THE PARALLEL HYDRODYNAMIC CODE FOR ASTROPHYSICAL FLOW WITH STELLAR EQUATIONS OF STATE

Igor Kulikov, Senior Researcher, High Performance Simulation Lab, ICMMG SB RAS

Igor Chernykh, Head of the High Performance Simulation Lab, ICMMG SB RAS

Vitaly Vshivkov, Chief Researcher, High Performance Simulation Lab, ICMMG SB RAS

Vladimir Prigarin, PhD Student, Novosibirsk State Technical University

Vladimir Mironov, Senior Researcher, Lomonosov Moscow State University

Alexander Tutukov, Chief Researcher, Institute of Astronomy RAS

The Life of Stars

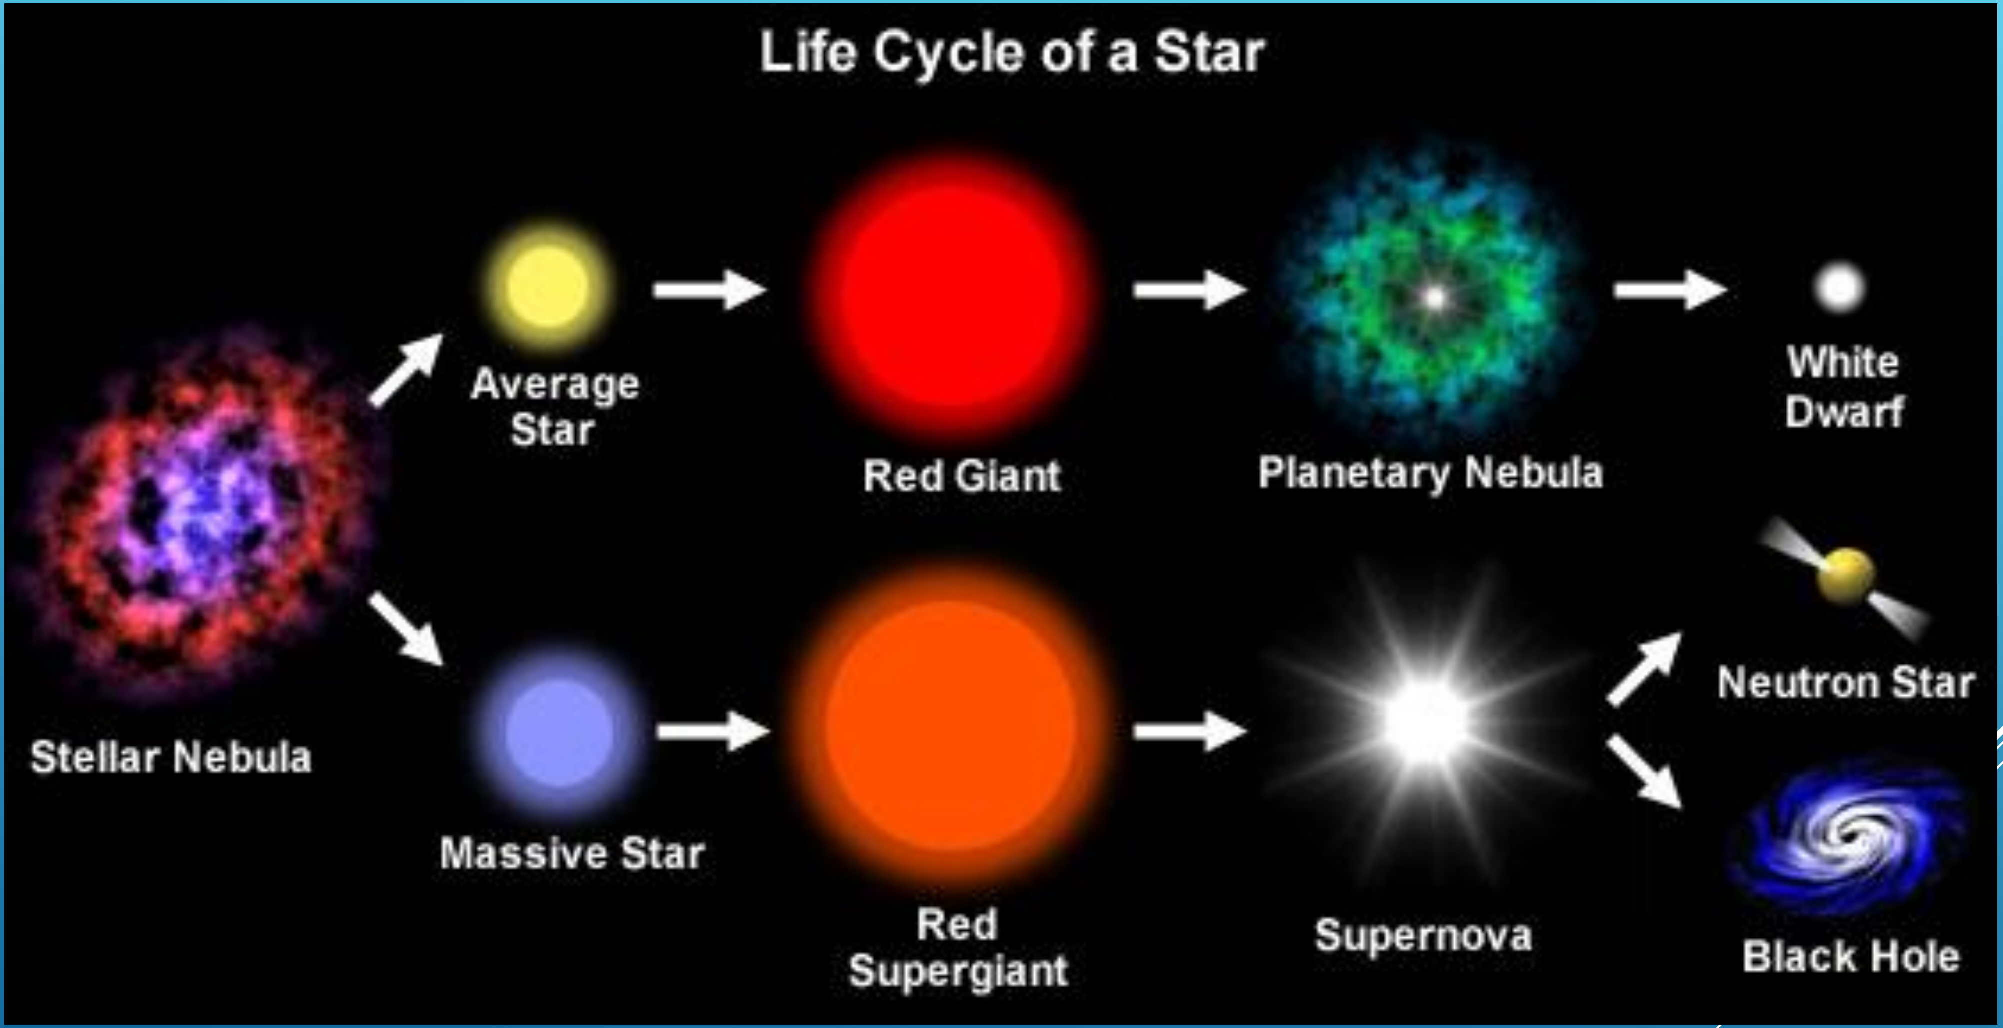

The Life of Stars

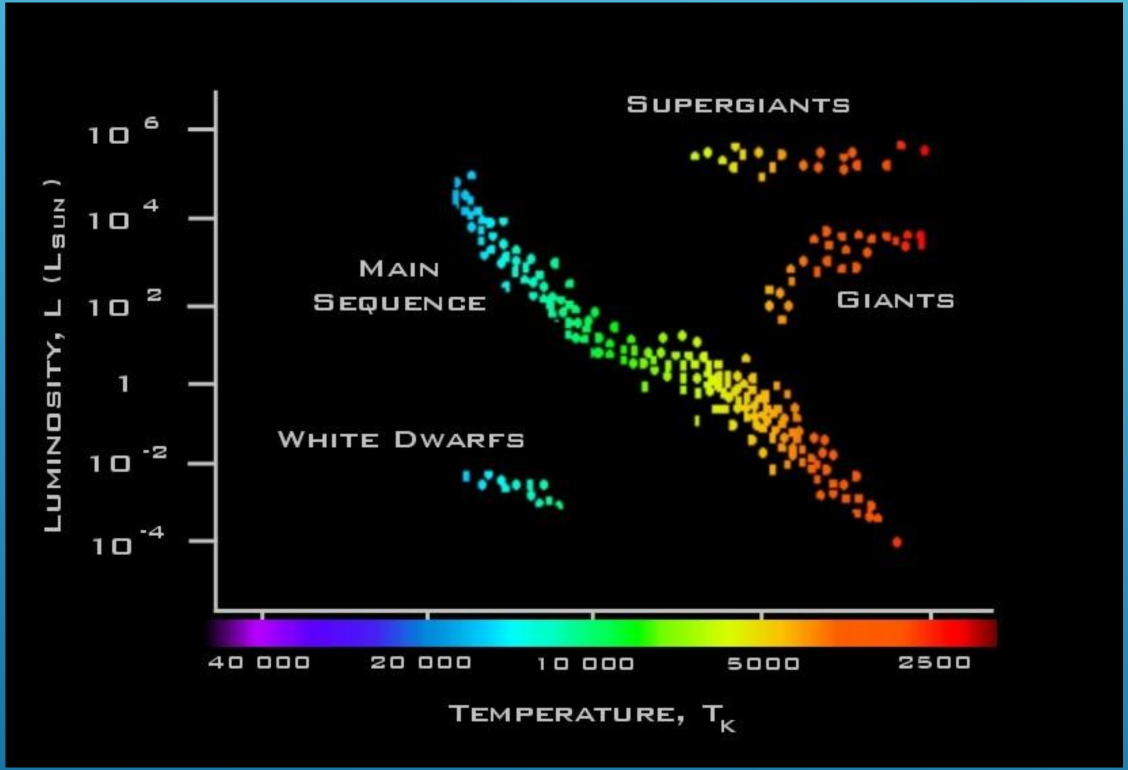

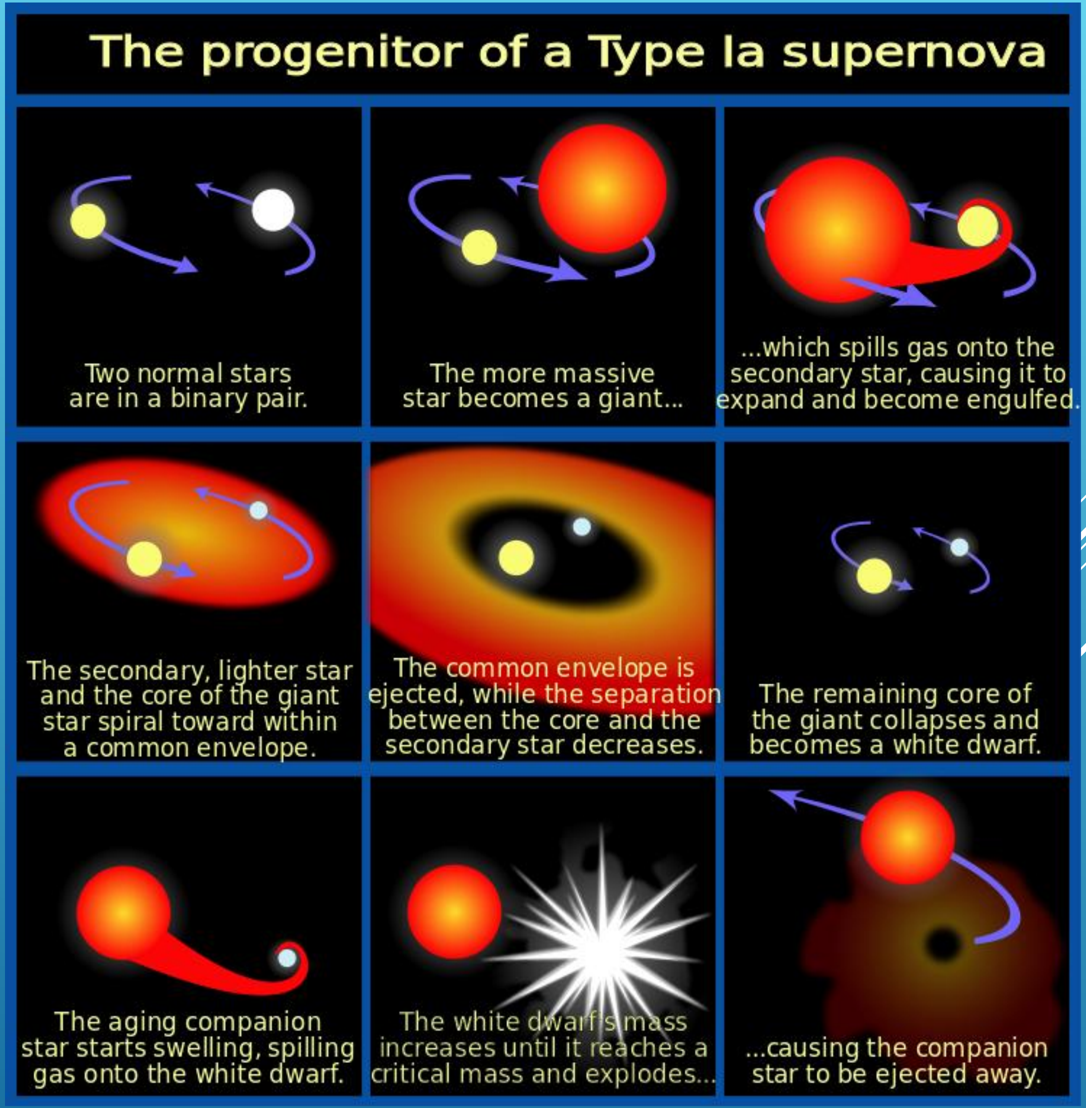

The progenitor of a Type Ia supernova

Two normal stars are in a binary pair.

The more massive star becomes a giant...

...which spills gas onto the secondary star, causing it to expand and become engulfed.

The secondary, lighter star and the core of the giant star spiral toward within a common envelope.

The common envelope is ejected, while the separation between the core and the secondary star decreases.

The remaining core of the giant collapses and becomes a white dwarf.

The aging companion star starts swelling, spilling gas onto the white dwarf.

The white dwarf's mass increases until it reaches a critical mass and explodes...

...causing the companion star to be ejected away.

The motivation

The supernova explosion enriches interstellar medium with the elements of life: *O, C, Fe, N, Si, Mg, Ca,…*

The mathematical & HPC challenges:

1. The numerical model construction
2. The numerical solver development
3. The efficiency parallel implementation

The Hydrodynamic Model of White Dwarf

- The Euler hydrodynamics equations
- The gravity
- The stellar equation of state:
  - Ideal gas for low temperature
  - Adiabatic (non)relativistic degenerate gas for high temperature
  - Radiation term
- The carbon burning 12C + 12C -> 23Na + p

Solvers in World

SPH approach

- 🟩 Robustness of the algorithm
- 🟩 Galilean-invariant solution
- 🟩 Simplicity of implementation
- 🟩 Flexible geometries of problems
- 🟩 High accurate gravity solvers

- 🟥 Artificial viscosity is parameterized
- 🟥 Variations of the smoothing length
- 🟥 The problem of shock wave and discontinuous solutions
- 🟥 Instabilities suppressed
- 🟥 The method is not scalable

AMR approach

- 🟩 Approved numerical methods
- 🟩 No artificial viscosity
- 🟩 Higher order shock waves
- 🟩 Resolution of discontinuities
- 🟩 No suppression of instabilities
- 🟩 Correct turbulence solution

- 🟥 The complexity of implementation
- 🟥 The effects of mesh
- 🟥 Problem of the minimal mesh
- 🟥 Not Galilean-invariant solution
- 🟥 The method is not scalable

The original numerical methods

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho\vec{v}\\ \rho E\\ \rho\varepsilon\end{pmatrix}+\nabla\cdot\begin{pmatrix}\rho\cdot\vec{v}\\ \rho\vec{v}\cdot\vec{v}\\ \rho E\cdot\vec{v}\\ \rho\varepsilon\cdot\vec{v}\end{pmatrix}=\begin{pmatrix}0\\ -\nabla p\\ -\nabla\cdot\left(p\vec{v}\right)\\ -p\nabla\cdot\left(\vec{v}\right)\end{pmatrix}$$

(1)

$$\frac{\partial}{\partial t}\begin{pmatrix}\rho\\ \rho\vec{v}\\ \rho E\\ \rho\varepsilon\end{pmatrix}=\begin{pmatrix}0\\ -\nabla p\\ -\nabla\cdot\left(p\vec{v}\right)\\ -p\nabla\cdot\left(\vec{v}\right)\end{pmatrix}$$

(2)

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

The Riemann problem

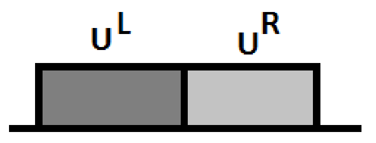

$$\frac{\partial u}{\partial t}+B\frac{\partial u}{\partial x}=0 \qquad B=R\Lambda L \qquad LR=I$$

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

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

The original numerical methods

The Riemann problem

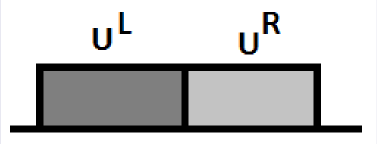

$$\frac{\partial u}{\partial t} + B\frac{\partial u}{\partial x} = 0 \qquad B = R\Lambda L \qquad LR = I$$

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

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

The piecewise-parabolic functions

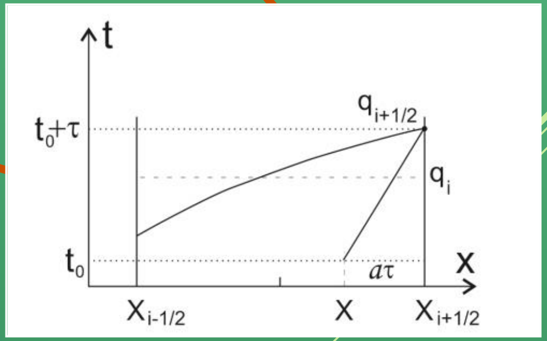

The parallel implementations

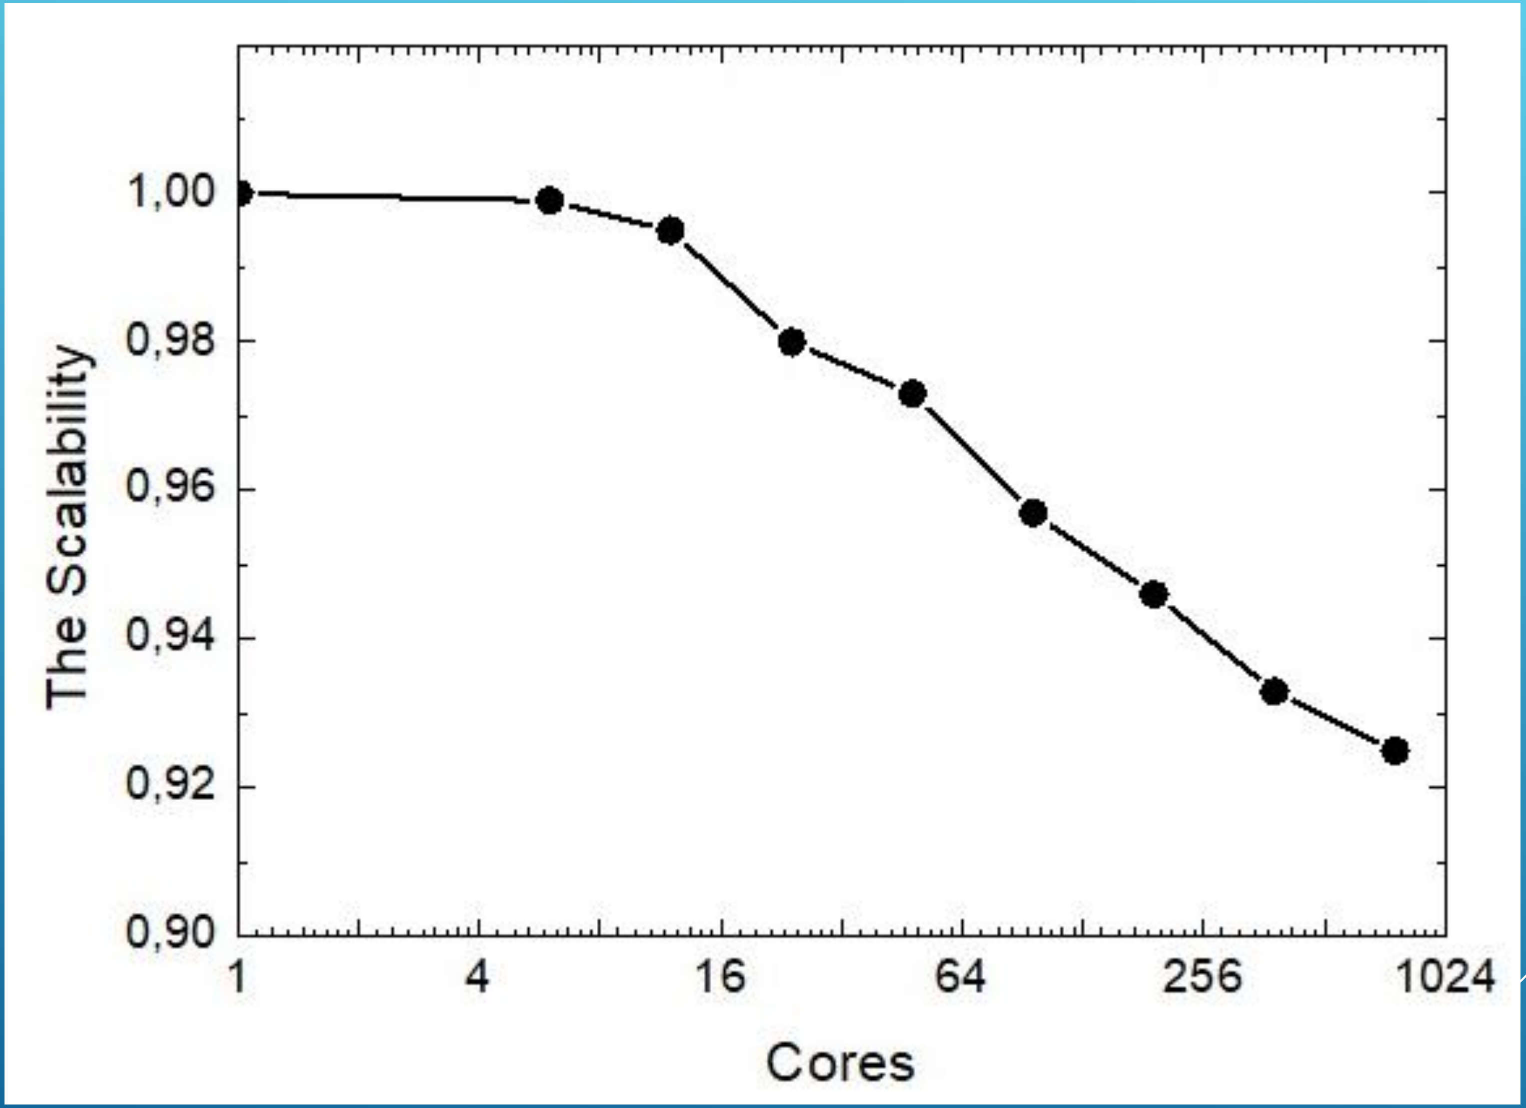

World’s Most Responsive Data Center SSD[1]

Delivering an **industry leading combination of low latency, high endurance, QoS and high throughput**, the Intel® Optane™ SSD is the first solution to **combine the attributes of memory and storage**. This innovative solution is optimized to **break through storage bottlenecks** by providing a new data tier. It accelerates applications for **fast caching and storage, increasing scale per server** and reducing transaction cost. Data centers based on the latest Intel® Xeon® processors can now also **deploy bigger and more affordable datasets** to gain new insights from larger memory pools.

[1] Responsiveness defined as average read latency measured at queue depth 1 during 4k random write workload. Measured using FIO 2.15. Common configuration - Intel 2U PCSD Server (“Wildcat Pass”), OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396GB DDR @ 2133MHz. Intel drives evaluated - Intel® Optane™ SSD DC P4800X 375GB, Intel® SSD DC P3700 1600GB, Intel® SSD DC P4600 1600GB. Samsung drives evaluated – Samsung® SSD PM1725a, Samsung® SSD PM1725, Samsung® PM963, Samsung® PM953. Micron drive evaluated – Micron® 9100 PCIe® NVMe™ SSD. Toshiba drives evaluated – Toshiba® ZD6300. Test – QD1 Random Read 4K latency, QD1 Random RW 4K 70% Read latency, QD1 Random Write 4K latency using fio-2.15.

INTEL® OPTANE™ SSD USE CASES

Fast Storage and Cache

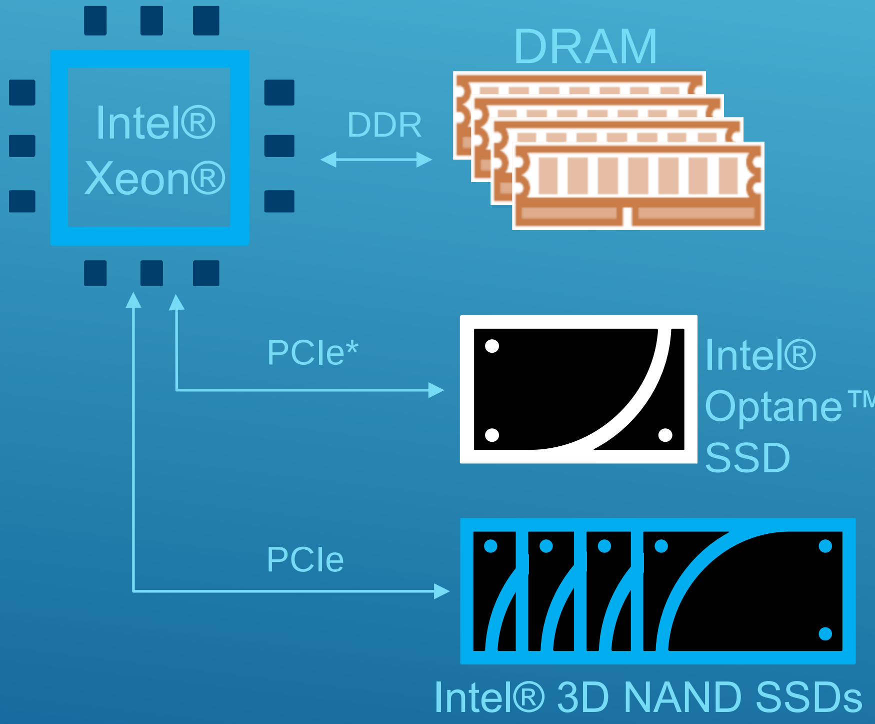

Extend Memory

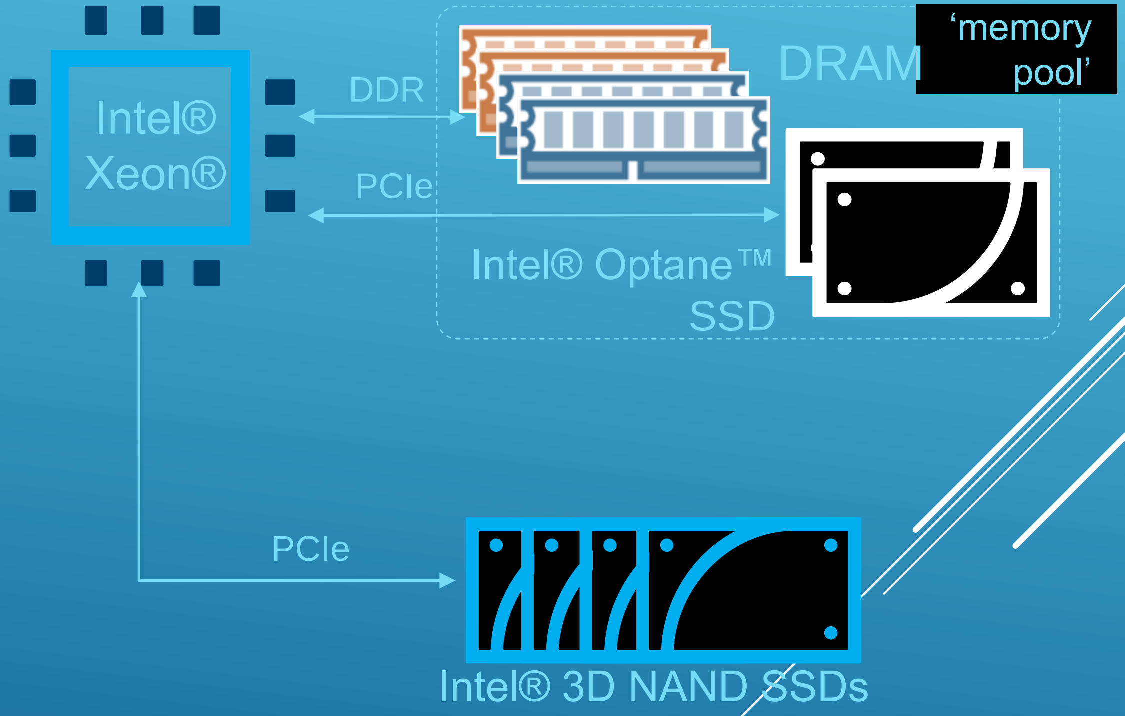

*Other names and brands names may be claimed as the property of others

HOW INTEL® MEMORY DRIVE TECHNOLOGY WORKS

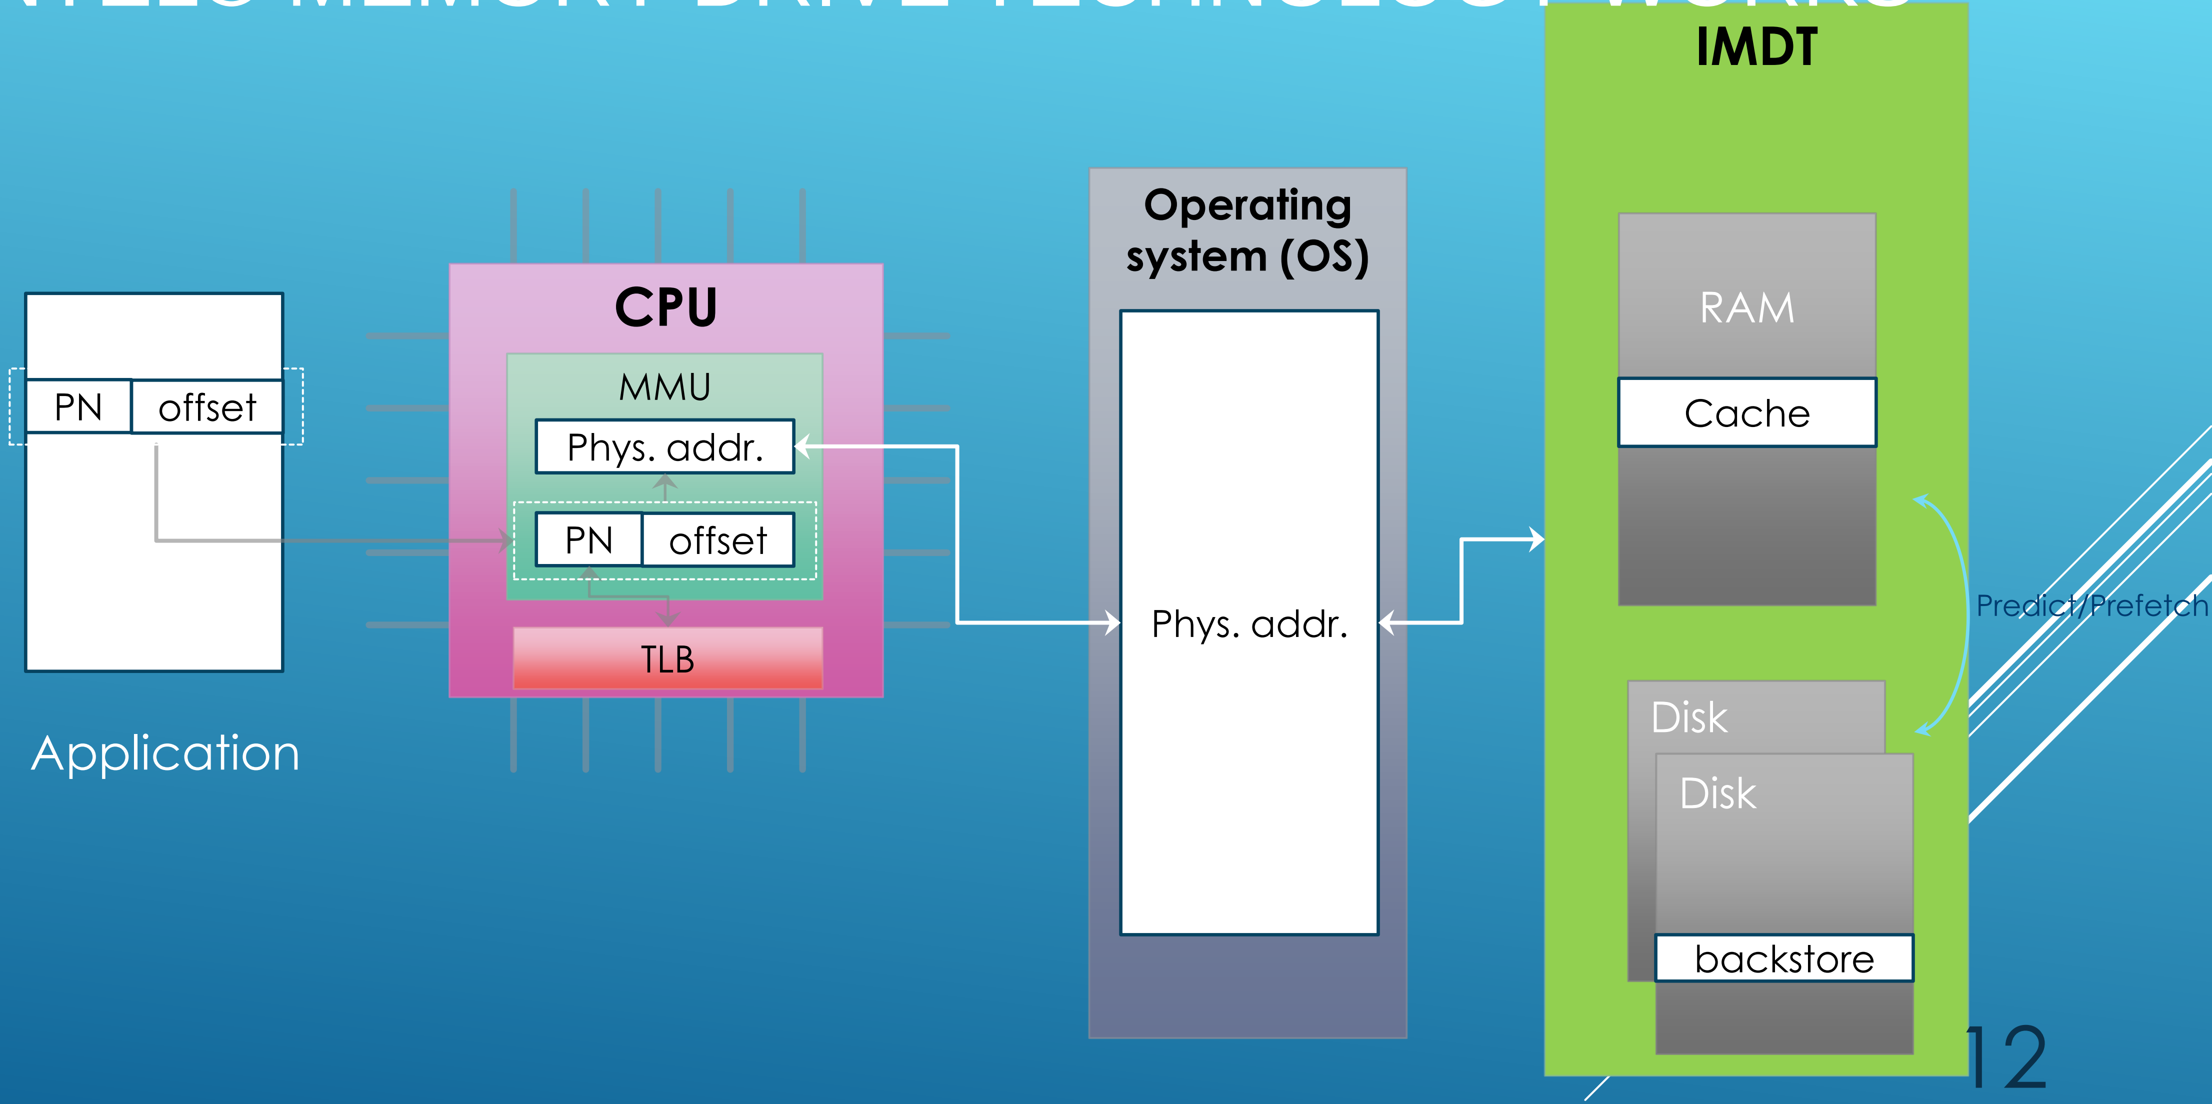

BENCHMARKS AND HARDWARE

- Used hardware:
  - Dual-socket Intel® Xeon® E5-2699 v4 (2x22 cores, 2.2 GHz)
  - First configuration (MDT):
    - 256 GB ECC DDR4
    - 4x320 GB Intel® Optane™ SSD (~10 GB/s aggregated bandwidth)
  - Second configuration (lot of DRAM):
    - 1536 GB ECC DDR4

- Used hardware:
  - Dual-socket Intel® Xeon® E5-2697A v4 (2x16 cores, 2.6 GHz)
  - First configuration (MDT):
    - 128 GB ECC DDR4
    - 2x350 GB Intel® Optane™ SSD (~10 GB/s aggregated bandwidth)

ASTROPHI

- The hyperbolic PDE engine
- Numerical 3D finite difference kernel
- Code is not currently optimized, opportunities for MDT optimization have been identified

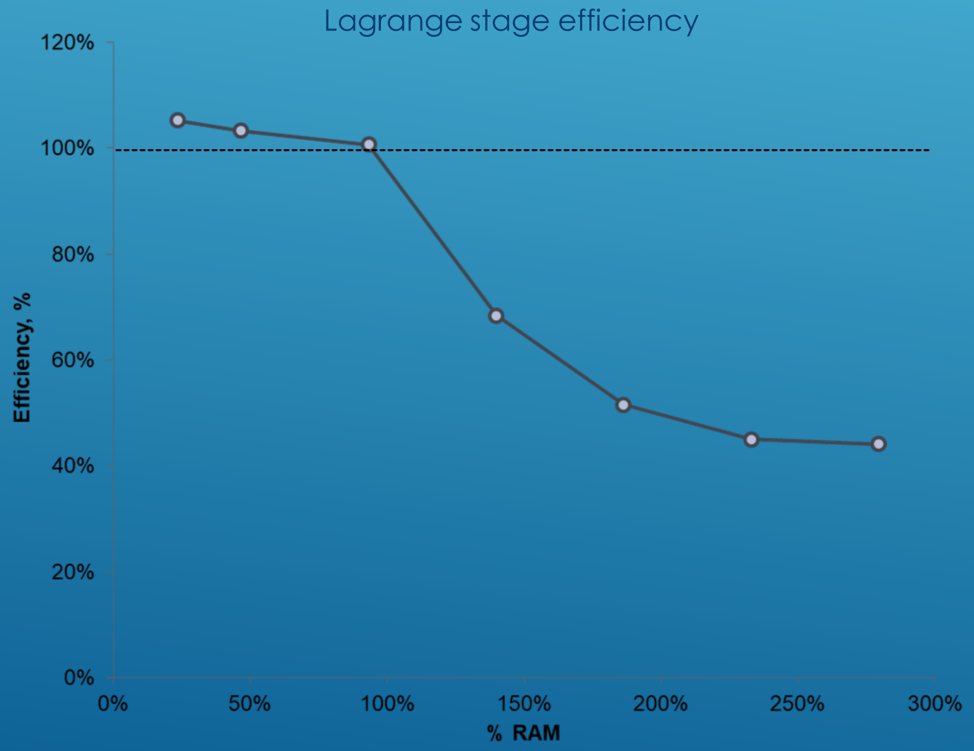

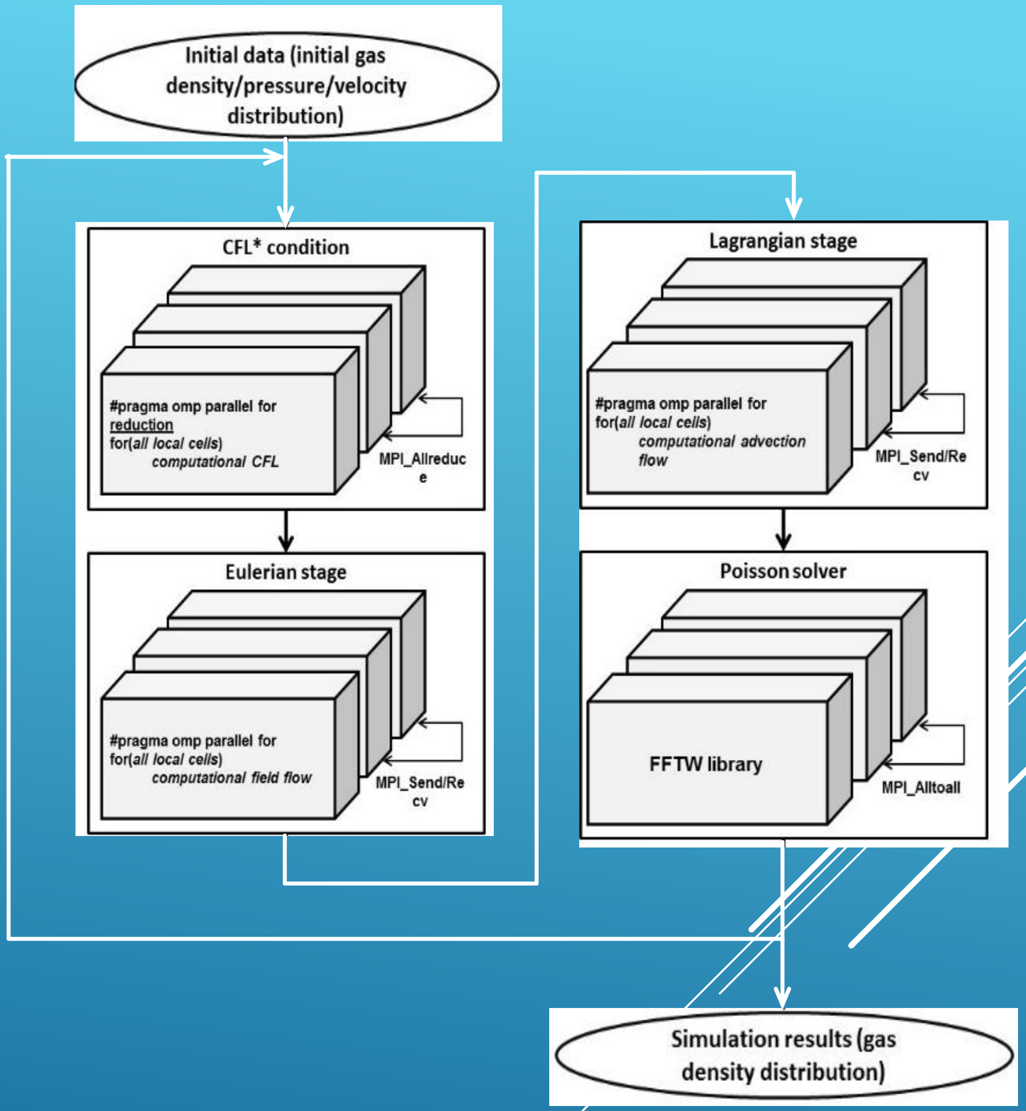

I.M.Kulikov, I.G.Chernykh, A.V.Snytnikov, B.M.Glinskiy, A.V.Tutukov, Comp. Phys. Comm., vol. 186, pp. 71-80, 2015.

The non-central explosion of white dwarf

NGC 6888

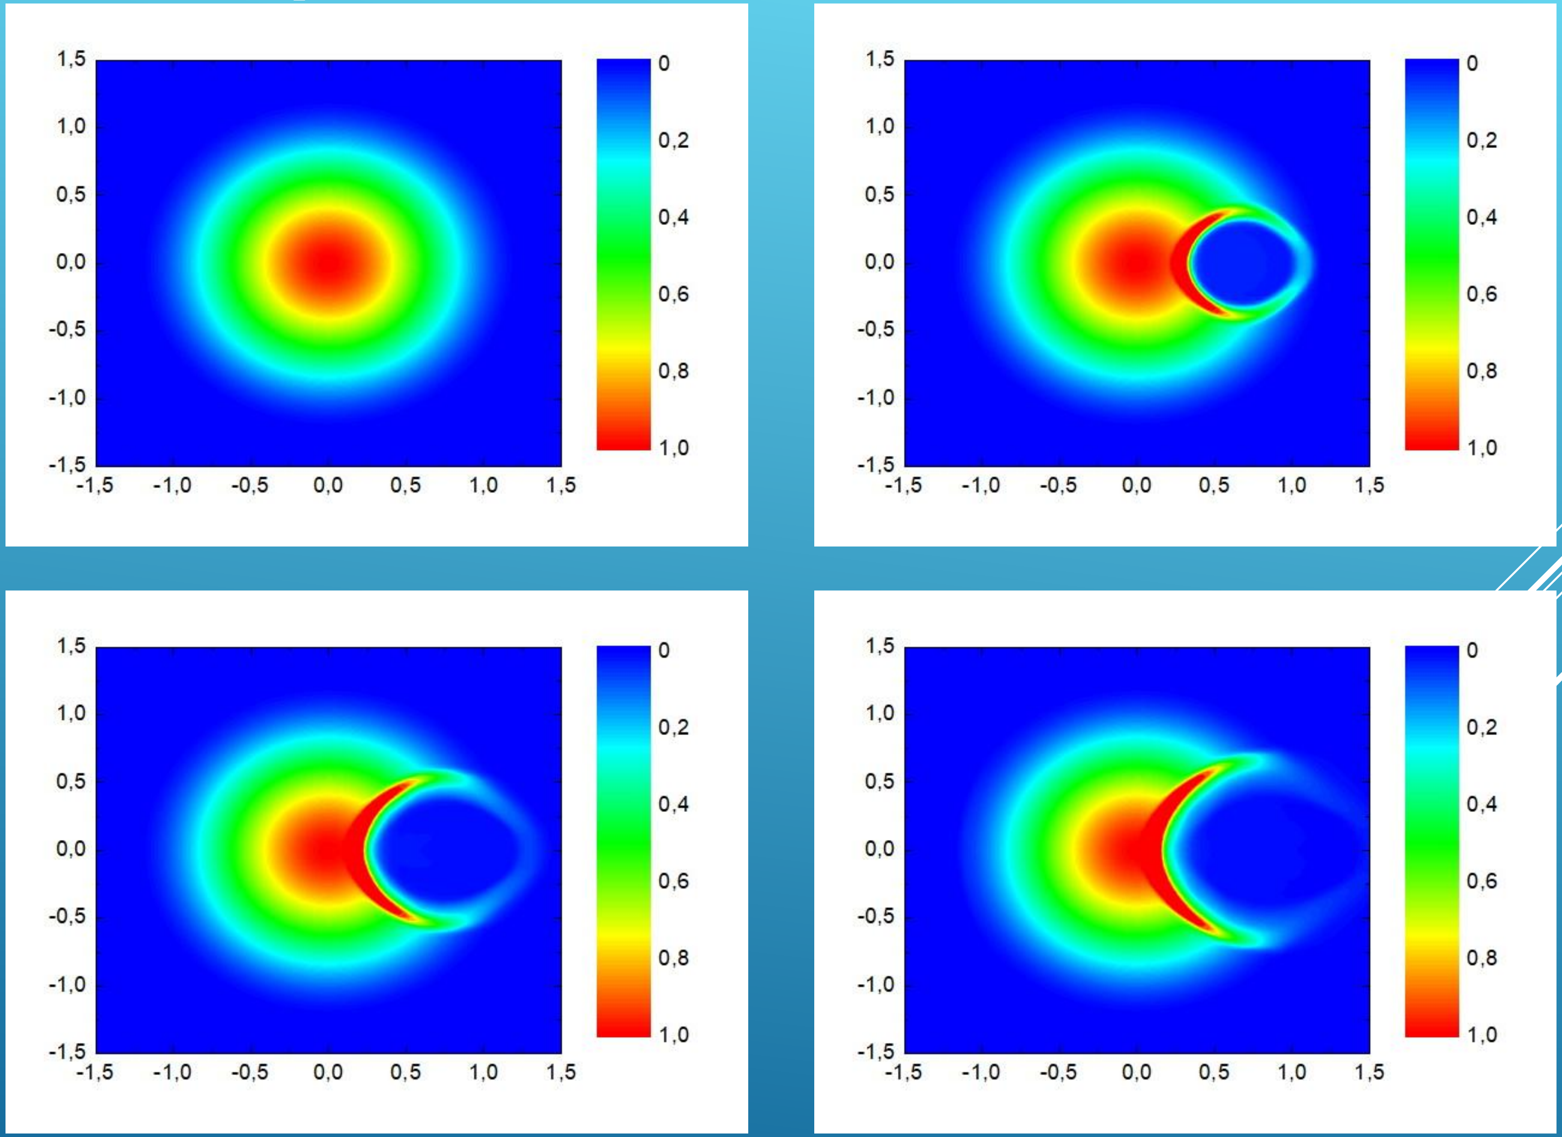

CONCLUSION

- A novel computation technique for numerical simulations of astrophysical flow at the supercomputers was described.
- We achieved more than 93 % weak scalability for 1024 CPU cores.
- For detailed numerical simulation of our problem, we need to use a large amount of RAM (more than 1TB) on each node.
- Optimal performance is expected on next generation of Intel® Optane™ SSDs

The publications

1. **Kulikov I. GPUPEGAS: A New GPU-accelerated Hydrodynamic Code for Numerical Simulations of Interacting Galaxies // The Astrophysical Journal Supplements Series. – 2014. – V. 214, 12. – P. 1-12**

2. **Kulikov I.M., Chernykh I.G., Snytnikov A.V., Glinskiy B.M., Tutukov A.V. AstroPhi: A code for complex simulation of dynamics of astrophysical objects using hybrid supercomputers // Computer Physics Communications. – 2015. – V. 186. – P. 71-80**

3. **Godunov S.K., Kulikov I.M. Computation of discontinuous solutions of fluid dynamics equations with entropy nondecrease guarantee // Computational Mathematics and Mathematical Physics. – 2014. – V. 54, № 6. – P. 1012-1024**

4. **Kulikov I., Vorobyov E. Using the PPML approach for constructing a low-dissipation, operator-splitting scheme for numerical simulations of hydrodynamic flows // Journal of Computational Physics. – 2016. – V. 317. – P. 318-346.**

ACKNOWLEDGEMENTS

- We thank:
  - The Russian Science Foundation (project 18-11-00044).
  - RSC Group for providing access to certain hardware
  - Intel Corporation for consulting and for template of presentation
  - All of You for Your attention!