 in the world; # 27 in EU (June 2021)

UNI-DEM parameters in the experiments

- Grid-versions:
  - coarse (96×96)
  - fine (480×480)
- Number of layers: 1
- Time period: 1 year
- Time step: 90 sec. (both in advection and chemistry stages)
- Cache utilization parameter: `NSIZE` $= 32$

UNI-DEM code parallelization and optimization details

- AMD compilers for the MPI code (mpifort, mpicc)
- Additional fine-grain parallelism by OpenMP of the performance-critical parts of the code, if appropriate (with *-fopemp* for invoking the compiler extensions and linking the necessary libraries)

Scalability of UNI-DEM on DISCOVERER (HPC, Sofia)

Time (T) in seconds and speed-up (Sp)
of UNI-DEM code (basic version) on DISCOVERER
$(96 \times 96 \times 1)$ grid, 35 species, CHUNKSIZE=32

NP	#	Advection		Chemistry		TOTAL		
(MPI)	NODES	T [s]	(Sp)	T [s]	(Sp)	T [s]	(Sp)	E [%]
1	1	439	(1.0)	902	(1.0)	1454	(1.0)	100 %
4	1	108	(4.1)	223	(4.0)	394	(3.7)	92 %
8	1	55	(8.0)	113	(8.0)	206	(7.1)	88 %
12	1	38	(11.6)	76	(11.9)	164	(8.9)	74 %
24	2	23	(19.1)	38	(23.7)	101	(14.4)	60 %
48	3	15	(29.3)	19	(47.5)	84	(17.3)	36 %

Time (T) in seconds and speed-up (Sp) (with respect to the number of MPI tasks, given in the first column) for running UNI-DEM (the basic MPI version) on the EuroHPC supercomputer DISCOVERER (in Sofia Tech Park, Bulgaria).

Scalability of UNI-DEM on DISCOVERER (HPC, Sofia)

Time (T) in seconds and speed-up (Sp)
of UNI-DEM code (basic version) on DISCOVERER
($480 \times 480 \times 1$) grid, 35 species, CHUNKSIZE=32

NP	#	Advection		Chemistry		TOTAL		
(MPI)	NODES	T [s]	(Sp)	T [s]	(Sp)	T [s]	(Sp)	E [%]
1	1	93022	(1.0)	81850	(1.0)	195036	(1.0)	100 %
4	1	23408	(4.0)	21615	(3.9)	48604	(4.0)	100 %
8	1	11830	(7.9)	11072	(7.4)	25045	(7.7)	96 %
12	1	7785	(11.9)	7112	(11.5)	17036	(11.3)	94 %
24	2	4075	(22.8)	3630	(22.5)	9150	(21.1)	88 %
48	3	2216	(42.0)	1845	(44.4)	4805	(40.1)	84 %
60	4	1790	(52.0)	1358	(60.3)	3638	(53.0)	88 %
96	6	1243	(74.8)	824	(99.3)	2701	(71.4)	74 %
120	8	1072	(86.8)	662	(123.6)	2394	(80.6)	67 %

Time (T) in seconds and speed-up (Sp) (with respect to the number of MPI tasks, given in the first column) for running UNI-DEM (the basic MPI version) on the EuroHPC supercomputer DISCOVERER (in Sofia Tech Park).

Additional speed-up via OpenMP on DISCOVERER

Time (T) in seconds and speed-up (Sp)
of UNI-DEM code (executing 96 parallel MPI tasks) on DISCOVERER
(480×480) grid, 35 species, CHUNKSIZE=32

NP (MPI)	# threads (OpenMP)	# NODES	Advection		Chemistry		TOTAL	
			T [s]	(Sp)	T [s]	(Sp)	T [s]	(Sp)
96	1	6	1243	(1.0)	824	(1.0)	2701	(1.0)
96	2	12	695	(1.8)	428	(1.9)	1925	(1.4)
96	4	24	403	(3.1)	246	(3.3)	1391	(1.9)
96	8	48	264	(4.7)	164	(5.0)	1225	(2.2)
96	16	96	177	(7.0)	109	(7.6)	1124	(2.4)

Time (T) in seconds and the speed-up (Sp) achieved by the OpenMP parallelization (the number of OpenMP threads given in the second column) for running UNI-DEM 2-level hybrid parallel code with fixed number (96) MPI tasks on the EuroHPC supercomputer DISCOVERER (using up to 1 rack / 96 nodes).

Conclusions

- The resent developments and a high performance implementation of the Unified Danish Eulerian Model (UNI-DEM) on the petascale EuroHPC supercomputer Discoverer were presented.

- The MPI parallel implementation proved to be efficient and scalable up to the level of granularity of the domain discretization (in dependence with the spatial grid size).

- The efficiency and speed-up is higher in the computationally-intensive stages. In particular, the chemistry stage (which does not need any communication between the tasks) has almost linear overall speed-up, even super-linear with respect to the number of MPI tasks (a cache-size effect due to decreasing size of processed data per task with increasing the number of MPI tasks).

- The advection stage scales pretty well too, taking into account that there is some unavoidable computational overhead due to overlapping subdomain boundaries of the partitioning.

Thank you for your attention!

Acknowledgments

This research is supported in part by the Bulgarian National Science Fund via the project *"Efficient methods for modelling, optimization and decision making"* - contract # KP-06-N52/5 (2021),

as well as by the EuroHPC JU via the Discoverer Fast Track & Benchmark Access project *"Installing, benchmarking, performance and scalability analysis of the UNI-DEM package for air pollution modelling"*.

References

1. V. Alexandrov, A. Sameh, Y. Siddique and Z. Zlatev, Numerical integration of chemical ODE problems arising in air pollution models, Env. Modeling and Assessment, 2 (1997) 365–377.

2. Ø. Hov, Z. Zlatev, R. Berkowicz, A. Eliassen and L. P. Prahm, Comparison of numerical techniques for use in air pollution models with non-linear chemical reactions, Atmospheric Environment 23 (1988), pp. 967–983.

3. Tz. Ostromsky, Z. Zlatev, Parallel Implementation of a Large-scale 3-D Air Pollution Model, Large Scale Scientific Computing (S. Margenov, J. Wasniewski, P. Yalamov, Eds.), LNCS-2179, Springer, 2001, pp. 309–316.

4. Tz. Ostromsky, Z. Zlatev, Flexible Two-level Parallel Implementations of a Large Air Pollution Model, Numerical Methods and Applications (I.Dimov, I.Lirkov, S. Margenov, Z. Zlatev - eds.), LNCS-2542, Springer (2002), pp. 545–554.

5. WEB-site of the Danish Eulerian Model, available at:
http://www.dmu.dk/AtmosphericEnvironment/DEM

6. Z. Zlatev, Computer treatment of large air pollution models, Kluwer (1995).

7. Z. Zlatev, I. Dimov, Computational and Numerical Challenges in Environmental Modelling, Elsevier, Amsterdam (2006).