

OpenFOAM GPU¹

<amonakov@ispras.ru>

OpenFOAM

GPU

(GPU).

GPU.

• OpenFOAM, CUDA, GPGPU,
GPU, GPU

1.

OpenFOAM —

[1].

MPI.

, OpenFOAM

OpenFOAM

OpenFOAM

icoFoam, pisoFoam, interFoam

OpenFOAM (

1

2007-2013

»,

07.514.11.4119

02

2011

<http://www.unicluster.ru>

OpenFOAM

OpenFOAM

OpenFOAM.

•

(PCG).

OpenFOAM

(PBiCG).

(

OpenFOAM-extend [2]

(BiCGStab).

PCG,

•

(GAMG).

MPI.

$Ax = b$

$x_1, x_2, \dots,$

x_0

(

10)

OpenFOAM

2.

$Ax = b$
 $AMx' = b, x = Mx'$
 M, AM
 M
 [3].

3.

OpenFOAM GPU.
 GPU.
 1. GPU:
 1. (PCG)
 2. (PBiCG, PBiCGStab)
 3.
 2. MPI- GPU- GPU:
 1. GPU- GPU
 2. GPU.
 GPU. (1.2)
 GPU

4.

GPU
 [4]:

```

r = b - Ax
p = z = Mr
v' = (r, z)
do {
  t = Ap
  a = v' / (p, t)
  x = x + ap
  r = r - at
  if (small (r))
    break;
  z = Mr
  v = (r, z)
  w = v / v'
  v' = v
  p = r + wp
}
( A — , M — , b —
, x — , r, z, p, t —
, v, v', a, w —
, (,) -
).
```

GPU
 1. BLAS-1:
 2.
 3.
 GPU
 GPU- GPU
 CUBLAS.
 GPU
 GPU
 PCI-Express
 GPU,

```

(
    ,
    r = r - at).

```

5.

5.1.

GPU

1. (ILU).

2. GPU.

(FSAI, AINV).

ILU,

GPU

AINV [5].

drop tolerance —

drop tolerance (droptol)

droptol = 0 AINV

(

),

,

```

droptol = 1.0

```

drop tolerance

GPU,

GPU,

CPU (

5.2.

(

drop tolerance).

1. (float)

2. AINV.

AINV

(

).

:

8.1.

1 Cavity3D.
icoFoam/cavity
16 ;

		End time	CPU, c	GPU, c	
20x20x20	0.001000	2.00000	22	42	0.52
40x40x40	0.000500	0.12500	71	56	1.27
80x80x80	0.000250	0.00775	284	91	3.12
100x100x100	0.000125	0.00050	105	48	2.18

. 1. OpenFOAM.
CPU PCG
DILU. GPU
drop tolerance, 0.1;
« CPU/GPU» icoFoam.
GPU
8000 , 64000

9.

1. :
2. -
3. AINV.
4. ,

10.

GPU
BiCGStab. ,
drop tolerance
MPI
IDR(s),
GPU.

[1] SGI, The OpenFOAM Foundation, <http://openfoam.org/>
[2] The OpenFOAM Extend Project, <http://www.extend-project.de/>
[3] M. Benzi, Preconditioning techniques for large linear systems: a survey. J. Comput. Phys., 128 (2002), 418–477
[4] Y. Saad, Iterative methods for sparse linear systems, SIAM, Philadelphia, 2003, 567
[5] R. Bridson, W.-P. Tang, Refining an approximate inverse, Journal on Computational and Applied Math, 123 (2000), Numerical Analysis 2000 vol. III: Linear Algebra, pp. 293-306.
[6] S. Pissanetzky, Sparse Matrix Technology, Academic Press, Waltham, 1984, 312
[7] G. Karypis, V. Kumar. METIS: Unstructured graph partitioning and sparse matrix ordering system, version 4.0, <http://www.cs.umn.edu/~metis>, 2009
[8] D. Gddke, R. Strzodka, S. Turek, Performance and accuracy of hardware-oriented native-, emulated- and mixed-precision solvers in FEM simulations, International Journal of Parallel, Emergent and Distributed Systems (IJPEDS), Special issue: Applied parallel computing, 22 (2007), 221–256

Optimizing OpenFOAM GPU Solvers

Alexander Monakov amonakov@ispras.ru, ISP RAS, Moscow, Russia

Abstract. The paper presents preliminary research on improving performance of CFD simulations in OpenFOAM via offloading parts of computations (specifically, solution of linear systems) to a graphics accelerator (GPU). We present a short review of OpenFOAM package and describe porting conjugate gradient method to the GPU architecture using CUDA programming model. Porting the basic algorithm is straightforward, however care should be taken to avoid unnecessary copying over PCI-Express bus. Efficient preconditioning on the GPU is then discussed. We use approximate inverse preconditioning, which can be implemented with good parallelism on the GPU. To amortize the cost of preparing the preconditioner, we allow reuse of preconditioners on the GPU and compute them on the CPU in a helper thread asynchronously. We mention several optimization opportunities: reordering the preconditioner to upper-left triangular form so that CUDA blocks multiplying by denser parts of preconditioner factors are scheduled first; using single-precision storage for the preconditioner to save memory bandwidth; reordering the mesh with nested dissection method from Metis library and using mixed-precision iteration for the conjugate gradient method. Preliminary performance testing results show performance improvement starting from 64000-cell meshes and reaching 2x for a 1-million cell mesh for a non-parallel run. As future work we mention support for parallel runs with MPI, research of other solvers such as multigrid, BiCGStab and IDR, and choosing drop tolerance automatically for the AINV preconditioner.

Keywords: OpenFOAM, CUDA, GPGPU, conjugate gradient technique, preconditioning in GPU, optimization for GPU

References

- [1]. SGI, The OpenFOAM Foundation, <http://openfoam.org/>
- [2]. The OpenFOAM Extend Project, <http://www.extend-project.de/>
- [3]. M. Benzi, Preconditioning techniques for large linear systems: a survey. J. Comput. Phys., 128 (2002), 418–477
- [4]. Y. Saad, Iterative methods for sparse linear systems, SIAM, Philadelphia, 2003, 567
- [5]. R. Bridson, W.-P. Tang, Refining an approximate inverse, Journal on Computational and Applied Math, 123 (2000), Numerical Analysis 2000 vol. III: Linear Algebra, pp. 293–306.
- [6]. S. Pissanetzky, Sparse Matrix Technology, Academic Press, Waltham, 1984, 312
- [7]. G. Karypis, V. Kumar. METIS: Unstructured graph partitioning and sparse matrix ordering system, version 4.0, <http://www.cs.umn.edu/~metis>, 2009
- [8]. D. Göddeke, R. Strzodka, S. Turek, Performance and accuracy of hardware-oriented native-, emulated- and mixed-precision solvers in FEM simulations, International Journal of Parallel, Emergent and Distributed Systems (IJPEDS), Special issue: Applied parallel computing, 22 (2007), 221–256