CHEMICAL NETWORK PROBLEMS SOLVED ON NASA/GODDARD'S MASSIVELY PARALLEL PROCESSOR COMPUTER

Seog Y. Cho and Gregory R. Carmichael
Department of Chemical and Materials Engineering
University of Iowa
Iowa City, Iowa 52242

David M. Cohen and Mehmet H. Oguztuzun
Department of Computer Science
University of Iowa
Iowa City, Iowa 52242

ABSTRACT

The single instruction stream, multiple data stream MPP processor consists of 16,384 bit serial arithmetic processors configured as a 128 x 128 array whose speed can exceed that of current supercomputers (Cyber 205). This paper presents and discusses the applicability of the MPP for solving reaction network problems including the mapping of the calculation to the architecture, and CPU timing comparisons.

Keywords: Chemical Net Work Problems, Eulerian Transport/Chemistry Model, Air Pollution Model, Coupled ODE-IVP.

INTRODUCTION

A detailed model which describes the transport and removal of photochemical oxidants, and acidic species and precursors in the troposphere has been under development for the past nine years. The present analysis consists of about 30 coupled three-dimensional time-dependent non-linear partial differential equations and about 50-100 coupled non-linear ordinary differential equations.

The model is representative of a number of comprehensive Eulerian transport/chemistry models being developed for regional air pollution problems. However, these models are only feasible when run on "supercomputers". Our model was developed on the NASA-Langley CDC-STAR computer and is currently being exercised on the NASA-Langley Cyber 205, the NCAR CRAY-1 and a FAC M240 in Nagoya University, Japan. The execution times are 0.025

$$\frac{\text{CPU-sec}}{\text{grids-time step}}$$ on the FAC M 240,

0.007 $\frac{\text{CPU-sec}}{\text{grids-time step}}$ on the

CRAY-1, and 0.70 $\frac{\text{CPU-sec}}{\text{grids-time step}}$ on a

VAX 11/780. Thus, a 24-hour simulation on the CRAY-1 for the eastern United States with 9500 grid points requires 100 CPU-minutes.

Our experience has shown that transport/chemistry models can execute about 70-100 times faster on the "supercomputers". However, 100 CPU-minutes/simulation-day is still too large for most applications. Since typical applications require simulations of seven to ten days. Therefore, to exercise these models various simplifying assumptions are used to decrease the CPU time. However, these assumptions add additional errors and uncertainties to the model results. Faster computers will enable the execution of the "best-science" model version.

Currently about 90% of the CPU time is

spent doing the chemistry calculations. The chemistry introduces the stiffness, the coupling, and the non-linearity into the model. Thus, the highest priority in continued model development is to search for ways in software and hardware to reduce the chemistry calculation. The purpose of this paper is to describe our attempts to exploit massively parallel computer architectures to accelerate the chemistry calculations.

MODEL OVERVIEW

The regional-scale combined transport/chemistry/deposition model is Eulerian and treats 50 chemical species. Thirty species are advected, while the remaining species are short-lived and are modeled using pseudo-steady state methods. The mathematical analysis consists of partial differential equations for the advected species and additional algebraic equations for the steady-state species. The advected species satisfy

$$\frac{\partial C_i}{\partial t} + \nabla\ (VC_i) = \nabla\cdot K\cdot\nabla C_i + R_i + S_i - G_i,\ i = 1,\ldots,30; \quad (1)$$

where C_i is the gas-phase concentration of the ith chemical species, V is the wind velocity vector, K is the eddy diffusivity tensor, R_i denotes the chemical reaction term, S_i is the source term, and G_i is used to describe the mass transfer between the gas and condensed phases. The algebraic equations for the gas-phase species assumed to be at steady state are written as

$$R_i(C_1, C_2,\ldots,C_{50}) = 0, \quad i = 31, \ldots,50. \quad (2)$$

These equations are representations of general chemically reactive flow problems.

Simulation of regional transport, chemistry and deposition as described by Equations (1) and (2) requires numerical integration. The method presently used is a combination of the concept of fractional time steps and one-dimensional finite elements. This is referred to as Locally One-Dimensional, Finite-Element Method (LOD-FEM). The LOD procedures (Mitchell, 1969) split the multi-dimensional partial differential equation into time dependent, one-dimensional problems which are solved sequentially. The transport equations are solved using a Crank-Nicolson Galerkin finite element technique. Chemistry and mass transfer equations are solved using an adaptation of the semi-implicit Euler method proposed by Preussner and Brand (1981).

SCOPING STUDIES

Test Problem 1: Chemical Network Problem

To evaluate the ability of the MPP to calculate chemical network applications a simple four species test problem was selected. The four species (C_1, C_2, C_3, C_4) are involved in the following chemical reactions:

$$C_1 + C_2 \xrightarrow{k_1} C_3 \quad (a)$$

$$C_3 \xrightarrow{k_1} C_1 + C_2 \quad (b)$$

$$2C_2 \xrightarrow{k_3} C_4 \quad (c)$$

$$C_4 \xrightarrow{k_4} 2C_2 \quad (d)$$

The transport equations describing this system is represented by Eq(1) with i=1,2,3 and 4.

As mentioned in the model overview section, one way of numerically solving complex transport chemistry network

problems is to split the equation into transport and chemistry parts. The chemistry calculations using this technique requires solving the set of equations

$$\frac{\partial C_i}{\partial t} = R_i \qquad i=1, \ldots, \text{\# of species} \tag{5}$$

at each grid point in the discretized space.

The use of the semi-implicit Euler method to solve Eq(5) results in the equations

$$\frac{dC_i}{dt_r} = -C_i \sum_j d_i^j \prod_{k=1} C_k + \sum_\ell P_i^\ell \prod_m C_m \qquad i=1,\ldots 4 \tag{6}$$

This set of ODE-IVP's is solved within each transport time step, i.e.,

$$t_o \leq t_r \leq t_{transport} .$$

Now consider the case when we have 16,384 grid points in the discretized spatial domain. Therefore each chemical calculation within each transport step requires the solution of 16,384 sets of Eq(6). To implement the solution of these equations on the MPP requires first the choice of how to map the equations to the architecture. In this case we have chosen to simply view each processor as a grid point in the discretized space, and to have each processor solve its own set of Eq(6). The algorithm for solution of Eq(6) is written in Parallel Pascal and resides on the VAX. The initial conditions and constants are distributed to each processor and the algorithm is executed on the MPP and output is sent back to the VAX.

The CPU time required for execution of 100 time steps on the MPP of this 4 species mechanism at 16,384 grid points is 0.293 CPU-seconds. The same problem was executed on the VAX-11/780 and required 138 CPU-seconds. Thus for this chemical network problem the MPP executed a factor of 470 times faster than VAX 11/780!!

Test Problem 2: NO_x Transport in Eastern United States

To test combined transport/chemistry network problems on the MPP a 3-dimensional test problem describing the transport and chemistry of NO, NO_2, O_3, and HNO_3 in the lower troposphere was selected. The governing set of equations is given by Eqn. (1). An oversimplified chemical mechanism is used in this test calculation, i.e.,

$$NO + O_3 \rightarrow NO_2 + O_2 \tag{e}$$

$$NO_2 + h\nu \xrightarrow[+O_2]{} NO + O_3 \tag{f}$$

$$NO_2 + OH \rightarrow HNO_3 \tag{g}$$

(The OH concentrations are given by an empirical formula and the O_2 concentrations are assumed constant.

Sample results for this test problem calculated on a VAX 11/780 are shown in Figures 1 and 2. Presented are the NO_x emissions for the eastern United States, and the 24-hour averaged predicted surface concentrations of NO, NO_2, and HNO_3. The meteorological conditions simulated are those of July 4, 1974. The grid system used in the simulation was 32 x 32 x 16, and a transport time step of 15 minutes and a chemistry time step of 1 second was used.

This combined transport/chemistry problem is currently being run on the MPP. There are two choices for mapping this problem to the MPP. One method is to perform the chemistry calculations on the MPP and the transport part of the calculation on the VAX. This

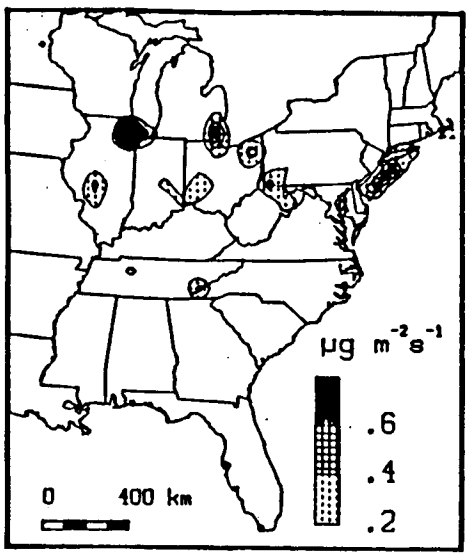

Figure 1. The emission of NO_x at July 4, 1974 at surface.

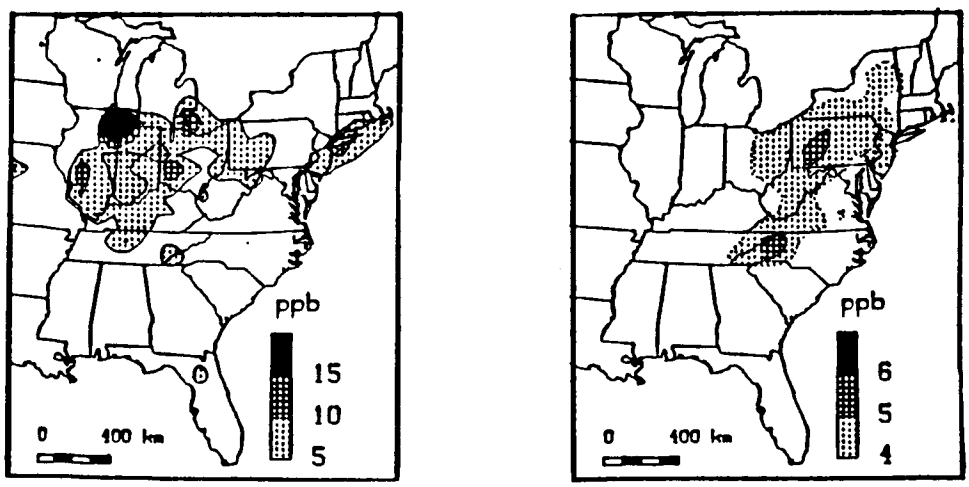

Figure 2. Averaged concentration of July 4, 1974 at surface.

method is currently being tested. The other method is to perform the entire calculation on the MPP. We are currently developing an algorithm to solve the sets of tridiagonal matrices on the MPP which arise from the transport processes.

The above test calculation indicates the MPP is well suited for chemical network problems where each node can hold the entire mechanism. Current memory restrictions limit the size of the chemical mechanism that can be solved in this fashion. At present each processor can hold 32 32-bit variables. (It is planned to increase the storage in the near future.) However, it is possible to handle larger chemical mechanisms. One way is to group processors together. For example if 128 words are required at each node then four processors can work together. This in turn would reduce the maximum number of grids possible by a factor of 4. Another way is to make use of the staging memory.

SUMMARY

The suitablity of the MPP computer for calculation of chemical network problems is under evaluation. To date the MPP has been used to calculate a test problem which represents one component of a sophisticated chemically reactive flow problem. Specifically the set of coupled ODE-IVP's describing the chemical reactions occuring at 16,384 spatial grid points was calculated. This problem is ideally suited for the MPP because by using operator splitting, the chemistry at each grid point acts independently from that of the other grids (within each transport time step). This test calculation showed that the MPP can perform a 100-time-step calculation 470 times faster than the same calculation on the VAX 11/780. Also since nearly 90% of CPU time of large chemically-reactive flow problems is spent doing the chemistry calculations, the MPP architecture offers great potential for CPU savings for model applications. Coupled transport-chemistry problems are now being tested on the MPP.

ACKNOWLEDGMENTS

The authors wish to thank the staff of NASA/Goddard, and specifically Jim Abeles, for their assistance in helping us use the MPP, the University of Iowa Graduate College for providing funds to carry out this project, and Chi-Chul Yang, a graduate assistant in the Department of Computer Science for assisting with MPP programming.

REFEREMCES

1. Gabriel, R. P. "Massively Parallel Computers: The Connection Machine and NON-VON", Science, **231**, 975-978, 1986.

2. Hwang, K. and Briggs, F. A. Computer Architecture and Parallel Processing, McGraw-Hill, New York, Table 6.3, 425, 1984.

3. Lerner, E. J. "Parallel Processing Gets Down to Business", High Technology, July, 1985, pages 20-28.

4. Mitchell A. R. Computational Methods in Partial Differential Equations, John Wiley, New York, 1969.

5. Preussner P. R. and Brand K. P. "Application of a Semi-Implicit Euler Method to Mass Action Kinetics", Chem Eng. Sci., 10, 1633-1641, 1981.