MAUSAM, 52, 1(January 2001), 191-200

551.509.313

On the parallelisation of weather and climate models

S. K. DASH

Centre for Atmospheric Sciences,
Indian Institute of Technology,
NewDelhi – 110016 (India)

सार - मौसम पूर्वानुमान और जलवायु के अध्ययनों के लिए प्रयुक्त किए गए सांख्यिकीय निदर्शों के लिए बहुत अधिक संख्या में आकलित साधनों की आवश्यकता पड़ती है। इस क्षेत्र में हाल ही में किए गए अनुसंधान से पता चला है कि सही पूर्वानुमानों के लिए निदर्शों का उपयोग करने के लिए अत्याधिक उच्च विभेदन, परिष्कृत ऑकड़ा समीकरण तकनीकों तथा भौतिक प्राचलीकरण योजनाओं एवं बहु निदर्श समुच्चय संघठनों की आवश्यकता पड़ती है। वस्तुतः सही पूर्वानुमान हेतु अपेक्षित स्थानिक विभेदन के लिए आकलित विद्युत की आवश्यकता पड़ सकती है जोकि किसी भी सुपर कंप्यूटर के एकल संसाधित्र में संसाधित विद्युत को मद्देनज़र रखते हुए निषेधात्मक रूप से अधिक है। पिछले दो दशकों में आकलन प्रौद्योगिकी में हुए विकास के परिणामस्वरूप कई संसाधकों सहित समानांतर कंप्यूटर का प्रादुर्भाव हुआ है जिनमें एकल कंप्यूटर की तुलना में अत्याधिक मात्रा में आकलित विद्युत की आपूर्ति करने की क्षमता होती है। अब विश्वस्तरीय जलवायु निदर्श को लंबे समय तक एकीकृत करने के लिए कार्यस्थलों के समूहों में अथवा व्यक्तिगत कंप्यूटरों में समानांतर प्रयुक्त किया जा सकता है। फिर भी, इस क्षेत्र में अधिकतम क्षमता को प्राप्त करने के लिए मार्ग में आए अवरोधों को अभी दूर करना है। विश्वस्तरीय मौसम और जलवायु निदर्शों के संबंध में अंतः संसाधन संचारण एक मुख्य समस्या है। इस शोध पत्र में मुख्यतः प्रचालनात्मक पूर्वानुमान और अनुसंधान के प्रमुख केन्द्रों में मौसम और जलवायु के निदर्शों में अनुरूपता की स्थिति, मौसम और जलवायु निदर्शों में अन्तर्निष्ट समान्तरता, अंतः संसाधित संचारण में आने वाली समस्याओं और अधिकतम संमानांतर क्षमता को प्राप्त करने के विभिन्न उपायों पर विचार विमर्श किया गया है।

ABSTRACT . The numerical models used for weather forecasting and climate studies need very large computing resources. The current research in the field indicates that for accurate forecasts, one needs to use models at very high resolution, sophisticated data assimilation techniques and physical parameterisation schemes and multi-model ensemble integrations. In fact the spatial resolution required for accurate forecasts may demand computing power which is prohibitively high considering the processing power of a single processor of any supercomputer. During the last two decades, the developments in computing technology show the emergence of parallel computers with a number of processors which are capable of supplying enormously large computing power as against a single computer. Today, a cluster of workstations or personal computers can be used in parallel to integrate a global climate model for a long time. However, there are bottlenecks to be overcome in order to achieve maximum efficiency. Inter-processor communication is the key issue in case of global weather and climate models. The present paper aims at discussing the status of parallelisation of weather and climate models at leading centres of operational forecasting and research, the inherent parallelism in weather and climate models, the problems encountered in inter-processing communication and various ways of achieving maximum parallel efficiency.

Key words – Numerical weather prediction, Parallel computing.

1. Introduction

The numerical weather prediction and climate research involve solving a number of nonlinear coupled mathematical equations by numerical techniques. Essentially, these are very complicated initial and boundary value problems, the accurate solution of which depends on many factors such as, the accuracy and resolution of initial observed time dependent meteorological parameters and slowly varying boundary conditions, sophisticated data assimilation techniques, numerical techniques to solve the nonlinear equations, physical parameterisation schemes to represent the sub-grid scale processes etc. Research in all these aspects need powerful and fast computers. Hence it has been observed that the progress in numerical weather prediction and computer technology go hand-in-hand. Fast and powerful computers have been very essential and helpful tools in

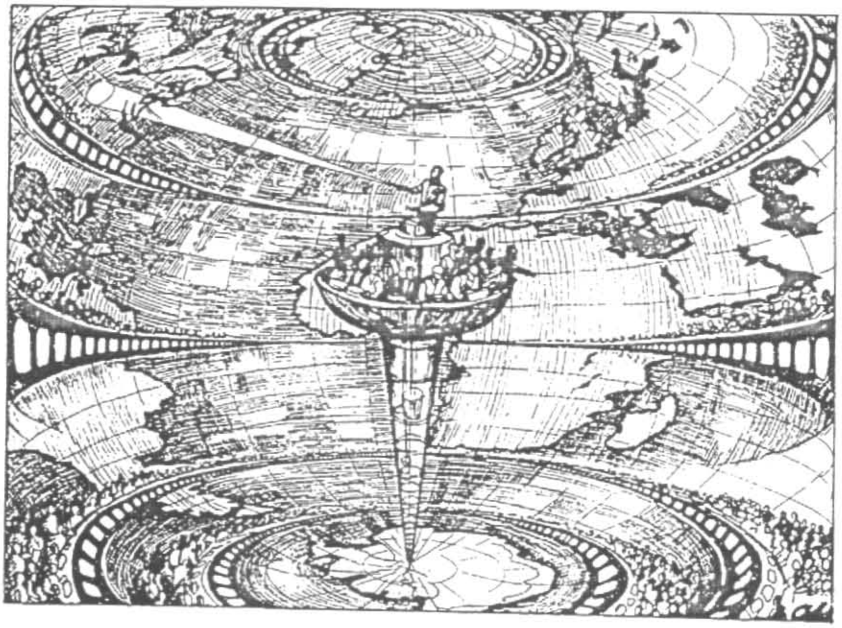

Fig. 1. Weather forecasting factory (Richardson 1922) source: Bengtsson (1984)

understanding some of the basic processes in atmospheric circulation and implementing those in numerical models for the improvement in operational weather forecasting and climate research. The advent of supercomputers was crucial for development of meteorological models in the early eighties. It was realized that to achieve performance of several Giga flops and beyond, one has to turn to multi-tasking systems with more than one Central Processing Unit (CPU).

The most commonly used models in numerical weather prediction (NWP) are spectral and grid point models. In grid models, the discrete versions of a closed set of nonlinear equations are repeatedly solved at different grids for getting the new values of wind, temperature, humidity, surface pressure and rainfall at future time. Hence the job for the whole globe can be divided into a number of tasks over different geographical portions of the earth by taking into consideration the boundaries between different zones adequately. For example two processors can work in parallel for the two hemispheres. The other widely used method in NWP is the spectral method where the time dependant parameters are expressed as spherical harmonics in the horizontal. In the spectral model, the computations are performed latitude wise. Hence one latitude can be assigned to one processor. Also there is no coupling between different spherical harmonics and hence the degree of parallelism is relatively high. Majority of the physical parameterisation schemes are horizontally independent and can be split up into sets of vertical columns, each set of which defines an independent task. Thus even at low level of parallelisation, the sequential codes of most widely used global spectral models and also grid point models have a lot of parallelism inherent in them. Since the meteorological models are three dimensional in space, in principle, by the help of domain decomposition method it is not difficult to use more than one processor in parallel. In fact, one may attribute this concept of domain decomposition and parallelisation to Lewis F.Richardson's original idea (Fig.1) of weather prediction by numerical processes (Richardson 1922), where he envisaged to divide the surface of the earth into 64,000 equal parts and employ one person with one computer to solve the mathematical equations and a man at the central stage to maintain uniform speed of progress of each sector and to synchronise the final results.

The sub-grid scale processes such as radiative transfer, condensation, small scale turbulent and convection portions are highly conditional depending on the presence of clouds and hence are non-parallel tasks. These type of computations give rise to load imbalance between the processors. There are also other non-parallel tasks such as input/output (I/O), inter-processor data communication and synchronisation in the existing sequential codes. These non-parallel tasks are the cause of concern and hence in early eighties it was felt that a small number of (8-16 processors as maximum) powerful processors were better for multi-tasking (ECMWF 1984). More number of processors create more overheads and problems in memory management and synchronisation. Domain decomposition method of achieving parallelism in NWP models is very straight forward, however, the advantage gained by decomposition may be lost through synchronisation, I/O and other overheads.

It is seen that both spectral and grid point techniques can be successfully used (Kreitz and Prior 1995) on distributed memory systems. Computer architecture does not seem to impose a clear preference for one or the other. It was believed that at high resolutions the spectral method is impractical (Foster *et al.* 1992), because of its impossibly high computational demands and communication requirements. Thus the grid point models seemed to have advantages over the spectral models. This was because of enormous imbalance between computational power and available communication bandwidth (Kreitz and Prior 1997). Currently, the bandwidth is more in proportion to the computing power and hence both spectral and grid point models can run equally effective on highly parallel machines.

During the last two decades, a lot of progress has been made in meteorological computing so far as the use of parallel processors is concerned. Although almost all the leading Centres of operational forecasting and research in NWP and climate studies use multiprocessors, the use of massively parallel computers is way behind. There are many problems encountered while trying to enhance the efficiency of parallelisation. The present paper aims at giving a comprehensive picture of the status of parallelisation of meteorological models in our country as well as outside, the problem of inter-processor communication of meteorological data, the methods adopted to overlap computations and communication in order to achieve enhanced efficiency and brief summary of consolidated efforts needed.

2. Inherent parallelism in spectral GCMs

It is well known that the fundamental equations governing the atmospheric circulation such as the horizontal momentum equations along the latitudinal and longitudinal directions, the continuity equation, the thermodynamic energy equation, the hydrostatic equation, the surface pressure tendency equation and the moisture equation (Haltiner 1971) are used in NWP model. However, the hydrostatic equation and the thermodynamic equation can be combined and similarly the surface pressure tendency equation and the continuity equation can be combined to give a set of five coupled equations in five time dependent meteorological parameters such as the two horizontal components of wind, temperature, moisture and surface pressure. These five nonlinear coupled equations constitute a closed system, which, in principle, can be solved at all future times from a given initial condition and with prescribed boundary conditions. In compact form the set of five equations can be written as

$$\frac{\partial X}{\partial t} = D(X) + P(X) \tag{1}$$

Where X is any model variable such as wind, temperature, humidity and surface pressure, D stands for dynamical processes such as advection, pressure forces etc. and P represents the physical processes such as radiation, condensation, surface boundary processes etc. These operators constitute very complex nonlinear partial derivative terms in the three spatial co-ordinates. Since it is not possible to get simple analytical solutions of the meteorological parameters from the above set of equations, it is a standard practice to resort to numerical techniques such as finite difference, finite element, spectral, semi-Lagrangian etc. to solve the above set of coupled equations.

The spectral method is one of the very commonly used methods in NWP because of various reasons including those of the easy way of handling the poles, less computational time requirements etc. In the spectral method, each of the meteorological parameters X in Eqn. (1) are represented by the following double series at any instant of time.

$$X(\lambda,\mu,\sigma,t) = \sum_{m=-M}^{M}\left[\sum_{n=|m|}^{N(m)} X_n^m(\sigma,t)P_n^m(\mu)\right]e^{im\lambda} \tag{2}$$

Here the spherical coordinates are λ, $\mu(=\sin\phi)$ and $\sigma(=p/p_s)$. λ, ϕ, p, p_s represent the longitude, latitude, pressure and surface pressure respectively. The coordinate t represents the time. X_n^m are the complex spectral coefficients and P_n^m are the associated Legendre functions of the first kind of order m and degree n. Also m represents the zonal wave number and n is often called the two dimensional index or the total wave number. M is the highest Fourier wave number included in the east-west representation and $N(m)$ is the highest degree of the associated Legendre function included in the north-south representation. As m and n increase, they correspond to increasing the model resolution in the horizontal and hence the number of mathematical operations. Theoretically speaking the upper limits of both the summations in Eqn.(2) should have been infinity, which is impossible to deal with and hence the truncation of both the series. There are two commonly used truncation schemes such as Rhomboidal and Triangular. In the Triangular scheme $N(m) = M$ and in Rhomboidal scheme $N(m) = |m| + M$. The inner sum in Eqn.(2) is done as a vector product over n and the outer is performed by Fast Fourier Transform (FFT). The spectral models in the Triangular and Rhomboidal truncations are represented by the letters T and R respectively followed by the maximum truncation number M and similarly the number of vertical levels in the model is represented by the letter L followed by the number of levels. Thus a model at resolution T126L31 represents the horizontal truncation at M=126 in the Triangular scheme and there are 31 levels in the vertical.

The physical processes are local in nature and hence conceptually radiation, condensation, evaporation, frictional drag etc. can not be represented as waves. These are to be computed at each grid point of the model. Similarly, the nonlinear terms in the coupled equations when represented in terms of summation of waves become messy and need large storage in terms of interaction coefficients. The transform method developed by Eliassen *et al*. (1970) and Orszag (1970) independently is largely used in global spectral models to reduce the storage considerably. In this method, the integrals arising out of the nonlinear terms and the physical parameterisation computations are completed at transform (Gaussian) grids and transformation from grid to spectral space is done by Gaussian quadratures. The spectral coefficients of meteorological parameters at future time are obtained by solving the equations in the spectral space and then the reverse transform is done from spectral to grid domain. The cycle goes on till the end of the model integration. Thus each time step of model integration involves transformation from spectral to grid and *vice-versa*. Thus spectral transform method involves multiple FFTs in the longitudinal direction and multiple Legendre Transforms (LTs) in the latitudinal direction.

In spectral transform method, the nonlinear terms of any order may be computed exactly within the assumed truncation if the transform grid is defined on a sufficient number of Gaussian lines of latitude and equally spaced lines of longitude. It has been shown that in practice it is sufficient to allow for an unaliased computation of only quadratic terms and hence for the triangular truncation, the number of Gaussian latitudes J should satisfy the condition,

$$J \geq (3M + 1)/2 \qquad (3)$$

Similarly in order to allow exact Fourier transform of quadratic terms, the number of points I in the east-west direction must satisfy the condition,

$$I \geq (3M + 1) \qquad (4)$$

According to Eqn.(3), the number of Gaussian latitudes identified for horizontal resolutions T21, T42, T63, T80, T126, T159, T319 are usually 32, 64, 96, 128, 192, 256 and 512 respectively. The number of grid points along each latitude line is just the double of the number of Gaussian latitudes for each resolution according to Eqn.(4). Thus one can easily calculate the total number of transformed grids at each vertical level of a model resolution. In case of one dimensional (1D) parallelisation along the latitude lines, it is convenient to assign the FFT computations of one north-south (N-S) pair (one Gaussian latitude in the northern hemisphere and the corresponding latitude in the southern hemisphere) of latitudes to one processor because the Legendre functions are symmetric with respect to the equator. Thus in 1D parallelisation along the latitude lines, one employs $J/2$ number of processors *i.e.* 16, 32, 48, 64, 96, 128 and 256 processors for T21, T42, T63, T80, T126, T159 and T319 model resolutions respectively. Since a GCM is three dimensional, in principle it is possible to adopt domain decomposition methods to achieve 3D parallelisation. Hence a lot of parallelism is inherent in a GCM.

3. Inter-processor communication

According to Eqn.(2), the LTs are done by multiplying the Legendre functions P_n^m corresponding to each N-S pair of Gaussian latitudes and the spectral coefficients of meteorological variables X_n^m to get the symmetric and anti-symmetric parts of the Fourier components X^m. The Fourier components for each latitude are calculated from the symmetric and anti-symmetric parts and then the FFTs are done for each latitude to get the values of the meteorological terms at all the transformed grids of the spectral model. All the non-linear terms are evaluated at this stage of the time step of model integration, which include the computations arising out of the physical parameterisation schemes.

In the spectral transform method, after the computations on the transformed grids are over, all the terms are converted back to their spectral equivalents at the next time step based on the following equation.

$$X_n^m = \int_{-1}^{1}\left[\frac{1}{2\pi}\int_{0}^{2\pi} X(\lambda,\mu)e^{-im\lambda}d\lambda\right]P_n^m(\mu)d\mu \qquad (5)$$

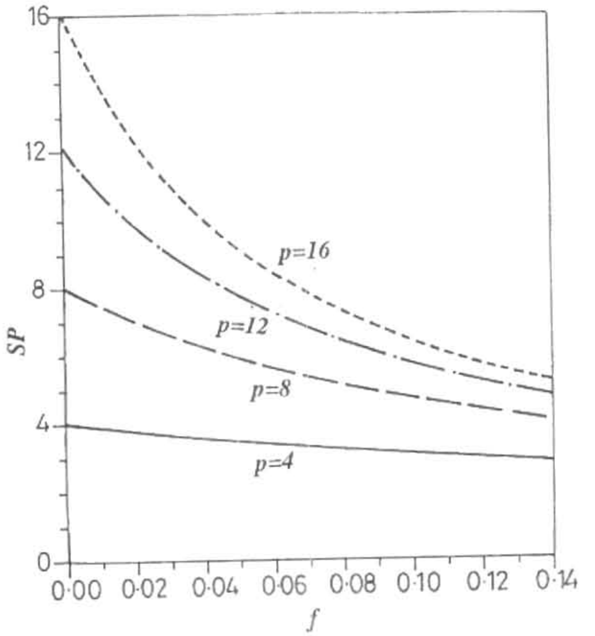

Fig. 2. Theoretical maximum speedup (*SP*) when using *p* processors for varying sequential parts (*f*) of the code (Amdahl 1967)

The inner integral represents the Fourier transform done by the FFTs and the outer integral is performed *via* the Gaussian quadrature,

$$X_n^m = \sum_{j=1}^{J} X^m(\mu_j) P_n^m(\mu_j) w_j \tag{6}$$

Here w_j is the Gaussian weight at the latitude number j. The summation in Eqn.(6) is over the total number of Gaussian latitudes from pole to pole. However, taking advantage of the symmetry of the Legendre functions and the Gaussian weights with respect to the equator, in practice, the summation is done over $J/2$ number of latitudes from the equator to the pole. Thus the contributions from each N-S pair of latitudes to the spectral coefficients X_n^m are computed in each processor and then summed over on any single processor or a group of processors to get the values of X_n^m finally. This involves inter-processor communication, which is the main cause of low performance of massively parallel computers. The total number of spectral coefficients S corresponding to any meteorological parameter or non-linear term for truncation M is given by,

$$S = (M+1)(M+2)/2 \tag{7}$$

It may be noted that the spectral coefficients are complex numbers and hence, in practice, there are $2S$ real numbers corresponding to S complex numbers. It is also worth noting that as the value of M is increased to enhance the resolution of a global model, the number of spectral coefficients increases by M^2. It is clear that when the number of processors used in parallel are increased to effectively integrate a high resolution GCM, the number of spectral coefficients increases enormously and hence the time for inter-processor communication.

4. Measure of parallelisation

The most commonly used parameters to measure the advantages of parallelisation are the theoretical maximum speedup and the parallel efficiency based on the famous Amdahl's law (Amdahl 1967). According to him every task contains some fixed fraction of critical work, that while being performed, requires all other sub tasks idle, awaiting either access or results. If a program runs for t units of time on a single processor and is modified to use p identical processors in parallel (p>1) then the fraction f of the total time t can be found during which the program will use only one processor (sequential part). Time for multitasked execution t_p is equal to the sum of the times for sequential and multitasked parts. Hence the theoretical maximum speedup (Amdahl 1967) can be written as,

$$SP = \frac{t}{t_p} = \frac{1}{f + (1-f)/p} \tag{8}$$

As shown in Fig.(2), the overheads enhance so much that there is practically no advantage in increasing the number of processors from 8 to 16.

The formulation of NWP and climate models suggests that it is not possible to parallelise these models completely. The most important fact for the sequential model not parallelised completely is the addition of contributions from each pair of N-S latitudes to get the final values of spectral coefficients of meteorological variables as represented in Eqn. (6). While this job is performed on a single processor or a group of processors, it is obvious that some of the processors remain idle. Also in the parallel code, the contributions from each pair of latitudes are available on individual processors and need to be communicated to a single processor or a group of processors for their sum, which was not so in case of the sequential code. Thus there is extra inter-processor communication time involved in the parallel code. Considering the inter-processor communication separately, one can modify (Dash and Jha 1996) Eqn. (8) to rewrite,

$$SP = \frac{1}{f + (1-f)/p + t_c/t} \tag{9}$$

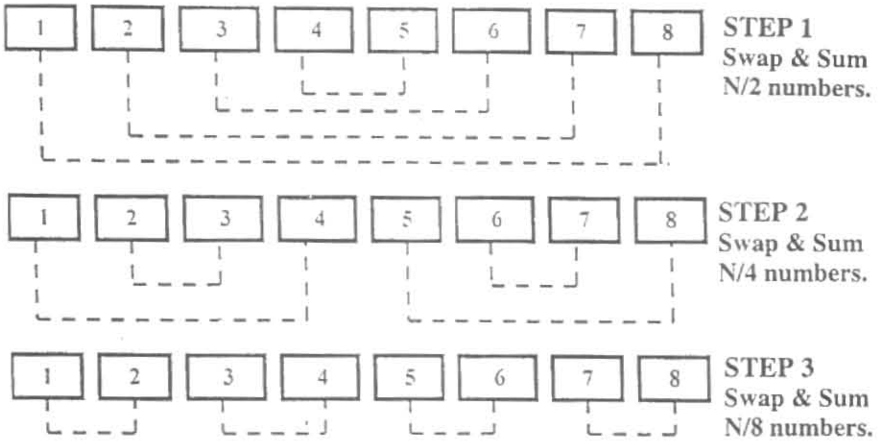

Fig. 3. Vector global sum technique for adding an array of N numbers on 16 processors (Dash 1998)

Here t_c is the communication overhead in p-processor system. The parallel efficiency is defined by,

$$E = SP / p \tag{10}$$

and is usually represented in percentage. The parallel efficiency measures the effectiveness of the implementation of parallelisation on a machine. If the problem size is sufficiently large, in most cases the parallel efficiency is reasonably good. In his summary of the best parallel performance achieved, Ashworth (1993) has shown very impressive values of parallel efficiency of 96% while increasing the number of Intel iPSC/860 processors up to 128. Similarly 16 processors of Cray C90 had 64% parallel efficiency. Ashworth (1993) also computed the machine efficiency and found its value 8.3% and 46% in the above two cases respectively. The machine efficiency is the ratio between the theoretical peak performance and the absolute speed of the machine. The absolute speed is computed by counting the floating point operations present in the code and dividing by the elapsed or wall-clock time. The machine efficiency, thus takes into account how well the code runs on a single processor in addition to the efforts involved in parallelisation.

5. Parallelisation of meteorological models in India

Realizing the importance of parallel computers, the Department of Science and Technology, Government of India has been encouraging various groups in the country for research in paralleisation of weather and climate models through sponsored projects since 1988, when negotiations were going on for acquiring the supercomputer Cray X-MP from USA. The C-DOT High Performance Parallel Processing System (CHIPPS) was built on Single Algorithm Multiple Data (SAMD) architecture exclusively for its use in weather forecasting and radio astronomy (Periasamy 1992, Dash and Misra 1993). The sequential code of the spectral GCM of IIT Delhi (Dash and Chkrapani 1989) has been parallelised at different horizontal resolutions such as T21 (Dash *et al.* 1993a & b, 1996), T42 (Dash *et al.* 1995a), T63 (Dash and Perisamy 1995) and T80 (Dash 1998) and implemented on 16,32,48 and 68 processors of the CHIPPS. The elapsed times for parallel computations and overheads such as communication and sequential calculations have been examined for different resolutions of the model. Results show that as the resolution of the parallel model is enhanced, the size of the spectral coefficients of weather parameters communicated across the processors becomes very large, so much so that it takes about 70% of the total time at resolution T42L15 implemented on 32 processors. Thus the theoretical maximum speedup saturates at about 9 when the number of processors are increased from 16 (in case of T21) to 32 (in case of T42). This result is in accordance with Amdahl (1967). With a view to reduce the inter-processor communication time, the computations and communication were overlapped by adopting the simple technique of global vector sum as illustrated in Fig.(3) for 8 processors. Algorithm (Dash *et al.* 1995b) has been developed so that it is not necessary to transfer the contributions (to spectral coefficients) from all the processors to one processor for addition. These contributions are exchanged suitably amongst the processors so that the addition goes on each processor for a particular group of coefficients depending on the total number of processors used for model integration. Results show that by adopting this technique, the theoretical maximum speedup and parallel efficiency in case of model resolution T63L15 integrated on 48 processors are 42 and 87% respectively. Similarly in case of T80L15 model integration on 64 processors, the theoretical maximum speedup and parallel efficiency are 50 and 78% respectively. The T80L18 model of the National Cenre for Medium Range Weather Forecasting (NCMRWF) has also been implemented on 128 processors of CHIPPS by adopting the domain decomposition method along the latitudes as well as along the vertical levels (2D parallelisation) and the speedup and efficiency are found to be 87 and 68% respectively (Dash 1997).

In 1992, several centres working on various aspects of parallel computing in India participated in the parallelisation of the T80 weather prediction code of NCMRWF (Basu 1998). The parallel computers involved are the CHIPPS developed by the Centre for Development of Telematics (C-DOT), Flosolver developed by National Aerospace Laboratories (NAL), Param developed by the Centre for Development of Advanced Computing (C-DAC) and Anupam developed by Bhabha Atomic Research Centre (BARC). According to Basu (1998), this excerise has been able to produce a sustain to peak speed ratio close to 6% mainly because of slow inter-processor communication.

Meteorological computations have been conducted on Flosolver since 1988. A number of atmospheric models starting from simple zonally symmetric model to 3D climate models have been successfully parallelised and implemented (Sinha and Nanjundiah 1997) on the Flosolver. The sequential code of the T80 model of NCMRWF is not modular enough for its efficient paralleisation and hence Nanjundiah and Sinha (1999) modified the code extensively exploiting the features of FORTRAN 90 and successfully reduced the size of the code. This modified code has been tested on various platforms of IBM SP2, SGI Origin, and Pentium based PCs. Venkatesh *et al.* (1998) examined two parallelisation strategies on four processors of SGI Power Challenge using the parallel code of NCMRWF T80 model. In one case the latitude loops Gloopa and Gloopb were parallelised with sequential part of about 4.74% and in another case, in addition to the above two loops some other subroutines used in the linear spectral domain were also parallelised with sequential part of only 0.34%. They demonstrated that the parallel efficiency achieved on four processors of SGI increased from 81% in the first case to 89% in the second case. They also calculated that in case of 64 processors, the theoretical value of parallel efficiency enhanced from 25% in the first case to 82% in the second case. This study demonstrates that parallelisation of NWP models at higher levels for an example in the spectral domain, will definitely be beneficial while using more number of processors.

C-DAC has parallelised and implemented atmospheric and oceanic models at different resolutions on its range of parallel computers Param 8600, Param 9000AA, Param 9000SS and Param 10000. The Param 10000 consists of 48 nodes containing 192 RISC processors of Ultra Sparc II. Each processor at 300 MHz gives 600 Mflops peak performance. The parallel version of NCMRWF T80 model performs efficiently (Kaginalkar and Purohit 1997) on Multiple Instructions Multiple Data (MIMD) type distributed memory Param computer. The parallel spectral GCM of IIT Delhi at horizontal resolutions of T21 and T42 has also been implemented on 8 processors of Param. The T21 model has been integrated on 8 processors of Param Openframe for three months and the values of theoretical maximum speedup and parallel efficiency are found to be 6.6 and 82.6% respectively (Dash 1998). C-DAC High Performance Computing and Communication (HPCC) software supports an ensemble of workstations, which may be viewed as independent workstations, cluster of workstations or as massively parallel processor system. It provides communication primitives on the ParamNet and Myrinet.

BARC has been actively engaged in building the parallel supercomputer Anupam based on Intel i860 and Digital Alpha microprocessors and a cluster of Pentium PCs using Asynchronous Transmission Mode (ATM) switch for communication. The Anupam system of parallel computers has a highly optimized parallel software environment which is easy to use and also efficient. It has distributed memory and uses message passing techniques for inter-processor communication. The operational NCMRWF T80 code has been successfully implemented on 16 nodes of Anupam. Work is in progress so as to install the operational model on a cluster of Digital Alpha workstations at NCMRWF.

6. Developments in parallelisation techniques

With the availability of multi-processors, almost all the NWP and climate modelling groups in the world have been using more that one processor for running their numerical models. Use of processors up to 16 in parallel is very common. When one goes to massively parallel processors up to 100 or more, problems arise due to the inter-processor communication. In the initial phase of parallel processing the problem of communication was essentially ignored because the scientists were busy in multi-tasking the sequential codes and found it convenient to assign large jobs to a small number of processors. However, the real challenge in the parallel processing is the use of very large number of cost-effective processors.

The inter-processor communication arises mostly because of the use of multiple FFTs in the longitudinal direction and multiple LTs in the latitudinal direction. When all the grid values X on a fixed latitude are in a processor, the FFT can be computed efficiently using single processor algorithms. Similarly when the arrays of X^m are distributed along a longitude, the LTs can be computed efficiently in each processor. Thus the inter-processor communication can be avoided. The transpose method is essentially based on the dynamical distribution of the data along the latitudes during FFTs and along the longitudes during the LTs. This method has been very popular since it avoids the inter-processor communication. When all the quantities stored in grid point space on a N-S pair of latitudes are allocated to the same processor, it is called 1D parallelisation as mentioned earlier. In the transpose method, it is basically 2D parallelisation where all the quantities necessary for the LTs are also present on the same processor. However in the triangular scheme, uniform allocation of quantities for LTs on all the processors is not made. This gives rise to load imbalance. One strategy for achieving a good load balance is to distribute the columns in such a way that the sum of the lengths of the columns X_n^m (Barros and Kauranne 1993) given to every processor is equal as illustrated in Fig. (4) or nearly equal. Besides that, it should be ensured that every processor receives the same number of columns. In order to fulfill these requirements, the first p columns of X_n^m are given in ascending order to processor 1 to p, p being the total number of processors. The next p columns are then allocated to processor p to 1 in reverse order. After the distribution of the first $2p$ columns, every processor has two columns whose lengths add to the same value. The next $2p$ columns are then allocated in the same way, till all the columns are distributed. This type of data distribution allows the LTs in parallel with no need for communication. Barros and Kauranne (1993) investigated the parallelisation of a global spectral shallow water model on parallel computers with distributed memory and by using the transpose method with proper load balance at resolution T127, they obtained efficiency up to 98% on 32 processors. The efficiency of 32 processors while parallelising T63 model was 93%. Thus efficiency enhanced with the increase in resolution of the model and hence the problem size.

To explain load balance further with reference to Fig. (4), let us consider a GCM at a horizontal resolution T15 running on four processors in parallel. The T15 GCM has 24 Gaussian latitudes from pole to pole and hence each processor has the computations of three pairs of latitudes in the grid space. After the FFTs are done, the values of X^m are transposed to the wave (numbers 0 to 15) domain along the vertical columns 1 to 16 and the load balance is maintained by assigning 16+9+8+1=34 spectral computations to the first processor, 15+10+7+2=34 spectral computations to the second processor, 14+11+6+3=34 spectral computations to the third processor and finally 13+12+5+4=34 spectral computations to the fourth processor.

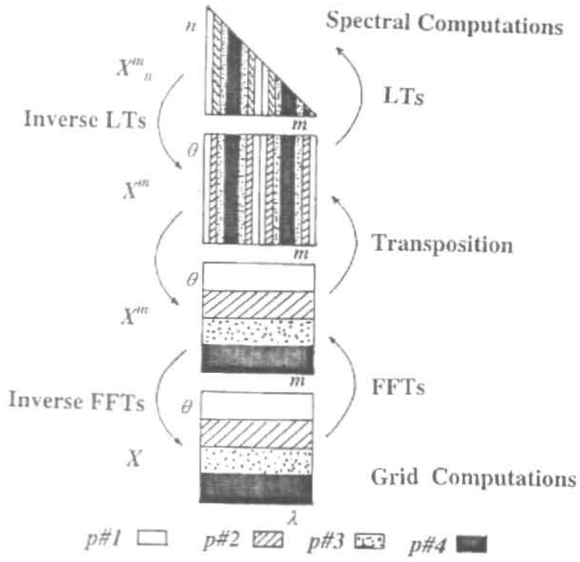

Fig. 4. Transpose method with local balancing technique. p represents the processor number (Barros and Kauranne 1993)

The global and regional spectral models of Florida State University (Christidis and Cocke 1997) have been parallelised on IBM SP distributed memory parallel system using message passing interface (MPI) libraries and the transpose methods. Elapsed times were obtained on 2,4,8,16,24 and 32 processors and the theoretical maximum speedup and parallel efficiency were computed. Results show very good efficiency close to the ideal case when the communication is zero.

The Integrated Forecasting System (IFS) model had been running operationally for many years at ECMWF on multi-processor share memory computer system from Cray Research using PARMACS and F77 features. During the last 4-5 years the IFS model has been running on distributed memory architecture of VPP700 Fujitsu system using MPI and F90 features. Detailed examination of the times taken for computations, load imbalance and message passing overheads for several model resolutions has been made using 46 VPP700 processors. Comparison shows that acceptably low overheads exist for model resolutions T213L31, T106L31 and T63L19. Load balancing efficiency has been maintained within radiation calculations at the expense of more complex code and message passing overheads. Performance measurements show that parallel scalability remains reasonably good even when using more than 40 VPP700 processors (Dent and Mozdzynski 1997).

The HIgh Resolution Limited Area Model (HIRLAM) at resolution 258x200x31 has been parallelised on Cray T3E system using SHMEM software which is Cray proprietary (Eerola *et al.* 1997). The scalability up to 64 processors of T3E does not show any sign of reduction. The data transposition method is found very efficient and the SHMEM software for inter-processor communication shows enhanced efficiency because of less overheads.

It is well known that considerable effort is required to produce efficient parallel codes for weather and climate models on distributed memory machines. It is very difficult to fully automate the paralleliation of the algorithms used in atmospheric models. Hence the job of parallelising existing sequential codes of the meteorological models is very tedious and time taking. Realizing this, a Parallelisation Agent (PA) has been developed which can run on MPI and Parallel Virtual Machine (PVM). The PA approach will facilitate management and evolution of complex codes (Kothari *et al.* 1997). With the advent of powerful and relatively inexpensive PCs, a cluster of PCs can be reasonably good medium for parallel computing with the help of PA.

7. Summary and conclusions

It is well known that the computing requirements for weather and climate models are increasing day by day because of the enhancement of model resolution and sophistications in physical parameterization schemes. The use of multi-model ensemble method for seasonal forecasting and the study of climate variability with the help of high resolution coupled ocean-atmosphere models need enormously large computing power. Most of the present day climate models are run for a large number of years at comparatively low resolution such as T42 and T63 mainly because of the high cost involved in computing. However, climate models should include the atmospheric processes at smaller spatial scales for accurate representation of the local climate changes. Thus meteorological computing, during the last two decades, has been largely based on more than one processor. At present many leading centers of operational forecasting and research in weather and climate modelling are using processors of the order of 16. This trend shows that the days of single processor are over. However, the bottleneck in using more processors is the inter-processor communication, which increases with the square of truncation wave number. Thus a small number of powerful processors have been preferred to a large number of less powerful processors. During the current decade lot of efforts have been going on to reduce the communication time by using higher bandwidth and trying to overlap computations and communication with the help of improved software. This has enabled the meteorological community to use up to 64 processors with reasonably good theoretical maximum speedup and parallel efficiency. Clusters of PCs and workstations are also showing good scalability with the use of improved communication switches and software.

One of the most important issues is that the existing sequential codes themselves are not usually optimized on a single processor. Parallelisation is done by scientists who have not written the sequential codes themselves, and hence it is very difficult on their part to look into all the avenues of inherent parallelism in weather and climate models. It is also very time taking. Some results indicate that by rewriting and optimizing a sequential code on a single processor and then parallelising at microlevels or in other words by resorting to higher levels of parallelism, it is possible to enhance the efficiency of a large number of processors. With improvements in communication technology, development of efficient compilers and optimised algorithms in meteorological models it would be possible to use cost-effective massively parallel processors at least up to a few hundreds in near future.

Acknowledgements

Some of the results discussed in the text are the outcome of research projects sponsored by the Department of Science and Technology, Government of India and All India Council for Technical Education. The author wishes to thank Mr. Pardhan Parth Sarthi for his help in preparing the manuscript.

References

Amdahl, G., 1967, "The validity of the single processor approach to achieving large scale computing capabilities", *AFIPS Conf. Proc, SJCC*, **3**, 483-485.

Ashworth, M., 1993, "Parallel processing in environmental modelling", Parallel supercomputing in atmospheric science, *Eds G-R Hoffmann and T Kauranne, World Scientific*, 1-25.

Barros, S. R. M. and Kauranne, T., 1993, "On the parallelisation of global spectral Eulerian shallow water models", Parallel supercomputing in atmospheric science, *Eds G-R Hoffmann and T Kauranne, World Scientific*, 36-43.

Basu, S. K., 1998, "Usability of parallel processing computers in numerical weather prediction", *Current Science*, **74**, 508-516.

Bengtsson, L., 1984, "Computer requirements for atmospheric modelling", Workshop on using multi-processors in meteorological models, ECMWF, 1-14.

Christidis, Z. and Cocke, S., 1997, "Optimization and parallelisation of the FSU tropical weather forecasting model on the IBM-SP", Proc. Symp. Regional weather predication on parallel computer environments, *Eds. G. Kallos, V. Kotroni and K. Lagouvardos, Unversity of Athens*, 257-265.

Dash, S. K., 1997, "Climate and weather modelling using different parallel computers", Making its Mark, *Eds G-R Hoffmann and N. Kreitz, World Scientific*, 284-289.

Dash, S. K., 1998, "Parallelisation of weather model and performance evaluation on varied number of processors", *Vayu Mandal: Bulletin of IMS*, **28**, 11-17.

Dash, S.K. and Chakrapani, B., 1989, "Simulation of winter circulation over India using a global spectral model", *Proc. Indian Acad. Sci. (Earth and Planet. Sci.)*, **98**, 189-205.

Dash, S. K. and Jha, B., 1996, "A Global numerical weather model integrated on transputer based parallel computer", *Comput. Syst. Sci. and Eng.*, **11**,93-98.

Dash, S. K. and Misra, V. K., 1993, "Parallel computing in numerical weather prediction", Horizons, IEEE Asia- Pacific, Jan-March, 41-44.

Dash, S. K., Jha, B., Selvakumar, S., Periasamy, M., Pitke, M. V. and Misra, V. K., 1993(a), "Implementation of a T21 global spectral model on 16 nodes of parallel computers", Parallel supercomputing in atmospheric science, *Eds G-R Hoffmann and T Kauranne, World Scientific*, 60-72.

Dash, S. K., Jha, B., Selvakumar, S., Periasamy, M., Padmalatha, K. and Misra, V. K., 1993(b), "Modification of T21 global circulation model for implementation on parallel processing system", *Indian J. Radio Space Physics*, **22**, 225-229.

Dash, S. K., Selvakumar, S. and Jha, B. 1995a, "Climate modelling using parallel processors", *Atmos. Env.*, **29**, 2001-2007.

Dash, S. K. and Periasamy, M., 1995, "Performance of a number of processors in running GCM at different resolutions", Coming of Age, *Eds G-R Hoffmann and N. Kreitz, World Scientific*, 424-430.

Dash, S. K., Periasamy, M., Selvakumar, S. and Jha, B., 1995b, "Efficiency of a number of processors in integrating a global weather model", High performance computing in the geo-sciences , Ed F-X. Le Dimet, Kluwer Academic Publishers, 257-261.

Dent, D. and Mozdzynski, G., 1997, "ECMWF operational forecasting on a distributed memory platform: forecast model", Making its Mark, *Eds G-R Hoffmann and N. Kreitz, Orld Scientific*, 36-51.

ECMWF, 1984, "Proc. Workshop on Using Multiprocessors in Meteorological Models", ECMWF, December, p363.

Eerola, K., Salmond, D., Gustafsson, N., Garcia-moya, J.A, Lonnberg, P., Jarvenoja, S., 1997, "A parallel version of the HIRLAM forecast model: strategy and results", Making its Mark, *Eds G-R Hoffmann and N. Kreitz, Orld Scientific*, 134-143.

Eliassen, E., Machenhaur, B. and Rasmussen, E., 1970, "On a numerical method for integration of the hydrodynamical equations with a spectral representation of the horizontal fields", Rep.No.2, Institute for Teoretisk Meteorologi, University of Copenhagen.

Foster, I., Gropp, W. and Stevens, R., 1992, "The parallel scalability of the spectral transform method", *Mon. Wea. Rev.*, **120**, 835-850.

Haltiner, J.G., 1971, "Numerical weather prediction", John Wiley & Sons, Inc. p 317.

Kaginalkar, A. and Purohit, S., 1997, "Benchmarking of medium range weather forecasting model on Param -A parallel machine", Making its Mark, *Eds G-R Hoffmann and N. Kreitz, Orld Scientific*, 461-472.

Kothari, S. C., Kim, Y. and Takle, E. S., 1997, "Parallelisation agent for atmospheric models", Proc. Symp. Regional weather predication on parallel computer environments, Eds. G. Kallos, V. Kotroni and K. Lagouvardos, Unversity of Athens, 287-294.

Kreitz, N. and Prior, P., 1995, "Experience of using parallel computers in meteorology", Coming of Age, *Eds. G-R. Hoffmann and N.Kreitz, World Scientific*, 551-555.

Kreitz, N. and Prior, P., 1997, "Experience of using parallel computing in meteorology", Making its mark, *Eds. G-R. Hoffmann and N. Kreitz, World Scientific*, 478-482.

Nanjundiah, R. S. and Sinha, U. N., 1999, "Impact of modern software engineering practices on the capabilities of an atmospheric general circulation model", *Current Science*, **76**, 1114-1116.

Orszag , S.A., 1970, "Transform method for calculation of vector coupled sums: Application to the spectral form of the vorticity equation", *J. Atmos. Sc.*, **27**, 890-895.

Periasamy, M., 1992, "C-DOT's data parallel computers: an alternative for the 1990s", Horizons, IEEE, Asia-Pacific, April-June, 41-44.

Richardson, L. F., 1922, "Weather prediction by numerical process", Cambridge University Press, p236.

Sinha, U. N. and Nanjundiah, R. S., 1997, "A decade of parallel meteorological computing on the Flosolver", Making its Mark, *Eds G-R Hoffmann and N. Kreitz, Orld Scientific*, 449-460.

Venkatesh, T.N., Sivaramakrishnan, R., Sarasamma, V. R. and Sinha, U. N., 1998, "Scalability of the parallel GCM T80 code", *Current Science*, **75**, 709- 712.