Introduction

The program of the Fourth ECMWF Workshop on the Use of Parallel Processors in Meteorology, held at ECMWF in Reading in November 1990, provides further evidence on the increasingly important rôle massively parallel computers are expected to play in weather and climate forecasting. Most operational weather models now routinely run multitasked on up to eight processors on vector supercomputers. However, none of the forecasting centres has yet ventured to implement a full model on a massively parallel computer, i.e. one with hundreds or thousands of processors. Yet precisely this is expected to happen before the end of the century.

The strongest indicator to the emergence of massively parallel supercomputers capable of running operational models has been the recent commitment of most current vector supercomputer vendors to producing massively parallel computers. Also, most major forecasting centres are running experiments on parallel, distributed memory computers with simplified models in order to study the changes new machine architectures will impose on model codes. A major effort (CHAMMP) has been launched in the United States to port climate models onto distributed memory parallel computers, with the aim of attaining performance levels needed to tackle the Grand Challenge problem of climate change.

The doubts in the past about whether in particular spectral weather models will be implementable on massively parallel computers at all because of fundamental hardware limitations seem to have dispersed to a large extent. Even though massively parallel computers are undoubtedly more limited in their capability to access large data structures, careful planning of the communication structure seems to be able to overcome much of the problems caused by this. Similar ideas have already proved crucial for an efficient multitasking implementation of weather models on vector supercomputers.

The main concern among workshop participants appeared to be the lack of programming tools and standards on massively parallel computers. In the present situation, embarking on a major code migration project entails large risks of losing this investment in case the programming method or the parallel computer targeted proves to be non-optimal in the long term. Such risks have proved serious even on vector supercomputers. Parallel computer manufacturers were strongly urged to agree on Fortran based programming standards to protect investments in code migration. This would make massively parallel computers far more attractive to customers.

On the other hand, meteorological users of supercomputers were encouraged to produce portable simple benchmark codes that could be implemented also on massively parallel computers with modest effort. The codes would still have to be complete enough to provide convincing evidence on the performance potential of each parallel computer in running operational models. Since the workshop, both of the above standardization efforts have got well underway.

The workshop provided a balanced collection of talks on implementing weather models on parallel computers and on various programming tools currently available on them. The diverse trends both in hardware and in programming tools were pulled together in two invited survey talks and in the open panel discussion on the last day of the workshop.