

Porting the chemical module of Calgrid Code on a SIMD massively parallel platform

R. Lorenzini^a, L. Delle Monache^a, G. Zanini^a, M. C. Cirillo^a,
M. Celino^b & P. Palazzari^b

*^aENEA, Italian Agency for New Technology, Energy and
Environment,*

C.R. Casaccia

Environment Department,

C.P. 2400 - 00100 Roma A.D., Italy

*^bENEA, Italian Agency for New Technology, Energy and
Environment*

C.R. Casaccia

HPCN Project,

C.P. 2400 - 00100 Roma A.D., Italy

Abstract

The integration of stiff, very coupled, ordinary differential equations which describe pollutant chemical reactions is the most intensive computational task of photochemical models, since it requires at least 70% of the total CPU time.

As a consequence of the local nature of these equations, the integration can be performed very efficiently by a SIMD architecture.

In this work we present the porting of QSSA (Quasi Steady State Approximation) chemical solver of CALGRID photochemical model on the SIMD massively parallel platform Quadrics/APE100 which offers, in QH4 configuration, a peak performance of 25 Gflops.

318 High Performance Computing

1. Introduction

Parallel processing is a feasible way to afford the study of photochemical air pollution models on very large domains. Running of the chemical module which implies the integration of chemical ordinary differential equations (ODEs), and requires at least 70% of the CPU time¹, can be easily implemented on a single instruction multiple data (SIMD) massively parallel machine because of the local nature of equations. Integration of the other parts of model requires the flexibility of multiple instruction multiple data (MIMD) platforms.

In this paper we present the porting of the chemical solver of the photochemical model CALGRID on a massively parallel SIMD machine. This is the first part of the porting of entire CALGRID model on the hybrid MIMD/SIMD platform GENESI installed at the ENEA Research Center in Rome.

In next sections we describe the CALGRID model and code structure (section 2), the QSSA chemical solver ported on Quadrics/APE100 (section 3), the SIMD parallel machine (section 4) and the structure of the parallel code (section 5). Conclusions will be drawn in section 6.

2. Description of model

CALGRIDⁱⁱ is an Eulerian photochemical model for the study of atmospheric pollutant transport in wide areas. Partial differential equations for pollutant chemical reactions are solved using the time-splitting methodⁱⁱⁱ. This approach involves time integration of advection-diffusion partial differential equations by the fractional step method; this method, based on Marchuk decomposition, "factorizes" the time development operator connecting concentrations at times $n\Delta t$ and $(n+1)\Delta t$ as:

$$C^{n+1} = A_x \left(\frac{\Delta t}{2} \right) A_y \left(\frac{\Delta t}{2} \right) A_z \left(\frac{\Delta t}{2} \right) A_c(\Delta t) A_z \left(\frac{\Delta t}{2} \right) A_y \left(\frac{\Delta t}{2} \right) A_x \left(\frac{\Delta t}{2} \right) C^n \quad (1)$$

where A_x and A_y are horizontal transport operators, A_z is the vertical transport, source injection and physical depletion (e.g. dry deposition) operator, and A_c is the operator containing all chemical conversion terms. This particular approach is second-order accurate in time due to the cancellation of first-order errors via the alternating forward/reverse operator application.

Following the operator splitting, CALGRID includes separate modules for horizontal advection, vertical transport and dry deposition, photochemical reactions.

A cubic spline-based scheme is implemented in CALGRID for solving horizontal advection and diffusion equation. This scheme conserves mass exactly, prohibits negative concentrations and exhibits very little numerical diffusion.

Several options are provided for the computation of horizontal diffusivities, including user-specified stability-dependent values. A vertical transport scheme is provided that permits arbitrary and spatially varying spacing of vertical levels.

Vertical diffusivities are based on convective scaling in daytime boundary layer and local or z-less scaling in the night-time boundary layer.

The inclusion of chemical reactions into airshed models requires the solution, at every grid-point, of a set of coupled nonlinear and stiff ODEs.

Due to local nature of chemical reactions the ODEs which describe the pollutant concentration evolution at a grid-point are independent of other grid-points concentrations.

The CALGRID chemical sub-model provides two chemical integrators: the user can select either a solver based on QSSA or a hybrid solver based on the second-order predictor, iterated corrector scheme developed by Young and Boris⁴. In the present work the QSSA solver has been adopted.

The chemical mechanism currently implemented in CALGRID is the version of the Carter, Atkinson, Lurmann, Lloyd (CALL) scheme, developed at the State Air Pollution Research Center, known as COND2234⁵, which involves 54 species and 129 reactions.

The CALGRID code is written in Fortran 77 and is structured in three main subroutines:

- the SETUP subroutine, which performs all initialisation operations; all input data (meteorological data provided by CALMET preprocessor, initial and boundary concentrations, emission and run control data) are read, and output files are opened.
- the COMP subroutine, which is the core of model since it contains the main time loop.
- the FIN subroutine where output files are written and closed.

The main time loop can be summarized as follows:

```
DO loop over the number of advection time steps  $\Delta t$ 
  CALL HORIZ
  DO loop over the number of grid-points in x-direction
    DO loop over the number of grid-points in y-direction
      CALL ZTRANS
      DO loop over the number of grid-points in z-direction
        CALL CHEM
      ENDDO
      CALL ZTRANS
    ENDDO
  ENDDO
  CALL HORIZ
ENDDO
```

320 High Performance Computing

The HORIZ, ZTRANS and CHEM subroutines perform the following tasks:

- HORIZ

Solution of the two-dimensional advection equation for all vertical layers for half Δt .

- ZTRANS

Solution of the vertical advection-diffusion equation, including dry deposition and emissions injection, for half time step.

- CHEM

Solution of the set of chemical coupled non linear ODEs using the selected chemical solver.

3. The QSSA integrator

Usually the system of ODEs which must be solved to describe the evolution of chemical species is non-linear, highly coupled and very "stiff". A ODEs system is defined "stiff" if the eigenvalues of the Jacobian matrix vary over many orders of magnitude.

The system of chemical ODEs to be integrated in photochemical models have eigenvalues ranging over 15 orders of magnitude; in order to obtain a stable and accurate solution the integration is usually performed using implicit methods with small time steps. An example are the Gear's methods⁶; these methods can be very accurate but they are very time consuming, since they require the solution of a large set of non-linear equations.

The QSSA method, which is one of the most used, uses asymptotic or "quasi-linear" assumptions to obtain a trade-off between solution accuracy and computational efficiency.

This technique is based on the chemical equations written, at each grid-point, as follows:

$$\frac{dC_i}{dt} = P_i(C_1 \dots C_{N_{\text{all}}}) - D_i(C_1 \dots C_{N_{\text{all}}})C_i \quad i = 1, \dots, N_{\text{all}} \quad (2)$$

where C_i is the concentration of the i^{th} chemical species, P_i and D_i are production and loss terms, and N_{all} is the number of species in the mechanism. In order to integrate these equations over the advection time step, Δt is divided into a set of "chemistry time steps" δt of variable length. If production and loss terms can be assumed constant over the chemistry time step equation (2) can be solved analytically for the concentration $C_{i,n}$ at time step δt_n as:

$$C_{i,n} = \frac{P_{i,n}}{D_{i,n}}(1 - e^{-D_{i,n}\delta t_n}) + C_{i,n-1}e^{-D_{i,n}\delta t_n} \quad (3)$$

where $C_{i,n-1}$ is the concentration of i^{th} chemical species at time δt_{n-1} and production and loss terms at time δt_n are evaluated from concentrations at time δt_{n-1} .

The accuracy of the QSSA method depends on the careful choice of δt assuring that production and loss terms do not vary excessively during the chemistry time step. δt is adapted during the computation according to:

$$\delta t_n = \varepsilon \min \left(\frac{C_{i,n-1}}{P_{i,n-1} - D_{i,n-1} C_{i,n-1}} \right) \quad \text{for all } i \quad (4)$$

where ε is the maximum permitted fractional change in concentration during a chemistry time step. It is worth noting that δt set will be different on every grid-point, and therefore each grid-point will require a different number of applications of the formula (3) to complete the integration of chemical equations over an advection step.

This version of QSSA includes also a single corrector step, involving the computation of equation (4) using production and loss terms averaged over values computed at the beginning and end of the chemistry time step.

4. Description of Quadrics Supercomputer

APE100/Quadrics⁷ is a SIMD architecture developed at the INFN in Pisa and Rome. The figure 1 depicts the Quadrics system: compilation and program submission, as well as I/O tasks are carried out on a front-end SUN workstation.




Figure 1. The Quadrics system

322 High Performance Computing

The so-called Z-CPU manages the instruction stream (synchronous for all computing nodes) and performs global integer arithmetics; also I/O operations are transacted here.

The local floating point operations are performed on the processing nodes that consist of memory, a floating point unit called MAD (Multiplier and Adder Device), and the communication chip.

The MAD is the core of the processing node. It is able to perform 2 floating point operations within a clock cycle, that is $(25 \text{ Mhz})^{-1}$. Thus one processor maximally can perform with 50 Mflops. Each node is linked with communication lines to 6 neighbor processors in 3 dimensions.

The smallest configuration of the Quadrics computer is a Q1, a unit of 8 basic computer nodes, called a board; larger machines, like the 128-nodes QH1 or the 512-node Quadrics QH4, are composed of several boards.

A 512-node Quadrics QH4 is typically configured as an $8 \times 8 \times 8$ grid with toroidal boundary conditions in each direction and can performs with 25 Gflops.

The programming language of Quadrics system is the TAO language, which is very similar to Fortran language, and has structures of message passing.

The final parallel version of CALGRID code will be implemented on the hybrid SIMD-MIMD platform whose MIMD part is formed by four Sun Sparc20 workstations (connected by an Ethernet network) exchanging information by means of PVM libraries.⁸

One of the MIMD nodes is connected to a Quadrics machine with a selectable MAD number, namely 128 and 512, via a HIPPI channel which provides a sustained throughput of 16 Mbytes/sec.

This particular hardware arrangement has been named GENESI (GEneralized NETworked SIMd)⁹ as it allows, in general, to connect several SIMD machines via a network of workstations.

5. Parallelization methodology

Air pollutants in a given grid-point of the spatial domain may interact chemically with each other, but do not interact chemically with pollutants in other grid-points. This means that the system of coupled ODEs describing the concentration evolution in a grid-point is independent of concentrations at other grid-points, so every system can be solved independently.

This characteristic can be exploited to develop a general strategy for mapping the N grid-points of the spatial domain onto a machine having N nodes.

The spatial grid-points are grouped in M sub-sets, M-1 of which contain the same number of elements, whereas the Mth contains a number of grid-points equal to N modulo M-1.

We aim to avoid the use of TAO conditional statements in our code because they inhibit high performance; for this reason, if the Mth set contains

fewer elements than others, we assign to it some “dummy” grid-points to equalize the elements number of each set.

Each node performs the integration of the ODEs system in a sequentially mode for all the grid-points of its subset; the synchronous work mode of processors assures that concentrations of the M grid-points at time are updating.

This approach is very efficient since, if the structure of the serial algorithm is not modified, a processor can integrate the ODEs system for every grid-point of its subset without communicating with other processors.

However, in order to adapt the code to the SIMD architecture, a modification requiring communication has been introduced.

As outlined in section 3, the CALGRID QSSA solver adapts the time step used in chemical computations in order to optimize the integration. This procedure is carried out independently on every grid-point. As a consequence, the integration is performed at every grid-point in a different number of steps.

This kind of approach is not suitable for implementation on a SIMD architecture, where the instructions stream is synchronous, and this algorithm has been slightly modified.

The approach used is to calculate the time step according to equation (4) for each of M grid-points whose chemistry equations are integrating simultaneously, and then to actually use the minimum of the obtained values.

This choice may increase the number of chemical time steps on some grid-points with respect to the serial code, but, on a SIMD machine, it increases the code efficiency without affecting the computational time which is equal, for each set of M grid-points, to the integration time of the slowest point.

Although this algorithm requires communication among processors, it is the best compromise in terms of computational efficiency because of the dramatic reduction of conditional statements.

As a result pollutant concentrations are calculated with precision equal or even better than required.

Our simulations have been carried out by using the input files of a CALGRID test case, which works on a grid with 20×20 horizontal points and 10 vertical levels.

This parallel version of QSSA integrator has been implemented on a QH1 Quadrics machine and, as a result, efficiencies higher than 30% of the peak performance have been reached.

6. Conclusions

In this paper we have described the porting of QSSA chemical solver, as used by the photochemical model CALGRID, on the parallel SIMD architecture Quadrics /APE100.

The local nature of chemical equations has been preserved by the data mapping onto Quadrics nodes.

324 High Performance Computing

In order to avoid as far as possible TAO conditional statements the algorithm for the adaptive choice of chemistry time step has been modified.

As a result ODEs integration can be performed in a very efficient way with few communications among processors and pollutant concentrations are calculated with accuracy equal or greater than required, without affecting the CPU time.

The parallel version of QSSA integrator implemented on a QH1 Quadrics machine exhibits efficiencies higher than 30% of the peak performance.

References.

- ¹ Dabdub, D. & Seinfeld, J. Parallel computation in atmospheric chemical modeling, *Parallel Computing*, 1996, **22**, 111-130.
- ² Yamartino, R.J., et al. The CALGRID mesoscale photochemical grid model-I. Model Formulation, *Atmospheric Environment*, 1992, **26A**, 1493-1512.
- ³ Yanenko The Method of Fractional Steps, Springer, New York, 1971.
- ⁴ Young, T.R. & Boris, J.P. A numerical technique for Solving Stiff Ordinary Differential Equations Associated with the Chemical Kinetics of Reactive-Flow Problems, *Journal of Physical Chemistry*, 1977, **81**, 2424-2427.
- ⁵ Carter, W.P. A detailed mechanism for the gas-phase atmospheric reactions of organic compounds, *Atmospheric Environment*, 1990, **24A**, 481-518.
- ⁶ Gear, C.W. Numerical Initial Value. Problems in Ordinary Differential Equations, Prentice-Hall, Englewood Cliffs, N.J., 1971.
- ⁷ Bartoloni, A., et al., *Int. J. Mod. Phys.*, 1993, **C 4**, 955.
- ⁸ A. Geist, A.Beguelin, J.Dongarra, W.Jiang, R.Manchek, V.Sunderam, PVM 3 user's guide and reference manual, Oak Ridge National Laboratory, Oak Ridge, Tennessee (1994).
- ⁹ In ENEA (Italian National Agency for New Technologies, Energy and Environment), a project, named GENESI (GEneralized Networked SImd) is in progress as basis for the PQE2000 project that will produce the hardware and the software of the tera-flop platform based on hybrid hyper-nodes SIMD-MIMD (for more information: <http://www.enea.it/~hpcn/hpcn01e.html>).