

AN EXPLICIT MULTIGRID SCHEME USING ARTIFICIAL COMPRESSIBILITY METHOD FOR THE SIMULATION OF UNSTEADY INCOMPRESSIBLE FLOWS ON MULTI-GPU CLUSTER

XIAO-LEI SHI ^{*}, CHAO-AN LIN [†]

^{*}[†] National Tsing Hua University (NTHU)
Department of Power Mechanical Engineering
30013 Hsinchu, Taiwan

^{*}E-mail: sx117@mails.tsinghua.edu.cn

[†]E-mail: calin@pme.nthu.edu.tw

Key words: Multigrid method, artificial compressibility method, Multi-GPU, Navier-Stokes equations; incompressible flows, Lid-driven cavity flow.

Summary. An explicit solver for unsteady incompressible Navier-Stokes equations working on multi-GPU cluster is presented, which is based on artificial compressibility method. An explicit FAS multigrid method is developed to accelerate the numerical convergence within each time step and 109 times maximum speedups are reported for the test cases. Fine-grained overlapping strategies for the multi-GPU code optimization are designed and more than 192 times accelerations (single grid) for the GPU code compared with one core of its CPU counterpart is obtained. A typical performance of this solvers is ~0.68s per time-marching using a resolution of 192^3 grids on 4 NVIDIA Tesla P100 (NVLink) GPUs. The numerical results for both steady and unsteady 3-D lid-driven cavity flows including turbulence statistical results for $Re = 3200$ are presented, which shows good agreement with experimental measurement. At last, the computation cost of artificial compressibility method and performance of FAS multigrid method are introduced.

1 NUMERICAL METHOD

The time-dependent, three-dimensional artificial compressibility equations can be written as:

$$\frac{1}{\beta} \frac{\partial p}{\partial \tau} + \frac{\partial u_i}{\partial x_i} = 0 \quad (1)$$

$$\frac{\partial u_i}{\partial \tau} + \frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} \quad (2)$$

where β is artificial compressibility coefficient, p is the kinematic pressure, ν is the kinematic viscosity, τ and t are pseudo and physical time respectively.

The governing equations are discretized on a staggered grid by finite-volume approach and all spatial derivatives are approximated using second-order central difference schemes. The discrete equations are integrated in time using a second order accurate, dual-time-stepping artificial compressibility approach¹. Euler forward differencing is adopted to integrate the equations in dual-time, such that:

$$\frac{\phi_{i,j,k}^{l+1,n+1} - \phi_{i,j,k}^{l,n+1}}{\Delta\tau} + \frac{3\phi_{i,j,k}^{l+1,n+1} - 4\phi_{i,j,k}^n + \phi_{i,j,k}^{n-1}}{2\Delta t} + R^*(\phi)_{i,j,k}^{l,n+1} = 0 \quad (3)$$

$$\frac{p_{i,j,k}^{l+1,n+1} - p_{i,j,k}^{l,n+1}}{\Delta\tau} + (\nabla \cdot \mathbf{V})_{i,j,k}^{l,n+1} = 0 \quad (4)$$

where ϕ donates the velocity components u , v and w in the x -, y - and z -directions respectively in Cartesian coordinate, R^* represents the residual in momentum equations (2), $\Delta\tau$ and Δt are the dual-time and physical time increasement respectively.

In this work, we have adopted the full-approximation storage (FAS) multigrid scheme to accelerate the convergence of Eqn. (3) and (4) in each physical time step. This is a non-linear version of the multigrid method² and we schedule the grids using classical V-cycle³.

2 MULTI-GPU IMPLEMENTATION AND PERFORMANCE

A two-dimensional domain decomposition along y - and z -direction was adopted in the current GPU implementation, which is the so called ‘‘pencil’’ decomposition. We rely on the CUDA-aware Message Passing Interface (MPI) to do inter-GPU communication. In this work we developed fine-grained overlapping strategies to optimize the primary function blocks for maximizing our code efficiency.

The performance of GPU kernels is examined by means of the roofline model⁴ as demonstrated in Fig. 1 (a). In order to compare the performance of GPU and its CPU counterpart, we execute the multigrid, level 1 to 4, V-cycle 25000, 5000, 2000, 1000 times for grid number 64^3 , 128^3 , 192^3 and 256^3 respectively, results shown in Fig. 1 (b).

The scalability of current GPU implementation on multiple Tesla P100 (NVLink) GPUs is reported in Fig. 2(a).

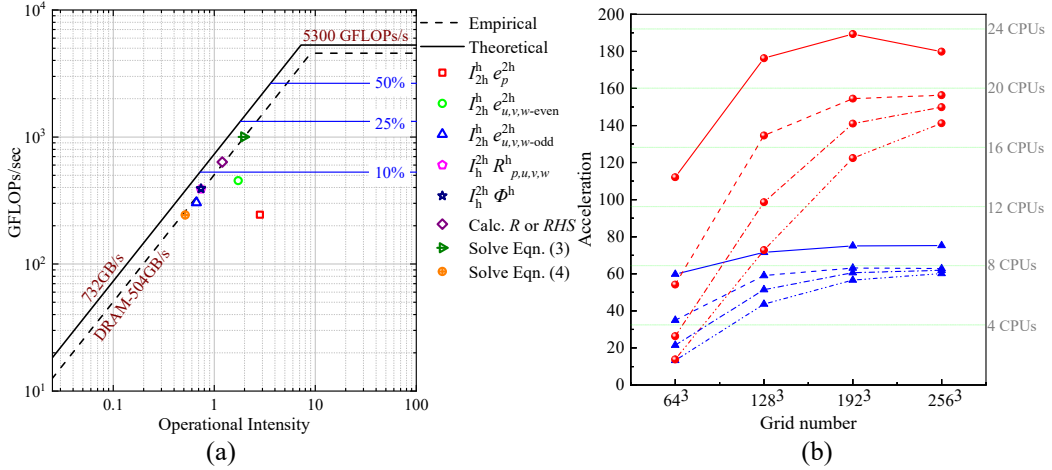


Fig. 1. (a) Roofline on Tesla P100 GPU with measured performances of each GPU kernel. (b) The acceleration of a NVIDIA® GPU compared with one core of Intel® Core™ i7-6900K CPU. ●: NVIDIA® Tesla® P100 GPU (PCI-e), ▲: NVIDIA® GeForce® GTX TITAN GPU; —: Single grid, - - - - -: FAS Lev. 2; - · - · - · -: FAS Lev. 3; · · · · · -: FAS Lev. 4; —: the number of Core™ i7-6900K CPU to yield the same performance with a single GPU (Acceleration/# of cores per CPU).

3 NUMERICAL RESULTS

The numerical results of three-dimensional lid-driven cavity flow at a Reynolds number of 3200 with a resolution of 192^3 is presented. The grid is symmetry clustered to the wall in order to capture the subtle flow structures. Fig. 2(b) shows the predicted mean and turbulence statistics, which correlate well with the measurements.

4 CONCLUSIONS

An explicit solver for unsteady incompressible Navier-Stokes equations working on multi-GPU cluster is developed, which is based on ACM. The governing equations are discretized in space using a second-central difference scheme on a staggered grid and dual-time stepping approach for the time derivatives, yields an implicit scheme. Computations were accelerated using an explicit FAS multigrid scheme up to four levels and up to 109 times speedups is reported for three dimensional LDC flows. Fine-grained overlapping strategies for the multi-GPU code optimization are designed, and we achieve more than 192 times accelerations (single grid) for the GPU code compared with one core of its CPU counterpart.




Fig. 2. (a) Scaling of multi-GPU implication on Tesla P100 GPU. Test case is lid-driven cavity flow with a resolution $384 \times 384 \times 512$ in x -, y - and z -direction. (b) Mean, rms and $u'w'$ profiles along the horizontal and vertical centerlines in the symmetry plane. The symbols are the experimental result of Prasad and Koseff⁵

REFERENCES

- [1] Soh WY, Goodrich JW. Unsteady solution of incompressible Navier-Stokes equations. *Journal of Computational Physics*. 1988;79:113-34.
- [2] Brandt A. Multilevel adaptive computations in fluid dynamics. *AIAA Journal*. 1980;18:1165-72.
- [3] Brandt A, Livne OE. *Multigrid techniques: 1984 guide with applications to fluid dynamics*, revised edition: Society for Industrial and Applied Mathematics; 2011.
- [4] Lo YJ, Williams S, Van Straalen B, Ligocki TJ, Cordery MJ, Wright NJ, et al. *Roofline Model Toolkit: A practical tool for architectural and program analysis*. Cham: Springer International Publishing; 2015. p. 129-48.
- [5] Prasad AK, Koseff JR. Reynolds number and end-wall effects on a lid-driven cavity flow. *Physics of Fluids A: Fluid Dynamics*. 1989;1:208-18.