



Project Summary

A Computer Architecture for Research in Meteorology and Atmospheric Chemistry

John McHugh, John Pierce, Don Rich, Janet Dunham, David McLin, and Nick Kanopoulos

This study examines the feasibility of constructing a peripheral hardware module that could be attached to a mini or mid-sized computer to accelerate the execution of large air pollution models, such as the EPA's Regional Oxidant Model (ROM). Crucial information necessary to design such an accelerator is acquired by running the ROM computer code under instrumentation which shows how the computational load is distributed within the model and the data transfer rates between each step of the model execution. These data reveal that a model such as the ROM is not amenable to acceleration using a vector-type architecture because the computational burden is too inhomogeneous in space and time. They also show that the most computationally intensive portions of the model involve little or no data communication with neighboring points in the grid mesh. Together, these facts suggest that the most efficient accelerator design is one that utilizes a system of loosely coupled processors. Two such designs are explored. In one, the host computer performs part of the model execution while the most computationally intensive portions are executed by a network of slave processors under control of the host machine. In the second design, called the tile machine, the spatial domain simulated by the model is subdivided into small pieces and a separate processor executes the full ROM code within each subdomain. Simulations show that an accelerator based on the tile machine architecture would be capable of executing the

ROM up to 100 times faster than the host machine working alone.

This Project Summary was developed by EPA's Atmospheric Sciences Research Laboratory, Research Triangle Park, NC, to announce key findings of the research project that is fully documented in a separate report of the same title (see Project Report ordering information at back).

Introduction

In the mid 1970's the EPA began the development of the Regional Oxidant Model (ROM) as a state-of-the-science tool for both research and regulatory applications. In the early 1980's similar efforts were begun to develop regional scale acid deposition and particulate models. All of these models are very large and complex and as a result they require large computer resources to operate. For example, on the EPA's VAX 8600 computer, the ROM requires about 24 hours of CPU time to simulate a 24-hour day. The same simulation requires only 2 hours on the Agency's IBM 3090 computer, but time on this machine is expensive and it is not available on a demand basis. The acid rain and particulate models are expected to require even longer execution times. In short, increases in the scale and complexity of air pollution models have outpaced the growth in the available computing capability. One possible solution to this problem is to build computer hardware specifically designed to perform the computations that these models entail. This report describes a study conducted by the Research Triangle Institute for the

EPA to determine the feasibility of constructing equipment of this type, capable of accelerating the execution of the ROM on the Agency's VAX 8600 and/or VAX 785 computers.

Procedure

The first step in this study was to understand the basic mathematical structure of the ROM. This model consists of 30 coupled nonlinear partial differential equations that are solved on a 3-dimensional grid containing some 7500 mesh points. Each equation describes the effects of winds, chemistry, sources, deposition and other physical processes on the concentration of each of 30 chemical species. The equations are solved as an initial value problem to simulate concentration patterns in space at 30 minute intervals over periods up to 1 month in length.

The technique employed in the ROM to solve the governing equations splits each of the 30 dependent variables into two components. One of these, represented mathematically by the Greek letter, Γ , embodies the effects of the horizontal winds only. Each of the 30 Γ components is the solution of a single, linear, 2-dimensional differential equation. The simplicity of the Γ equations is offset somewhat by the fact that the value of Γ at each grid point is coupled to the corresponding values at 36 neighboring grid points. This is an important factor in the design of a model accelerator.

The second component of the dependent variables, represented mathematically by γ , embodies the effects of chemistry, sources, deposition and other processes. Unlike Γ , the γ components are governed by a system of nonlinear equations in which the γ 's of all 30 species are coupled together. However, for the most part the value of any one of the γ components at any given grid point is independent of the γ values at neighboring grid points. This fact suggests that the γ equations would be amenable to treatment by parallel processors.

The second stage of the feasibility study was to run the ROM code under instrumented conditions to determine how much machine time is spent on the Γ computations, how much is devoted to γ , and how these times vary from grid point to grid point and time step to time step during the model execution. These data, plus information on the data transfer rates within the model, provide

the specifications for feasible accelerator designs.

Conclusions

The analyses performed on the ROM computer code during an actual model run revealed a number of important features. First, the ROM is totally dominated by floating point computations. Consequently, an essential requirement of a model accelerator is high speed floating point capability. Second, an average 10 percent of the model execution time is used by the Γ computations and the remaining 90 percent is devoted to the γ calculations. This lopsided split reflects the heavy computational burden created by the nonlinear chemical processes represented by the γ equation. Third, the CPU time required to derive γ varies by more than a factor of seven from grid point to grid point and time step to time step.

Taken together these facts suggest that the most efficient design of a ROM accelerator would be one consisting of a set of loosely coupled processors. The investigators proposed and simulated two separate designs. In one system, referred to as the Γ/γ architecture, the host machine performs the Γ computations while the γ equations are solved

by a set of slave processors which together comprise the accelerator module. Simulations of this design indicated that with an accelerator consisting of about 25 Micro VAX-II class processors, the VAX 8600 could execute the ROM about 10 times faster than it can acting alone. This would make the ROM run times on the VAX comparable to those achievable on the IBM 3090.

In the second proposed system, referred to as the tile machine, the spatial domain simulated by the ROM would be subdivided into a number of small areas, or tiles. In this system a separate processor would handle both the Γ and γ equations in each tile while the host machine would take care of data communication and load leveling chores. Analysis of this design indicated that with a large enough number of processors, an accelerator based on this architecture could potentially increase ROM execution speeds by a factor of 100. Computing capacity of this magnitude would open up new areas of model applications and research that are presently infeasible due to prohibitive computer time requirements.

The report also addresses the issues of fault tolerance, the use of array processors, hardware interfaces, cost vs performance trade-offs, and other topics.

J. McHugh, J. Pierce, D. Rich, J. Dunham, D. McLin, and N. Kanopoulos are with Research Triangle Institute, Research Triangle Park, NC 17709.

Robert G. Lamb is the EPA Project Officer (see below).

The complete report, entitled "A Computer Architecture for Research in Meteorology and Atmospheric Chemistry," (Order No. PB 88-145 313/AS; Cost: \$19.95, subject to change) will be available only from:

National Technical Information Service

5285 Port Royal Road

Springfield, VA 22161

Telephone: 703-487-4650

The EPA Project Officer can be contacted at:

Atmospheric Sciences Research Laboratory

U.S. Environmental Protection Agency

Research Triangle Park, NC 27711




United States
Environmental Protection
Agency

Center for Environmental Research
Information
Cincinnati OH 45268

Official Business
Penalty for Private Use \$300

EPA/600/S3-87/049



0000329 PS

U S ENVIR PROTECTION AGENCY
REGION 5 LIBRARY
230 S DEARBORN STREET
CHICAGO IL 60604