Porting a Limited Area Numerical Weather Forecasting Model on a Scalable Shared Memory Parallel Computer

Roberto Ansaloni[1], Paolo Malfetti[2], and Tiziana Paccagnella[3]

[1] Silicon Graphics SRL, Milano, Italy, Roberto.Ansaloni@sgi.com
[2] CINECA, Casalecchio di Reno, Bologna, Italy, Malfetti@cineca.it
[3] ARPA, Bologna, Italy, T.Paccagnella@smr.arpa.emr.it

Abstract. As new parallel machines become popular commercially, it is important to understand if they are adequate for performing and scaling well for a range of applications. This paper describes the porting of a weather forecast limited area model from a vector machine to a recent SMP-multiprocessors and evaluates the performances obtained.

1 Introduction

Recently, ccNUMA multiprocessors have been shown to deliver good performance on many applications. Weather forecast limited area models run commonly on vector machines or on MPP systems. It is important to know that the ccNUMA systems perform and scale well even for the above mentioned models, and a relatively simple parallelization effort due to the programming model available on these system is needed.

In Section 2 we begin by briefly describing the limited area model used and, in Section 3, the memory and communication architecture of the SGI Origin 2000 (Origin). In Section 4 we address the problem of porting the code from a vector machine, like the Cray C90 (C90), to a scalable shared memory processor machine (SMP), like the Origin, while in Section 5 we evaluate the performances of the code on SMP multiprocessor. Finally, Section 6 briefly explains the operational suite and Section 7 concludes the paper.

2 Model Description

LAMBO, Limited Area Model BOlogna, is a grid-point primitive equations model, based on the 1989 and 1993 versions of the ETA model, used at the National Centre for Environmental Prediction of Washington [2].

LAMBO is running operationally since 1993 at Agenzia Regionale Prevenzione Ambiente - Servizio Meteorologico Regionale (S.M.R.).

As already mentioned before, LAMBO is a grid-point, primitive equations limited-area model: in such models the only basic approximation, well justified by the scale analysis of the vertical component of the momentum equation, is the

P. Amestoy et al. (Eds.): Euro-Par'99, LNCS 1685, pp. 1435–1438, 1999.

hydrostatic approximation, which assumes pressure at any point simply equal to the weight of the unit cross-section column of air above that point. Atmospheric motion is predicted by applying the principles of conservation of momentum, energy and mass and using the law of ideal gases. Such set of differential equations constitutes the initial and boundary value problem, the solution of which provides the future state of the atmosphere.

The equations of motion are solved in practice using finite difference methods and all model variables are defined on the so-called Arakawa E-type grid.

Particular numerical schemes were developed to integrate on the E-grid the part of the equations related to adiabatic processes [1].

LAMBO has a full complement of physical parameterizations. Parameterizations are modules of the model code which attempt to represent the effects on model prognostic (dependent) variables of those processes which cannot be explicitly resolved during model integration. Such processes include in general moist processes, vertical turbulent exchanges, radiative exchanges, lateral diffusion, surface exchanges of moisture, heat and momentum and so on.

3 LAMBO on the SGI Origin 2000

A serial version of LAMBO has been running in production on CINECA C90 since September 1993. The purpose of this work was to implement a parallel version of LAMBO on a Origin, following few basic criteria:

- the parallel version of the code had to run on the 16 Origin processors at least in the same time as the serial C90 version
- the code modifications had to be kept at a minimum level in order to retain code readability and portability

The Origin version has been written in a shared memory programming model, by far the most natural and efficient way to implement parallel code on Origin systems. Parallelism has been achieved by exploiting the auto-parallelizing compiler features and by the insertion of SGI parallel directives.

The migration of the code to the parallel Origin system has been articulated in four major steps: porting, single processor tuning, parallelization, performance analysis.

The porting process has been really straightforward: the only important issue was related to the numerical precision required, due to the different default variable size on C90 (64 bits) and on Origin (32 bits).

In order to efficiently run the LAMBO vector code on the cache-based Origin architecture, aggressive optimization compiler flags had to be turned on, in particular for the loop nesting and cache prefetching analysis.

The major problem arose when considering the code parallelization. The original version of LAMBO made a large use of EQUIVALENCEd variables, to save memory and increase code readability, but the presence of an EQUIVALENCEd variable in a loop inhibits its parallelization. Thus, in order to achieve a significant level of parallelism, it has been necessary to remove most of the EQUIVALENCE statements, thus reducing the code readability for the authors.

Different parallelization schemes have been applied to different subroutines, always choosing the best approach according to the algorithm implemented.

Another advantage given by the programming model chosen, is the possibility to adopt an incremental code parallelization approach.

Eventually it turned out that 10 subroutines have been manually parallelized by the insertion of compiler directives and 6 have been automatically parallelized by the compiler. In the case of the radiation package, the parallelization has been achieved at a higher level, by parallelizing the main loop in the driver routine which calls the other radiation routines.

4 Parallel Performance

In order to evaluate the parallel performance, the results are compared with those foreseen by Amdhal's law which represents the parallel execution time T(p) as a function of the number of processor p and of the parallel fraction fp of the serial time (see Table 1). S(p), the speed-up function definition, is used to evaluate the parallel performance; in the ideal case, when all the code is perfectly parallel (fp=1) , the speed-up function is the linear function S(p)=p. Since only the routines that account for the 96% of the total execution time were considered for parallelization, the Amdhal's curve corresponding to fp=0.95 should be considered. Due to imperfect load balance, cache misses and data contention between processors that fraction is positioned between 90% and 95%. LAMBO performance is in good agreement with those predicted by the Amdhal's law model as shown in Figure 1.

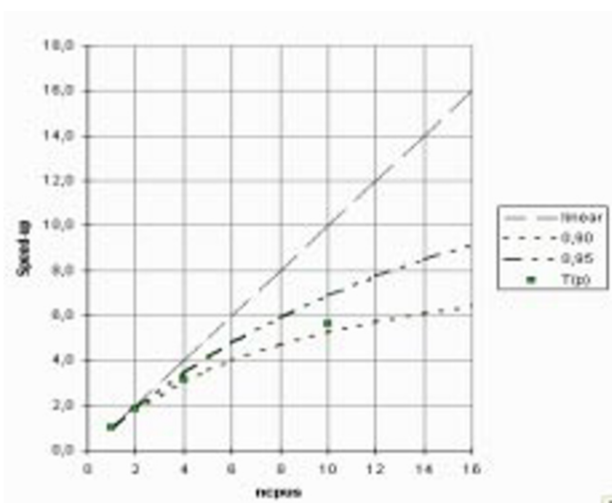

Fig. 1. Linear speedup, theoretical speedup for f_p

Pes	1	2	4	6	8	10	12	14	16
T(p)	275	150	87	66	56	50	46	43	40

Table 1. Parallel execution times on Origin

5 The Operational Suite

LAMBO runs daily at two different horizontal and vertical resolutions. The first run, 40 km horizontal resolution with 20 vertical sigma levels, is performed taking initial and boundary conditions from European Center for Medium range Weather Forecast (E.C.M.W.F.) 00GMT analysis and forecast. The second run, 20 km horizontal resolution with 32 vertical sigma levels, is nested in the previous run using a one way nesting procedure developed at S.M.R.. Both the runs are computed on integration domains covering all the Italian territory. E.C.M.W.F. products required by LAMBO are obtained through a link with the Servizio Meteorologico Aeronautica Militare, the National Weather Service.

After one month of testing phase, LAMBO is operative on CINECA Origin since the 1st July 1998, and, on average over a six months period, on 10 Origin processors, the first run takes about 5' while the second takes about 32'. This should be compared with the 10' and 50', respectively required by the previous C90 runs.

6 Conclusions

In the recent years it has been a common perception that only vector processor machines were adequate for running limited area weather forecast models. Few recent models were developed on MPP systems, others were ported using the message passing paradigm. We have shown that a limited area model on a scalable shared memory parallel computer can give good performances, providing satisfactory operational forecasts.

Moreover it should be stressed that these results have been obtained by a relatively simple parallelization effort due to the programming model available on the Origin system: a much higher effort would have been necessary in the case of a message passing implementation.

References

[1] Janjic, Z., 1979: Forward-backward scheme modified to prevent two-grid interval noise and its application in sigma coordinate model. Contr. Atm. Phys., 52, 69-84.

[2] Mesinger F., and Z. I. Janjic, S. Nickovic, D. Gavrilov and D. G. Deaven, 1988: The step-mountain coordinate: Model description and performance for cases of Alpine lee cyclogenesis and for a case of Appalachian redevelopment. Mon. Wea. Rev., 116, 1493-1518.