Computers & Fluids xxx (2010) xxx–xxx

Contents lists available at ScienceDirect

Computers & Fluids

journal homepage: www.elsevier.com/locate/compfluid

From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation

Yuri Feldman *, Alexander Yu. Gelfgat

School of Mechanical Engineering, Faculty of Engineering, Tel-Aviv University, Ramat Aviv 69978, Israel

ARTICLE INFO

Article history:
Received 20 April 2010
Received in revised form 30 June 2010
Accepted 11 August 2010
Available online xxxx

Keywords:
Multigrid
Finite volume
Lid-driven flow
Incompressible flow
Benchmark data
High perfrmance computing

ABSTRACT

A parallel implementation of a fully pressure–velocity coupled multigrid solver based on analytical solution accelerated coupled line Gauss Seidel (ASA-CLGS) smoother with grid partitioning is carried out. The parallelized algorithm is characterized by an enhanced scalability that results from a formulation enabling an intermediate analytical solution for the entire row (column) of control volumes. General strategies of applying single- or multigrid approach depending on flow characteristics are discussed. Performance of the parallelized algorithm is studied for up to 2048 processors. The developed approach is applied to analysis of a time-dependent three-dimensional incompressible lid-driven cavity flow. The steady state results of benchmark quality are reported for $Re = 10^3$, 1.5×10^3 and 1.9×10^3. A new benchmark case of a fully 3D flow in a cubic cavity driven by the lid moving at 45° relatively to its lateral boundaries is proposed and the corresponding data is reported.

1. Introduction

This note presents a study of scalability properties of a fully pressure–velocity coupled CFD solver that implements the multigrid approach with the ASA-CLGS smoother developed in [1]. The original sequential version of the solver was successfully verified on a time-periodic supercritical flow in a differentially heated two-dimensional cavity of height/length ratio 8, and steady flow inside three-dimensional cubic box. Then the parallelized algorithm was implemented for the study of the steady – oscillatory transition of 3D lid-driven flow in a cube [2]. The analysis was performed on massively parallel platform and involved up to 2048 CPU cores. This study was primarily dedicated to the physical problem leaving out of scope examination of the scalability characteristics of the developed solver. The latter is reported in the present paper along with the description of some important properties of the applied numerical methodology. It also outlines general solution strategies, which are applicable for a wide range of various CFD approaches and are not restricted to specific examples reported here.

Commonly used multigrid methods employ relaxation schemes acting upon local data (e.g., Jacobi or Gauss Seidel iteration) whose single-grid scalability is well studied [3]. However when they are used as an internal multigrid iteration which is parallelized in a telescope manner by the grid partitioning approach, the scalability of the whole algorithm is expected to degrade [4]. This is due to increase of communication/computation ratio as the multigrid algorithm descends to the coarser levels while the main part of the communication time is spent in updating overlapping regions. Degani and Fox [5] were the first who analytically estimated the uppermost limit of a three-dimensional multi/single grid efficiency degradation factor. Then a number of studies were performed to investigate scalability of the grid partition approach applied to multiple semi-coarsening algorithms [6,7], to segregated methods where algorithms employing various pressure–velocity decoupling approaches were used [8–10], as well as to pressure–velocity coupled algorithms [11] with Vanka type smoother [12]. A more detailed overview of parallelization of geometric multigrid methods applied to computational fluid dynamics can be found in [13]. An attractiveness of grid partition parallelization with Vanka type smoother was outlined recently by Manservisi [14] who formally proved its monotonic convergence for the steady Stokes problem.

To examine the scalability properties of the parallelized ASA-CLGS multigrid CFD solver [1,2] we consider a classical benchmark problem of lid-driven flow in a cubic cavity. To make the three-dimensionality of the flow equally emphasized along every spatial direction, we alter the classical formulation and consider the lid moving along the diagonal, so that its velocity forms the angle of 45° with the side wall. Following the results in [2], benchmark-quality data is extracted by applying Richardson extrapolation to the results obtained on two finest grids.

* Corresponding author.
E-mail address: yurifeld@post.tau.ac.il (Yu. Feldman).

0045-7930/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.compfluid.2010.08.009

Please cite this article in press as: Feldman Yu, Gelfgat AY. From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation. Comput Fluids (2010), doi:10.1016/j.compfluid.2010.08.009

2. Governing equations

A cubic lid-driven cavity with the side of length L is considered. The upper boundary moves with a constant velocity U while all other boundaries are stationary. The no-slip boundary conditions are applied on all the boundaries. The flow is described by the continuity and momentum equations

$$\nabla \cdot \mathbf{u} = 0 \tag{1}$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \frac{1}{Re}\nabla^2 \mathbf{u} \tag{2}$$

where dimensionless variables are velocity $\boldsymbol{u} = (u, v, w)$, pressure p, and time t. The equations are rendered dimensionless using the scales L, U, $t = L/U$ and $P = \rho U^2$ for length, velocity, time and pressure, respectively. Here ρ is the fluid density. The Reynolds number is defined by $Re = UL/\nu$, where ν is the kinematic viscosity. The flow region is defined in the dimensionless Cartesian coordinates x, y, and z each of which varies between -0.5 and 0.5.

3. Parallelization technique

Parallelization of the ASA-CLGS multigrid algorithm is performed by utilizing a standard telescoping approach, i.e. the number of unknowns per level reduces as the level number is reduced. Taking advantage of the developed ASA-CLGS smoother [1] based on the intermediate analytical solution for the entire column (row) of finite volumes we utilize the two-dimensional virtual topology for the three-dimensional configuration. Therefore, at a specific grid level each CPU is responsible for its sub-domain consisting of $L \times M$ columns in x and z directions, respectively, as shown in Fig. 1. The parallel implementation of the smoother, PAR-ASA-CLGS, is developed for distributed memory machines using MPI approach. Developing the parallel version of the algorithm we tried to minimize the difference from its scalar version.

For simplicity, only a top view of a typical segment of computational domain distributed on a specific processor is shown in Fig. 2. This implies that all mathematical and data communication operations performed on a specific finite volume cell are performed on the corresponding three-dimensional column of finite volumes.

Due to the fact that ASA-CLGS smoother is based on a block relaxation concept applied to unknowns distributed on a staggered grid, the i, $i+1$ and j, $j+1$ velocities are updated simultaneously during a single relaxation sweep. Therefore for a given spatial location a standard finite volume discretization of convective terms [15] entering into Eq. (2) depends on velocity values located at the grid points $i, i \pm 1, i \pm 2, j, j \pm 1, j \pm 2$ in x and z directions respectively. Thus, a parallel implementation of the ASA-CLGS smoother utilizing Gauss–Seidel iteration cannot be constructed using usual zebra ordering. To avoid these dependences we use the three-plane smoother approach as was proposed by Oosterlee et al. [16]. Implied by the two-dimensional virtual topology it restricts the minimal redistributed sub-domain to the size of 3×3 in x and z directions respectively. It also determines the topology of the overlapping region as a band of a single finite volume width attached to the sub-domain along its whole perimeter (see Fig. 2) and defines the maximal number of cores involved in the multigrid mode, which is strictly bounded by the coarsest grid resolution. Consider, for example, a typical multigrid configuration consisting of four subsequent grids with 25, 50,100 and 200 nodes in each direction. In this case a maximal number of the sub-domains positioned in the each direction is $25{:}3 \approx 8$ and therefore the maximal number of cores that can be involved is $8^2 = 64$. Use of larger number of cores will lead to significantly off balance work load and communications between CPUs. In contrast, for the single grid $200{:}3 \approx 66$ that allows one to utilize simultaneously up to $66^2 = 4356$ cores.

For a given distributed sub-domain a relaxation sweep at a specific grid level is propagated by a zigzag (dotted line in Fig. 2), which results from the double loop in the i- and j-directions. In one relaxation sweep of ASA-CLGS the pressure is updated once and the velocity twice. The data communication between overlapping regions of the neighboring processors is implemented by using synchronous send-and-receive commands and is carried out each time the corresponding value of pressure or velocity lying in the overlapping region has been changed. Note also that the data communication is performed explicitly between all neighbor processors including the diagonal neighbors. Such conservative data-communication approach is essential for two main reasons. The first one is the proper implementation of the Gauss–Seidel iteration, which implies the use of all the updated values for the forthcoming computations. The second reason is related to the zigzag propagation of the relaxation sweep responsible for the need of explicit data communication between the diagonal located sub-domains. It should be also noted that even a minor difference between the dimensions of the distributed sub-domains on the coarsest grid may finally result in a significant difference between the corresponding sub-domain dimensions on the finest grid. This is because of utilizing telescoping multigrid approach with doubling of number of grid points at each subsequent grid level. Therefore, use of the blocking communication approach is mandatory to ensure a proper Gauss–Seidel process, and therefore good convergence of the parallelized algorithm.

Another important parallelization issue relates to the restriction and prolongation operators acting on the distributed sub-domains. Using Brandt's [17] correction scheme (CS) the data restricted to

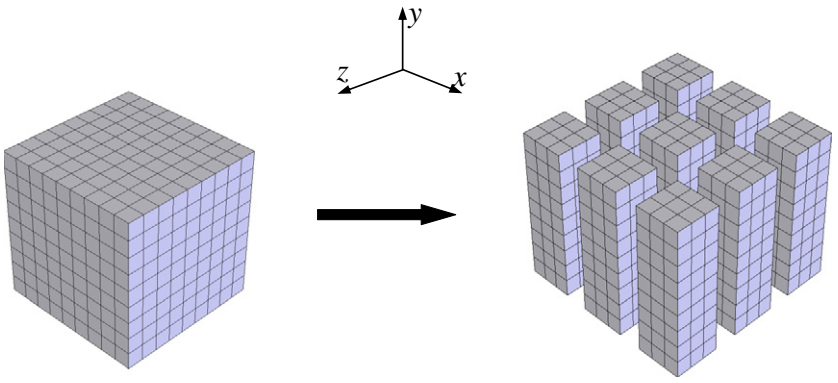

Fig. 1. The whole and the distributed computational domain.

Please cite this article in press as: Feldman Yu, Gelfgat AY. From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation. Comput Fluids (2010), doi:10.1016/j.compfluid.2010.08.009

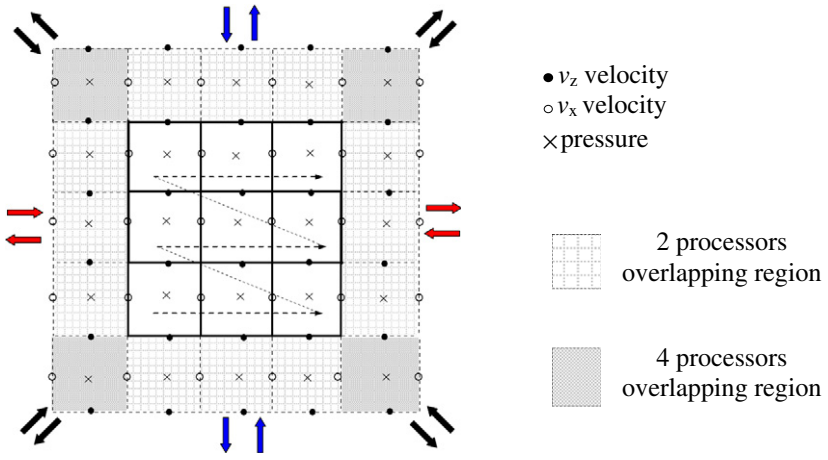

Fig. 2. A top view of typical distributed sub-domain with its overlapping region.

the coarser grid level contains only residual values, while all initial values of pressure and velocity corrections are set to zero. Thus, the data from the overlapping regions is involved implicitly via the residual updating process. In contrast, the prolongation process utilizes explicitly the pressure and the velocity values possessing to the overlapping regions. Furthermore, the prolongation operator acting at the sub-domain corners uses the pressure and the velocity values that are not involved in the relaxation procedure, but still must be updated by an explicit communication between the corresponding neighbor sub-domains. Therefore, data communication between the overlapping regions should be performed at each relaxation sweep.

Fig. 3 demonstrates the scalability properties of the parallelized algorithm estimated both for multi- and single-grid approaches. The multigrid approach is characterized by a speed up proportional to $Np/\ln(Np)$ while a speed up of the single-grid approach is proportional to $Np/\log_{10}(Np)$ as was also stated by Douglas [4]. According to this estimate the single-grid scalability is $\ln(Np)/\log_{10}(Np) \approx \ln(10) \approx 2.3$ times more efficient than that of the multigrid. Our numerical experiments show that single-grid/multigrid scalability ratio grows from 1.9 for 16 CPUs to 2.7 for 64 CPUs and continue to grow with further increase of the CPUs number. This is due to the relative increase of number of communications performed on the coarser grids, as well as due to additional CPU time consumed for the restriction and prolongation procedures. However, if there are large changes in the velocity and/or pressure in consecutive time steps, a large number of single-grid relaxation sweeps will be needed to attain the required accuracy. This may obliterate the better efficiency of single-grid approach resulting in very long computational times in spite of the use of a large number of CPUs working in parallel. A representative example of such a scenario may be the case when a fully developed non-linear flow regime is calculated by rapidly increasing the Reynolds number. In such situations the multigrid approach is preferable. A relatively low number of involved cores whose number is restricted by the resolution of the coarsest grid will be compensated by a high convergence rate of the multigrid. Once a stable asymptotic regime, i.e., steady or oscillatory flow state is approached, the differences between two consecutive time steps for all variables become relatively small, so that further calculations can be performed on the finest grid only, thus, utilizing the maximal possible number of CPU cores in the most effective way.

4. Numerical studies

The developed parallelized algorithm was verified by calculation of a steady state lid-driven flow in a cube at $Re = 10^3$. The steady state flow was obtained by time integration using stretched staggered grids with equal number of finite volumes in each direction that was either 152 or 200. It was assumed that the steady state solution is attained when the maximum point-wise absolute difference between two consecutive time steps is less than 10^{-7}. The zero-grid-step limit of the results was then calculated using the Richardson extrapolation for the both grids. The results are in a good comparison with the benchmark solutions of Albensoeder and Kuhlmann [18], as shown in Table 1. A good agreement is observed for the entire range of the center line velocity and pressure values, which verifies the present calculations.

We complement the benchmark case of [18] by two additional steady state solutions computed for larger Reynolds numbers,

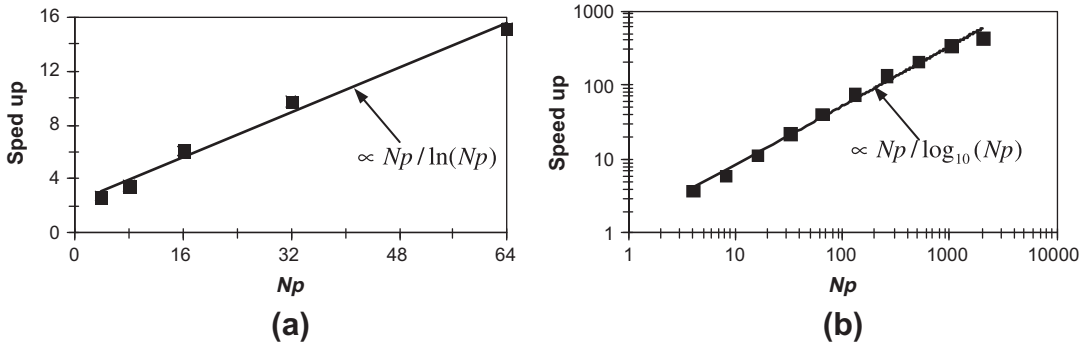

Fig. 3. Scalability of the parallelized algorithm based on version PAR-ASA-CLGS smoother: (a) multigrid approach containing four telescopic grids; (b) single-grid approach.

Please cite this article in press as: Feldman Yu, Gelfgat AY. From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation. Comput Fluids (2010), doi:10.1016/j.compfluid.2010.08.009

Table 1
Pressure and velocity values along the centerlines (0, y, 0) and (x, 0, 0): comparison between the reference [18] and present (Richardson extrapolation from 152^3 and 200^3 grids) solution.

$v_x \times 10^3$			$p \times 10^4$			$v_y \times 10^3$		$p \times 10^3$	
y	[18]	Present	[18]	Present	x	[18]	Present	[18]	Present
0.5	10^3	10^3	86.45	86.25	0.5	0	0	16.58	16.61
0.4766	589.6	589.5	79.85	79.82	0.4688	−188.6	−188.7	17.52	17.52
0.4688	484.4	484.6	77.55	77.52	0.4609	−240.9	−241.0	17.54	17.55
0.4609	398.2	397.9	75.31	75.28	0.4531	−290.3	−290.2	17.40	17.39
0.4531	331.7	332.0	73.22	73.20	0.4453	−335.1	−335.1	17.05	17.04
0.3516	121.8	121.8	51.61	51.61	0.4063	−434.2	−434.2	12.30	12.29
0.2344	73.34	73.30	25.75	25.76	0.3594	−311.2	−311.2	5.326	5.318
0.1172	39.05	39.02	6.760	6.782	0.3047	−152.2	−152.3	1.560	1.556
0	8.018	8.021	0	0	0	36.74	36.74	0	0
−0.0469	−6.119	−6.131	1.250	1.259	−0.2656	169.9	169.9	8.512	8.511
−0.2187	−110.0	−110.0	41.71	41.72	−0.2734	175.8	175.8	8.976	8.972
−0.3281	−251.6	−251.6	161.8	161.7	−0.3437	229.2	229.2	14.04	14.04
−0.3984	−272.9	−272.9	279.2	279.1	−0.4062	244.1	244.0	18.96	18.96
−0.4297	−237.0	−236.9	315.1	315.0	−0.4219	235.0	235.0	19.93	19.93
−0.4375	−222.8	−222.8	321.2	321.1	−0.4297	227.5	227.4	20.34	20.33
−0.4453	−206.2	−206.2	326.2	326.1	−0.4375	217.4	217.4	20.69	20.69
−0.5	0	0	334.0	334.0	−0.5	0	0	21.76	21.76

1.5×10^3 and 1.9×10^3, beyond which the flow becomes oscillatory unstable. A series of time-dependent computations on different grids allowed us to estimate the critical Reynolds number of the steady – oscillatory transition as $Re_{cr} \approx 1914$ [2]. The steady state stopping criterion was exactly the same as in the previous case. Tables 2 and 3 present quantitative distribution of pressure and velocity components along two centerlines of the cubic cavity. All the reported velocity and pressure values obtained on the 152^3, 200^3 grids coincide to within 1%, thus establishing grid independence of the results. We believe that the Richardson extrapolation of the obtained results yields benchmark-quality data as it was already shown for $Re = 10^3$. The characteristic (wall-clock) computation time needed for calculation of a steady state solution at $Re = 1.9 \times 10^3$ on 200^3 grid employing 2048 cores in single grid mode, taking the steady state at $Re = 10^3$ as an initial condition, was about 3 h.

Another case that in our opinion should be added to the standard cubic lid-driven cavity benchmark [18] is a steady flow driven cavity by the lid moving at 45° to the cube wall (see Fig. 4). This flow has no any two-dimensional similarities, and the velocity components are equally large in x- and z-directions. This flow exhibits interesting and not yet studied three-dimensional flow patterns, as is shown in Fig. 4 for $Re = 10^3$. This problem is reflection symmetric with respect to the plane passing through the opposite edges (−0.5, y, −0.5) and (0.5, y, 0.5). As can be seen from Fig. 4 the flow pattern is fully three-dimensional and is characterized by one primary and one secondary downstream eddy whose centers are located on the symmetry plane. It is noteworthy that in this case there is no evidence of the secondary upstream eddy observed in the "classical" lid-driven cavity model. The main vortex is elongated in the direction of moving lid, while the secondary downstream eddy is located close to the cavity bottom corner. Table 4 reports velocity and pressure values along the cavity centerline (0, y, 0). Note that due to the symmetry v_x and v_z velocity components are equal. Another consequence of the symmetry is the fact that fluid particles located on the symmetry plane never leave it. At the same time, the particles located close to the symmetry plane near the centers of the main and secondary eddies attain a significant motion in a direction perpendicular to it, moving along a spiral. In this case the time necessary to calculate the steady state solution on 200^3 grid using 2048 cores is about 5 h, in single grid mode taking the steady state at $Re = 100$ as the initial condition.

Table 2
Pressure and v_x velocity along centerline (0, y, 0) obtained on 152^3, 200^3 grids and their Richardson extrapolation.

y	$Re = 1500$						$Re = 1900$					
	$v_x \times 10^3$			$p \times 10^5$			$v_x \times 10^3$			$p \times 10^5$		
	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.
0.5	10^3	10^3	10^3	627.7	628.3	629.1	10^3	10^3	10^3	519.3	519.9	520.8
0.4766	528.0	528.1	528.2	583.5	583.7	584.0	488.6	488.7	488.9	482.3	482.5	482.8
0.4688	420.4	420.8	421.4	568.1	568.3	568.7	381.9	382.5	383.2	470.1	470.4	470.7
0.4609	339.5	339.5	339.5	553.9	554.1	554.4	306.1	306.2	306.3	459.4	459.6	459.9
0.4531	281.2	281.6	282.1	541.2	541.5	541.9	254.3	254.8	255.4	450.1	450.4	450.8
0.3516	118.3	118.5	118.6	394.8	395.1	395.6	117.1	117.3	117.5	330.0	330.3	330.8
0.2344	68.76	68.84	68.94	191.6	191.8	192.1	67.59	67.69	67.84	154.8	154.9	155.1
0.1172	33.89	33.91	33.95	48.84	48.89	48.96	31.27	31.30	31.34	36.50	36.51	36.52
0	7.188	7.206	7.229	0	0	0	5.888	5.892	5.898	0	0	0
−0.0469	−2.950	−2.931	−2.907	4.830	4.930	5.063	−2.871	−2.866	−2.860	4.650	4.750	4.883
−0.2187	−68.83	−68.76	−68.66	171.1	170.7	170.3	−53.54	−53.47	−53.39	116.0	116.0	115.9
−0.3281	−211.3	−211.4	−211.6	775.3	7.74.0	772.1	−183.7	−183.7	−183.7	487.9	486.6	484.9
−0.3984	−293.1	−293.6	−294.3	1770	1771	1774	−294.8	−295.3	−296.1	1318	1319	1320
−0.4297	−272.2	−272.5	−272.9	2177	2180	2185	−278.1	−278.5	−279.0	1720	1723	1726
−0.4375	−259.5	−259.8	−260.2	2253	2257	2261	−265.2	−265.6	−266.0	1797	1800	1804
−0.4453	−243.4	−243.6	−243.9	2316	2320	2325	−248.7	−249.0	−249.3	1861	1864	1869
−0.5	0	0	0	2419	2423	2428	0	0	0	1967	1970	1974

Please cite this article in press as: Feldman Yu, Gelfgat AY. From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation. Comput Fluids (2010), doi:10.1016/j.compfluid.2010.08.009

Table 3
Pressure and v_y velocity along centerline $(x, 0, 0)$ obtained on 152^3, 200^3 grids and their Richardson extrapolation.

x	$Re = 1500$						$Re = 1900$					
	$v_y \times 10^3$			$p \times 10^4$			$v_y \times 10^3$			$p \times 10^4$		
	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.
0.5	0	0	0	124.0	124.2	124.4	0	0	0	101.3	101.4	101.5
0.4688	−248.4	−248.7	−249.0	129.2	129.4	129.7	−287.6	−288.0	−288.4	104.2	104.3	104.5
0.4609	−310.8	−311.2	−311.6	127.2	127.5	127.8	−352.7	−353.3	−354.0	100.9	101.1	101.3
0.4531	−364.3	−364.4	−364.7	128.5	123.3	116.3	−402.8	−403.1	−403.5	95.32	95.44	95.60
0.4453	−405.4	−405.7	−406.2	116.6	116.8	117.0	−434.0	−434.5	−435.2	87.51	87.62	87.76
0.4063	−399.0	−399.5	−400.2	64.91	64.88	64.83	−346.0	−346.4	−346.9	40.67	40.57	40.44
0.3594	−209.2	−209.1	−209.0	24.03	23.96	23.85	−157.5	−157.3	−157.1	14.81	14.77	14.71
0.3047	−98.41	−98.35	−98.28	8.806	8.800	8.792	−79.42	−79.44	−79.46	5.707	5.703	5.698
0	28.76	28.78	28.81	0	0.0	0.0	27.23	27.28	27.35	0	0	0
−0.2656	137.4	137.4	137.3	51.64	51.60	51.54	115.7	115.6	115.5	38.98	38.96	38.94
−0.2734	143.7	143.6	143.6	54.79	54.74	54.66	121.5	121.5	121.3	41.32	41.29	41.24
−0.3437	211.4	211.6	211.9	94.00	93.99	93.97	194.5	194.7	194.9	72.45	72.43	72.40
−0.4062	250.4	250.9	251.5	142.7	142.9	143.0	248.5	249.0	249.8	118.3	118.4	118.5
−0.4219	246.9	247.3	247.9	154.0	154.1	154.3	245.0	245.6	246.3	129.3	129.5	129.7
−0.4297	241.9	242.3	242.9	158.9	159.1	159.3	239.5	240.1	240.8	134.2	134.3	134.5
−0.4375	234.3	234.8	235.3	163.3	163.5	163.8	231.4	231.9	232.7	138.4	138.6	138.9
−0.5	0	0	0	177.1	177.4	177.7	0	0	0	150.7	151.0	151.3

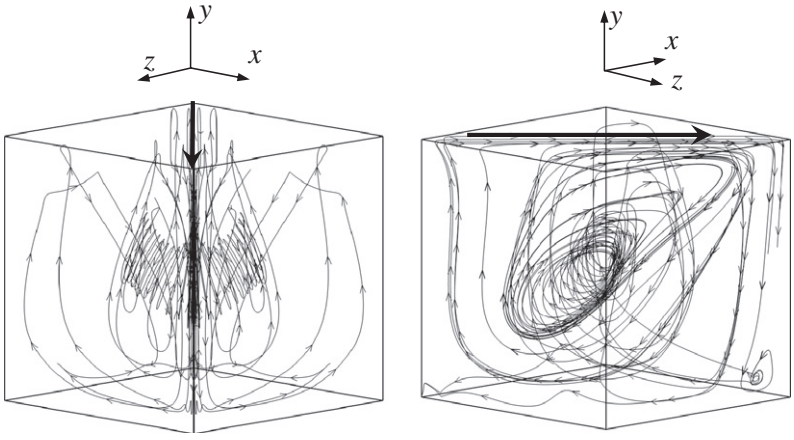

Fig. 4. A particle trajectories in a cubic lid-driven cavity with lid moving at 45° to the x axis, $Re = 10^3$.

Table 4
Pressure and velocity components along centerline $(0, y, 0)$ obtained on 152^3, 200^3 grids for $Re = 10^3$ and their Richardson extrapolation .

y	v_x, $v_z \times 10^3$			$v_y \times 10^3$			$p \times 10^4$		
	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.	Grid 152^3	Grid 200^3	Rich. extrap.
0.5	707.1	707.1	707.1	0	0	0	56.59	56.63	56.69
0.4766	417.7	417.8	417.9	5.378	5.357	5.328	51.59	51.59	51.57
0.4688	341.3	341.3	341.4	8.803	8.774	8.736	49.87	49.87	49.86
0.4609	277.1	277.2	277.4	12.50	12.47	12.42	48.24	48.24	48.23
0.4531	226.6	226.6	226.6	16.07	16.03	15.98	46.77	46.77	46.77
0.3516	76.74	76.82	76.92	30.36	30.13	29.83	35.18	35.22	35.28
0.2344	62.50	62.56	62.65	22.59	22.28	21.86	21.85	21.88	21.93
0.1172	41.78	41.77	41.76	5.790	5.439	4.975	9.711	9.717	9.725
0	−1.398	−1.649	−1.981	−33.95	−34.41	−35.03	0	0	0
−0.0469	−31.54	−31.93	−32.45	−64.70	−65.23	−65.95	−2.517	−2.518	−2.520
−0.2187	−130.7	−131.0	−131.4	−160.2	−160.5	−160.9	28.24	28.31	28.39
−0.3281	−134.7	−134.7	−134.7	−137.9	−138.0	−138.0	107.0	107.1	107.1
−0.3984	−143.1	−143.0	−143.0	−86.78	−86.68	−86.56	186.5	186.5	186.6
−0.4297	−158.8	−158.8	−158.9	−52.87	−52.73	−52.55	220.0	220.0	220.0
−0.4375	−161.8	−161.9	−161.9	−44.06	−43.93	−43.75	226.9	226.9	226.8
−0.4453	−162.2	162.2	−162.3	−35.52	−35.40	−35.23	232.9	232.9	232.8
−0.5	0	0	0	0	0	0	253.6	253.6	253.6

Please cite this article in press as: Feldman Yu, Gelfgat AY. From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation. Comput Fluids (2010), doi:10.1016/j.compfluid.2010.08.009

5. Conclusions

A parallelized version of multigrid algorithm based on PAR-ASA-CLGS smoother for time integration of incompressible Navier–Stokes equations on staggered grids has been developed and successfully verified using the lid driven cubic cavity flow benchmark problem. The scalability properties of the algorithm were studied for up to 2048 cores running in parallel. The speed up of the algorithm for its single- and multi-grid modes are found to be $O(Np/\ln(Np))$ and $O(Np/\log_{10}(Np))$, respectively. An explicit data communication between the diagonal processors is crucial for a proper implementation of a Gauss–Seidel process and this has a strong effect on the convergence rate of the parallelized algorithm. Use of the smoother based on the Gauss–Seidel iteration together with the two-dimensional virtual topology restricts the minimal redistributed sub-domain to the size of 3×3 in x and z directions respectively. The maximum number of cores involved in the analysis is restricted mainly by the coarsest grid size. Therefore, the multigrid process described is efficient only for those problems that exhibit a slow convergence to a stable asymptotic regime, e.g., steady or oscillatory state. When the convergence is fast or an initial guess can be chosen close to the solution, the single-grid approach becomes preferable. An example of the latter is time integration when the current result is used as the initial guess for the next time step.

Using the developed approach we extended the standard lid-driven flow in a cube benchmark to larger Reynolds number, as well as to a new fully three-dimensional configuration where the lid moves at the angle of 45° with respect to the cube vertical wall. These extensions can be important for further verification of 3D CFD codes.

Acknowledgments

This research was supported by Ministry of Science and Technology, Israel, Grant No. 3-5689. We are grateful to the Leibniz-Institute for Crystal Growth for giving us the chance to run the calculations on the HLRN-II supercomputer (see www.hlrn.de) and to Dr. Buttari for his help in parallelization of the code.

References

[1] Feldman Yu, Gelfgat AYu. On pressure–velocity coupled time-integration of incompressible Navier–Stokes equations using direct inversion of Stokes operator or accelerated multigrid technique. Comput Struct 2009;87:710–20.
[2] Feldman Yu, Gelfgat AYu. Oscillatory instability of a 3D lid-driven flow in a cube. Phys Fluids, accepted for publication.
[3] Gropp W, Lusc E, Skjellum A. Using MPI. The MIT press: Cambridge; 1999.
[4] Douglas CC. Multigrid methods in science and engineering. IEE Comput Sci Eng 1996;3:55–68.
[5] Degani AT, Fox GC. Parallel multigrid computation of the unsteady incompressible Navier–Stokes equations. J Comput Phys 1996;128:223–36.
[6] Oosterlee CW. The convergence of parallel multiblock multigrid methods. Appl Numer Math 1995;19:115–28.
[7] Ham FE, Lein FS, Strong AB. Multiple semi-coarsened multigrid method with application to large eddy simulation. Int J Numer Methods Fluids 2006;50:579–96.
[8] Durst F, Schäfer M. A parallel block-structured multigrid method for the prediction of incompressible flows. Int J Numer Methods Fluids 1996;22:549–65.
[9] Wang P, Ferraro RD. Parallel multigrid finite volume computation of three-dimensional thermal convection. Comput Math Appl 1999;37:49–60.
[10] Henniger R, Obrist D, Kleiser L. High-order accurate solution of the incompressible Navier–Stokes equations on massively parallel computers. J Comput Phys 2010;229:3543–72.
[11] John V, Tobiska L. Numerical performance of smoothers in coupled multigrid methods for the parallel solution of the incompressible Navier–Stokes equations. Int J Numer Methods Fluids 2000;33:453–73.
[12] Vanka S. Block-implicit multigrid calculation of two-dimensional recirculating flows. Comput Methods Appl Mech Eng 1986;59(1):29–48.
[13] Wesseling P, Oosterlee CW. Geometric multigrid with applications to computational fluid dynamics. J Comput Appl Math 2001;128:311–34.
[14] Manservisi S. Numerical analysis of Vanka-type solvers for steady Stokes and Navier–Stokes flows. SIAM J Numer Anal 2006;44:2025–56.
[15] Patankar SV. Numerical heat transfer and fluid flow. McGraw-Hill: New York; 1980.
[16] Oosterlee CW, Gaspar FJ, Washio T, Wienands R. Multigrid line smoothers for higher order upwind discretizations of convection-dominated problems. J Comput Phys 1998;139:274–307.
[17] Brandt A. Multi-level adaptive solutions to boundary-value problems. Math Comput 1977;31(138):333–90.
[18] Albensoeder S, Kuhlmann HC. Accurate three-dimensional lid-driven cavity flow. J Comput Phys 2005;206:536–58.

Please cite this article in press as: Feldman Yu, Gelfgat AY. From multi- to single-grid CFD on massively parallel computers: Numerical experiments on lid-driven flow in a cube using pressure–velocity coupled formulation. Comput Fluids (2010), doi:10.1016/j.compfluid.2010.08.009