


Numerical simulation of astrophysical problems on Intel Xeon Phi based supercomputers

I. Kulikov, I. Chernykh
ICMMG SB RAS

June 9, 2016


Saint Petersburg

Cosmological simulation




The EAGLE project

Cosmological simulation



Physics:

- Cooling/Heating
- Stars formation
- Supernova explosion
- Chemical reactions



Cosmological simulation

Hydrodynamics

$$\frac{\partial \rho}{\partial t} + \frac{1}{a} \nabla \cdot (\rho \vec{u}) = 0$$

$$\frac{\partial \rho_i}{\partial t} + \frac{1}{a} \nabla \cdot (\rho_i \vec{u}) = s_i$$

$$\frac{\partial \rho \vec{u}}{\partial t} + \frac{1}{a} \nabla \cdot (\rho \vec{u} \vec{u}) = -\frac{1}{a} \nabla p - \frac{a'}{a} \rho \vec{u} - \frac{\rho}{a^2} \nabla \Phi$$

$$\frac{\partial \rho \varepsilon}{\partial t} + \frac{1}{a} \nabla \cdot (\rho \varepsilon \vec{u}) = -\frac{1}{a} (\gamma - 1) \rho \varepsilon \nabla \cdot (\vec{u}) - 2 \frac{a'}{a} \rho \varepsilon + \Gamma - \Lambda$$

$$\frac{\partial \rho E}{\partial t} + \frac{1}{a} \nabla \cdot (\rho E \vec{u}) = -\frac{1}{a} \nabla \cdot (p \vec{u}) - 2 \frac{a'}{a} \rho E - \frac{1}{a^2} (\rho \vec{u}, \nabla \Phi) + \Gamma - \Lambda$$

Hubble's law

$$\frac{da}{dt} = H \sqrt{\Omega_M (a^{-1} - 1) + \Omega_\Lambda (a^2 - 1) + 1}$$

Cosmological simulation

Collisionless component (stars and dark matter)

$$\frac{\partial n}{\partial t} + \frac{1}{a} \nabla \cdot (n \vec{v}) = 0$$

$$\frac{\partial n \vec{v}}{\partial t} + \frac{1}{a} \nabla \cdot (n \vec{v} \vec{v}) = -\frac{1}{a} \nabla \cdot (\Pi) - \frac{a'}{a} n \vec{v} - \frac{n}{a^2} \nabla \Phi$$

$$\frac{\partial \Pi_{xx}}{\partial t} + \frac{1}{a} \nabla \cdot (\Pi_{xx} \vec{v}) = -\frac{2}{a} \Pi_{xx} \frac{\partial v_x}{\partial x} - 2 \frac{a'}{a} \Pi_{xx}$$

$$\frac{\partial \Pi_{yy}}{\partial t} + \frac{1}{a} \nabla \cdot (\Pi_{yy} \vec{v}) = -\frac{2}{a} \Pi_{yy} \frac{\partial v_y}{\partial y} - 2 \frac{a'}{a} \Pi_{yy}$$


$$\frac{\partial \Pi_{zz}}{\partial t} + \frac{1}{a} \nabla \cdot (\Pi_{zz} \vec{v}) = -\frac{2}{a} \Pi_{zz} \frac{\partial v_z}{\partial z} - 2 \frac{a'}{a} \Pi_{zz}$$

$$\frac{\partial n W}{\partial t} + \frac{1}{a} \nabla \cdot (n W \vec{v}) = -\frac{1}{a} \nabla \cdot (\Pi \vec{v}) - 2 \frac{a'}{a} n W - \frac{1}{a^2} (n \vec{v}, \nabla \Phi)$$

$$\Delta \Phi = 4\pi G \left((\rho + n) - (\rho + n)_0 \right)$$

Numerical Method

$$\frac{\partial Q}{\partial t} + \nabla \cdot (Q \cdot \vec{v}) = \mathfrak{F}(Q, \nabla Q) \quad \xrightarrow{(1)} \quad \frac{\partial Q}{\partial t} = \mathfrak{F}(Q, \nabla Q)$$




$$\frac{\partial Q}{\partial t} + \nabla \cdot (Q \cdot \vec{v}) = 0$$

Riemann problem

$$\frac{\partial u}{\partial t} + B \frac{\partial u}{\partial x} = 0 \quad B = R\Lambda L \quad LR = I$$

$$L \frac{\partial u}{\partial t} + LR\Lambda L \frac{\partial u}{\partial x} = 0 \quad w = Lu$$

$$\frac{\partial w}{\partial t} + \Lambda \frac{\partial w}{\partial x} = 0 \quad w(x, t) = w(x - \Lambda t) \quad u = Rw$$



piecewise-parabolic functions

(*) Kulikov, et al., LNCS, 2009; APJS, 2011, 2014; AAABS, 2013; CPC 2015; NewA 2016
 (**) Ustyugov, Popov, Comp. Math. & Math. Phys., 2007, 2008, Comp. Phys., 2009

Overdefined equation system



$$\rho E = \rho \varepsilon + \frac{\rho v^2}{2} \quad \begin{matrix} \nearrow \\ \searrow \end{matrix} \quad \begin{aligned} \|V\| &= \sqrt{2(E - \varepsilon)} \\ \rho \varepsilon &= \left(\rho E - \frac{\rho v^2}{2} \right) \end{aligned}$$

1. Renormalization of velocity vector *at the boundary area between gas and vacuum**.
2. Correction of entropy**

This modification helps to control detailed energy balance, and entropy nondecrease guarantee.

*) Vshivkov V., Lazareva G., Snytnikov A., Kulikov I., Tutukov A. Computational methods for ill-posed problems of gravitational gasdynamics // J. Inv. Ill-Posed Problems, 19, 2011, 151-166




















***) Godunov S., Kulikov I. Computation of Discontinuous Solutions of Fluid Dynamics Equations with Entropy Nondecrease Guarantee // J. Comp. Math & Math. Phys., 54, 2014, 1012-1024

Numerical method




- ❑ High-order method
- ❑ The absence of artificial viscosity
- ❑ Galilean-invariant solution
- ❑ Entropy nondecrease guarantee
- ❑ Possible extension of the model by other hyperbolic equations (for example MHD equations)
- ❑ Simple parallelization
- ❑ «Potentially infinite scalability»

Chemical reactions




(Grassi et al., MNRAS, 2013)

Chemical reactions



Cosmological simulation



walls (pancake) $\lambda_1 \ll \lambda_2 \sim \lambda_3$

cluster $\lambda_1 \sim \lambda_2 \sim \lambda_3$

filaments $\lambda_1 \sim \lambda_2 \ll \lambda_3$


Parallel realization

134x speedup on 240 Intel Xeon Phi cores



RSC PetaStream (JSCC RAS)


MICs	64 x Intel Xeon Phi
Logical cores	15 360 (64 x 240)



RSC PetaStream (Polytechnic)

MICs	256 x Intel Xeon Phi
Logical cores	61 440 (256 x 240)

Parallel realization



Configuration	MPI/OpenMP Communications
Intel Core-i7 3820 (4 cores, 128 ³ grid size)	2.37 %
Intel Xeon E7-4870 (4 cores, 256 ³ grid size)	1.65 %
Intel Xeon E5-2650v2 (16 cores, 512 ³ grid size)	2.03 %
Intel Xeon Phi 7120D (60 MICs, 256 ³ grid size)	8.52 %
Intel Xeon Phi 7120D (60 MICs, 512 ³ grid size)	7.13 %


Number of accelerators	Efficiency (MVS-10P)	Efficiency (Polytechnic)
1	1,0000	1,0000
2	1,0345	1,0058
4	1,0335	0,9864
8	0,9771	0,9857
16	0,9417	0,9591
32	0,9345	0,9101
64	0,9235	0,8659
96	—	0,8326
128	—	0,8101
192	—	0,7711
224	—	0,7565

TABLE 2

Efficiency on a different accelerators number for MVS-10P and Polytechnic supercomputers.

73% efficiency (weak scalability) on 256x Intel Xeon Phi accelerators, 15360 cores (61440 threads) used

Vectorization



Initial code	Vectorized code
<p>Data operations:</p> <pre>void EulerStage (...) {... #pragma omp parallel for default(none) shared(...)private(...) num_threads(MIC_THREADS) ... R_diff = ...; for(curr_ind=0 ; curr_ind <NX;curr_ind++) { RVx[curr_ind] += - taumic*R_diff[curr_ind]/2/hmic; } ...}</pre>	<p>Data operations:</p> <pre>void EulerStage (...) {... #pragma omp parallel for default(none) shared(...)private(...) num_threads(MIC_THREADS) ... R_diff = ... ; __m512d taumic = __mm256_set1_pd(tau / 2.0 / hmic) for(curr_ind=0 ; curr_ind <NX;curr_ind+=8) { RVxv = __mm512_load_pd(RVx+curr_ind); Rdifv = __mm512_load_pd(R_diff+curr_ind); RVxv = __mm512_fmadd_pd(taumic,Rdifv,RVxv); __mm512_store_pd(RVx+curr_ind,RVxv); } ... }</pre>

Vectorization






Fig 2. Optimization of memory operations, commands for working with unaligned memory

Vectorization



```

rvppp = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vppp, zero, _CMP_LT_OS), rvppp, next);
rvpmp = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vpmp, zero, _CMP_LT_OS), rvpmp, next);
rvmpp = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vmpp, zero, _CMP_LT_OS), rvmpp, next);
rvmmp = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vmmp, zero, _CMP_LT_OS), rvmmp, next);
rvppm = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vppm, zero, _CMP_GT_OS), rvppm, prev);
rvpmm = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vpmm, zero, _CMP_GT_OS), rvpmm, prev);
rvmpm = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vmpm, zero, _CMP_GT_OS), rvmpm, prev);
rvmmm = _mm512_mask_blend_pd(_mm512_cmp_pd_mask(vmmm, zero, _CMP_GT_OS), rvmmm, prev);
FZP = _mm512_mul_pd(_mm512_add_pd(_mm512_add_pd(_mm512_mul_pd(vppp, rvppp), _mm512_mul_pd(vpmp, rvpmp)),
_mm512_add_pd(_mm512_mul_pd(vmpp, rvmpp), _mm512_mul_pd(vmmp, rvmmp))), four);
FZM = _mm512_mul_pd(_mm512_add_pd(_mm512_add_pd(_mm512_mul_pd(vppm, rvppm), _mm512_mul_pd(vpmm, rvpmm)),
_mm512_add_pd(_mm512_mul_pd(vmpm, rvmpm), _mm512_mul_pd(vmmm, rvmmm))), four);
__m512d aVect = _mm512_loadu_pd(a + i*NZ*NY + k*NZ + l);
__m512d dmvVect = _mm512_set1_pd(dmv);
__m512d result = _mm512_sub_pd(aVect,
_mm512_mul_pd(_mm512_add_pd(_mm512_add_pd(_mm512_sub_pd(FXP, FXM), _mm512_sub_pd(FYP, FYM)),
_mm512_sub_pd(FZP, FZM)), dmvVect));
  
```

Vectorization






Fig. 3. Acceleration of the AstroPhi code via vectorization. Stage 1 – algorithm optimization. Stage 2 – optimization of arithmetic operations. Stage 3 – optimization of memory operations. Stage 4 – reducing a number of memory load operations. Stage 5 – transition to FMA commands.

Future work



- Global MHD simulation of the solar wind interaction with comets, natural satellites and planets (the movement of galaxies through the intergalactic gas)
- A supernova explosion with chemical kinetics(178 "chemical" and nuclear reactions) in the gas-dynamic and relativistic hydrodynamics model
- Implementation of the astrophysical code to state-of-art level (Top10 astrophysical code in 2020)



Thank you for your attention