Parallel Computational Fluid Dynamics Conference (ParCFD2013)

Performance analysis and optimization of PalaBos on petascale Sunway BlueLight MPP Supercomputer

Tian Min[a], Gu Weidong[a], Pan Jingshan[a], Guo Meng[a]*

[a]*Shandong Provincial Key Laboratory of computer Network, Shandong Computer Science Center, No19 Keyuan Road, Jinan, 250101, PR China*

Abstract

We present some results concerning the computational performances of the open source general purpose CFD code PalaBos, in terms of scalability and efficiency, on the petascale Sunway BlueLight MPP system. Based on the numerical simulated program of 3D cavity lid driven flow, the optimization methods in I/O, communication, memory access, etc, are applied in debugging and optimization of the parallel MPI program. Experimental results of large scalar parallel computing of 3D cavity lid driven flow show that, the parallel strategy and optimization methods are correct and efficient. The parallel implementation scheme is very useful and can shorten the computing time explicitly.

Nomenclature

MPP	Massively Parallel Processor
HPC	High Performance Computing
CFD	Computational Fluid Dynamics
LBM	the lattice Boltzmann method

1. Introduction

As the importance of High Performance Computing (HPC) in Computational Fluid Dynamics (CFD) is increasing, industries are more and more interested in its applications. However, the cost of license of commercial CFD codes is proportional to the number of cores used and thus running large simulations in parallel on multi-core systems maybe economically prohibitive unless open source software like PalaBos is used. By studying computational performance requirement of industrial interest, we found that the performance capabilities of high-end computational resources have increased rapidly over recent years. In particular, the introduction of petascale systems has brought with it massive increases in the number of processing units, where it is now common to have many tens of thousands of cores available for users' codes. This development raises a number of significant challenges for the parallel performance of CFD applications.

As a software tool for classical CFD, particle-based models and complex physical interaction, PalaBos offers a powerful environment for fluid flow simulations based on the lattice Boltzmann method (LBM) [1]. Recently, new parallelization and optimization techniques have been introduced to PalaBos in order to address these challenges at several different stages of

the calculation [2-4]. The introduction of LBM method and the hardware architecture are described in Section 2. In Section 3 we introduce the optimization methods in I/O, communication, memory access, etc. In Section 4, the impacts on scalability and performance of these new features have been analyzed on a range of prototype petascale system.

2. Method and Implementation

2.1. The lattice Boltzmann method

The lattice Boltzmann method is a numerical technique for the simulation of fluid dynamics, and in particular the numerical solution of the incompressible, time-dependent Navier-Stokes equation [5]. Its strength is however based on the ability to easily represent complex physical phenomena, ranging from multiphase flows to chemical interactions between the fluid and the borders. Indeed, the method finds its origin in a molecular description of a fluid and can directly incorporate physical terms stemming from a knowledge of the interaction between molecules. For this reason, it is an invaluable tool in fundamental research, as it keeps the cycle between the elaboration of a theory and the formulation of a corresponding numerical model short.

Compared to other CFD approaches, lattice Boltzmann might at first sight seem quite resource consuming: the discrete probability distribution functions described by the model require more memory for their storage than the hydrodynamic variables used by a classical solver of the Navier-Stokes equation (nine real valued quantities per node against three for 2D incompressible solvers).

This is however never a real issue, especially on modern computers, and it is greatly compensated by an outstanding computational efficiency [6]. Thanks to its explicit formulation and exact advection operator, the lattice Boltzmann scheme involves only a very limited amount of floating points operations per computational node. Furthermore, thanks to the locality of its algorithm, the lattice Boltzmann method is particularly well suited for computations on various parallel architectures, even on those with slow interconnection networks.

The lattice Boltzmann method is a very successful tool for modeling fluids in science and engineering [7]. Compared to traditional Navier Stokes solvers, the method allows an easy implementation of complex boundary conditions and due to the high degree of locality of the algorithm—is well suited for the implementation on parallel supercomputers.

2.2. 3D lid-driven cavity flow

The lid-driven cavity problem has long been used a test or validation case for new codes or new solution methods. Lid-driven cavity flows are not only technologically important, but also they are of great scientific interest. These flows display many kinds of fluid mechanical phenomena, including corner eddies, Taylor-Gortler-like (TGL) vortices, transition, turbulence and so on. Simple geometrical settings and easily posed boundary conditions have made cavity flows become popular test cases for computational schemes.

As a classic benchmark, the 2D lid-driven cavity flows have been extensively studied with numerical methods. However, the pioneering experimental work of Koseff & Street and coworkers in the early 1980s clearly showed that cavity flows were inherently 3D in nature. With the increase of computing capability in recent years, the 3D lid-driven cavity problems have matured as a standard Re-dependent benchmark. This problem has been solved as both a laminar flow and a turbulent flow, and many different numerical techniques have been used to compute these solutions. Since this case has been solved many times, there is a great deal of data to compare with.

2.3. Hardware architecture

Sunway BlueLight MPP Supercomputer is the first publicly announced PFLOPS supercomputer using ShenWei processors solely developed by the People's Republic of China. It ranked #2 in the 2011 China HPC Top100, #14 on the November 2011 TOP500 list, and #39 on the November 2011 Green500 List. The machine was installed at National Supercomputing Jinan Center in September 2011 and was developed by National Parallel Computer Engineering Technology Research Center and supported by Technology Department 863 project. The water-cooled 9-rack system has 8704 ShenWei SW1600 processors (For the Top100 run 8575 CPUs were used, at 975 MHz each) organized as 34 super nodes (each consisting of 256 compute nodes), 150 TB main memory, 2 PB external storage, peak performance of 1.07016 PFLOPS, sustained performance of 795.9 TFLOPS, LINPACK efficiency 74.37%, and total power consumption 1074 kW.