Funded by the European Union

Co-ordinated by ECMWF

ESCAPE

ESCAPE2

These projects have received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671627 (ESCAPE) and No 800897 (ESCAPE2)

Funded by the European Union

Co-ordinated by ECMWF

ESCAPE 2

ESCAPE1+2: Energy-efficient Scalable Algorithms for Weather and Climate Prediction at Exascale

Andreas Mueller, Giovanni Tumolo, Willem Deconinck, Nils Wedi, Peter Bauer, et al.

ECMWF (European Centre for Medium-Range Weather Forecasts)

ESCAPE: ECMWF, MeteoSwiss, DMI, RMI, NVIDIA, METEO FRANCE Toujours un temps d'avance, ICHEC, Deutscher Wetterdienst Wetter und Klima aus einer Hand, Optalysys, Bull atos technologies, PSNC, Loughborough University

ESCAPE2: ECMWF, DKRZ Deutsches Klimarechenzentrum, Max-Planck-Institut für Meteorologie, MeteoSwiss, Barcelona Supercomputing Center Centro Nacional de Supercomputación, cea, Loughborough University, RMI, Politecnico Milano, MOX, DMI, cmcc Centro Euro-Mediterraneo sui Cambiamenti Climatici, Bull atos technologies

These projects have received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671627 (ESCAPE) and No 800897 (ESCAPE2)

Funded by the European Union

Dwarf concept: identify key patterns in terms of computation and communication

- Phil Colella introduced in a presentation in 2014 seven dwarfs of algorithms for high-end simulation in the physical sciences.
- later extended to 13 Berkeley dwarfs
- 1. Dense Linear Algebra
- 2. Sparse Linear Algebra
- 3. Spectral Methods
- 4. N-Body Methods
- 5. Structured Grids
- 6. Unstructured Grids
- 7. MapReduce
- 8. Combinational Logic
- 9. Graph Traversal
- 10. Dynamic Programming
- 11. Backtrack and Branch-and-Bound
- 12. Graphical Models
- 13. Finite State Machines

Concept behind ESCAPE

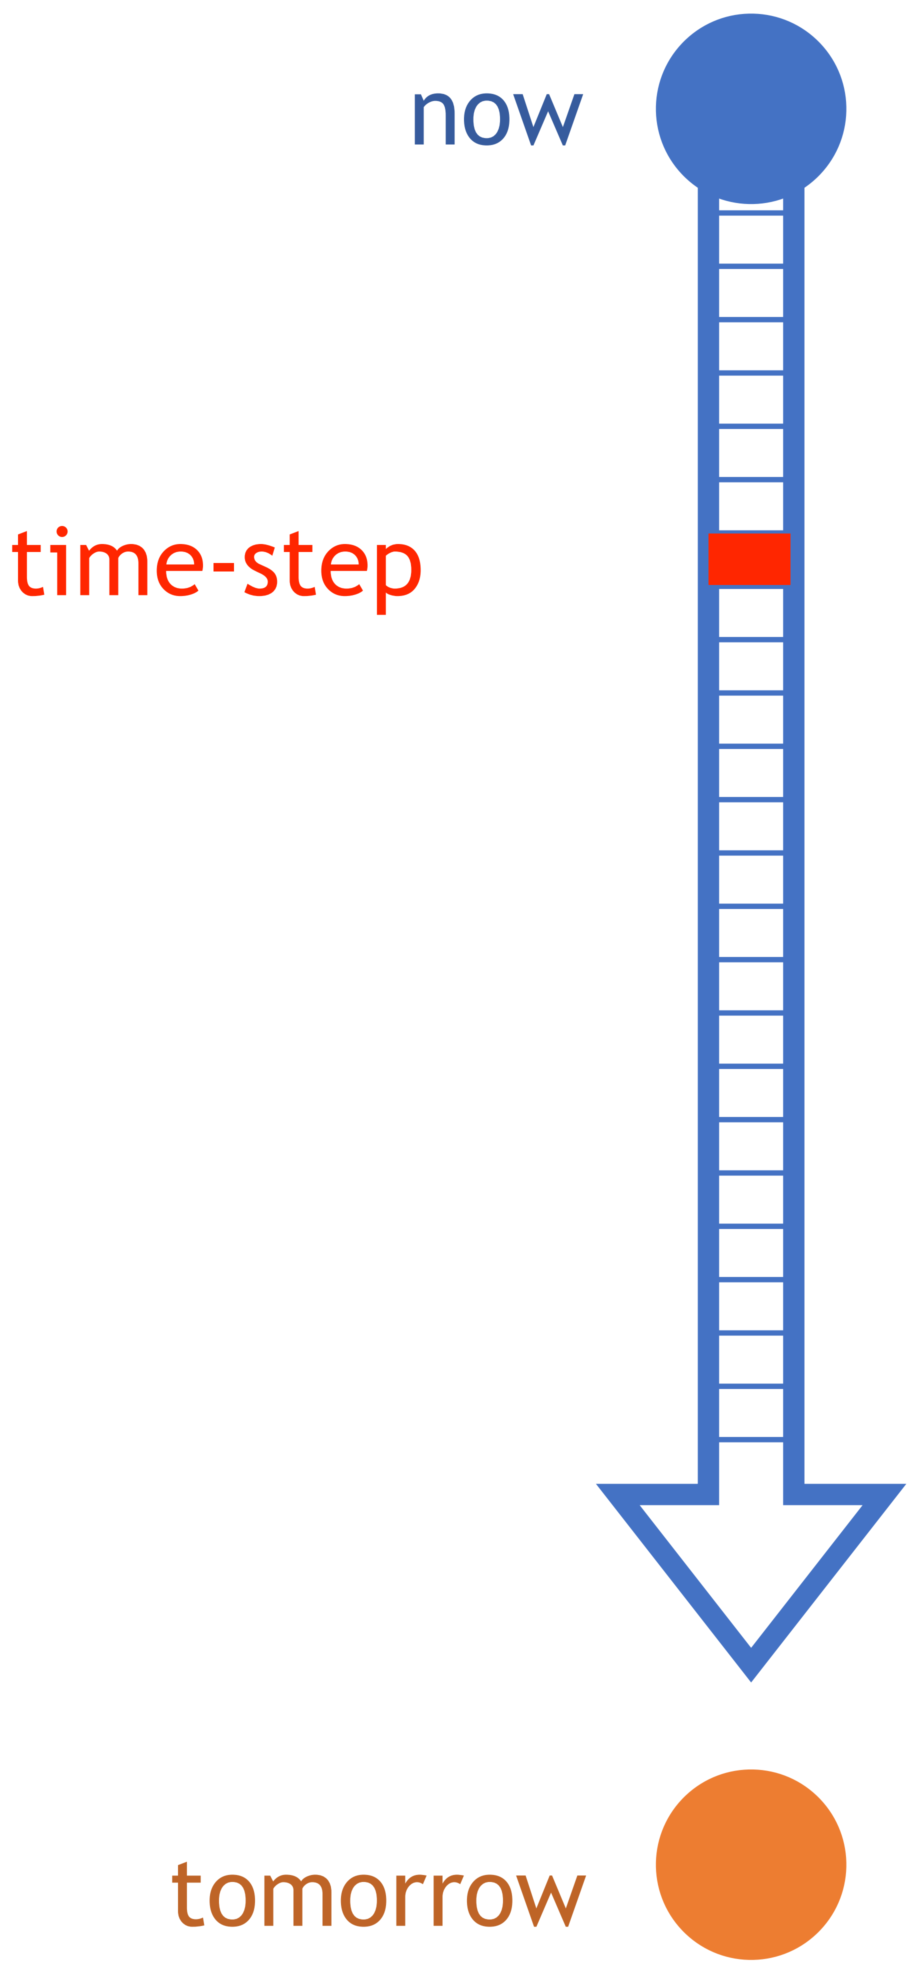

Funded by the European Union

Concept behind ESCAPE

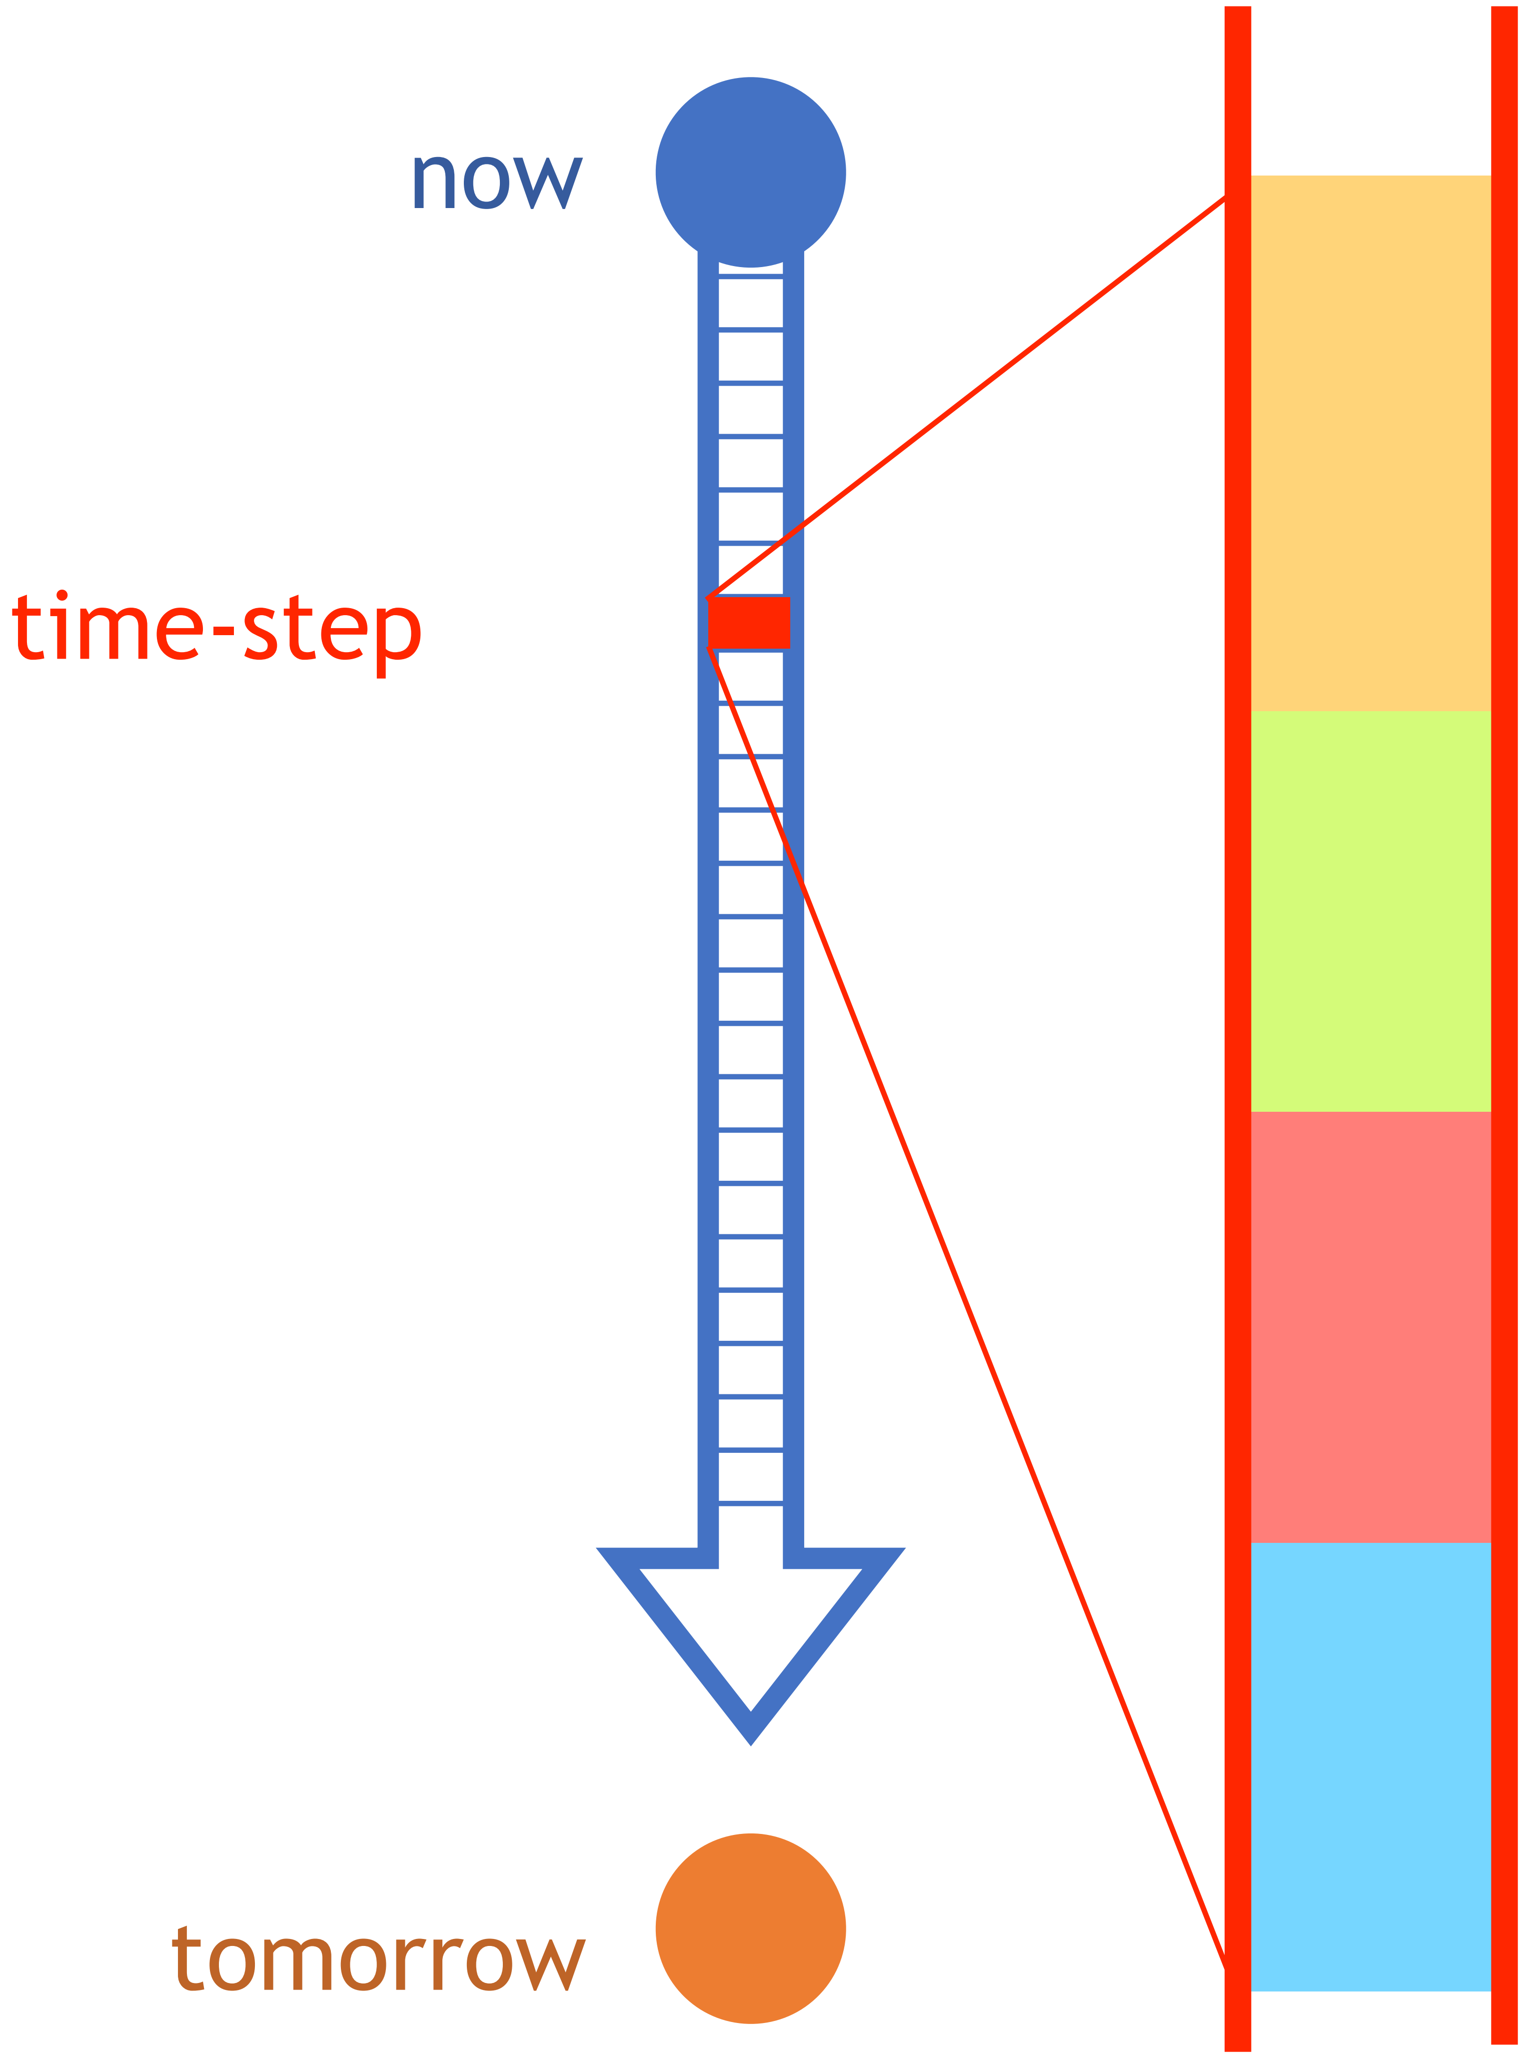

components:

time-stepping

advection

gradient computation

physics

Funded by the European Union

Concept behind ESCAPE

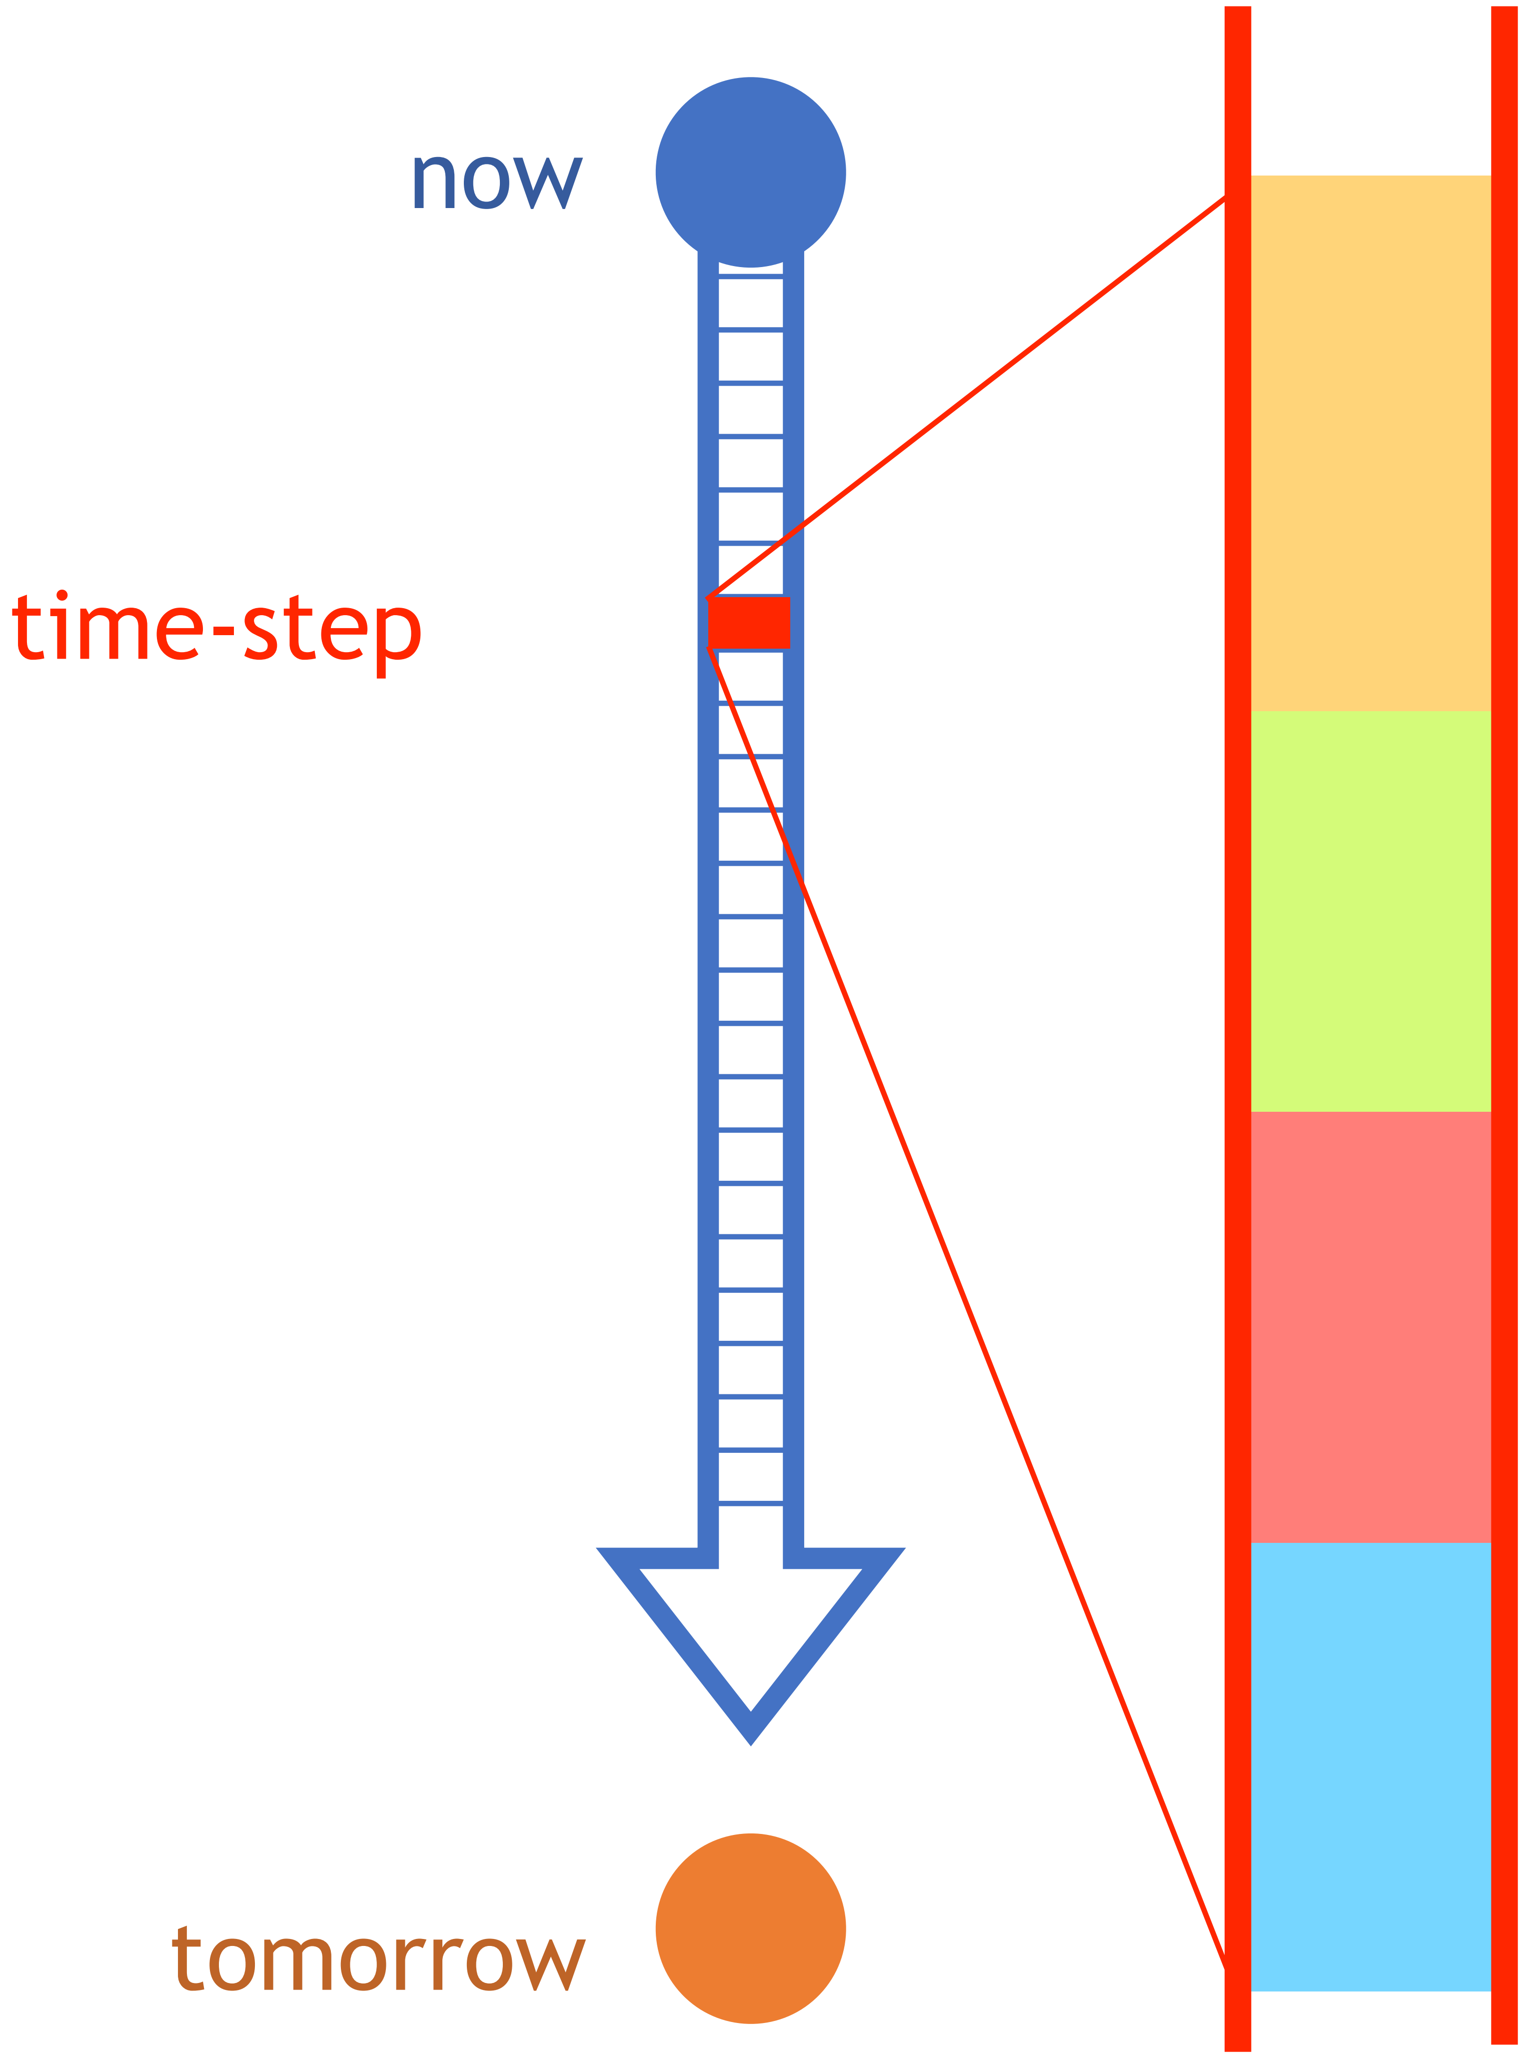

components:	**challenges:**
time-stepping	communication, memory
advection	halo-communication
gradient computation	expensive calculations
physics	expensive calculations

ESCAPE 2

Funded by the European Union

Concept behind ESCAPE

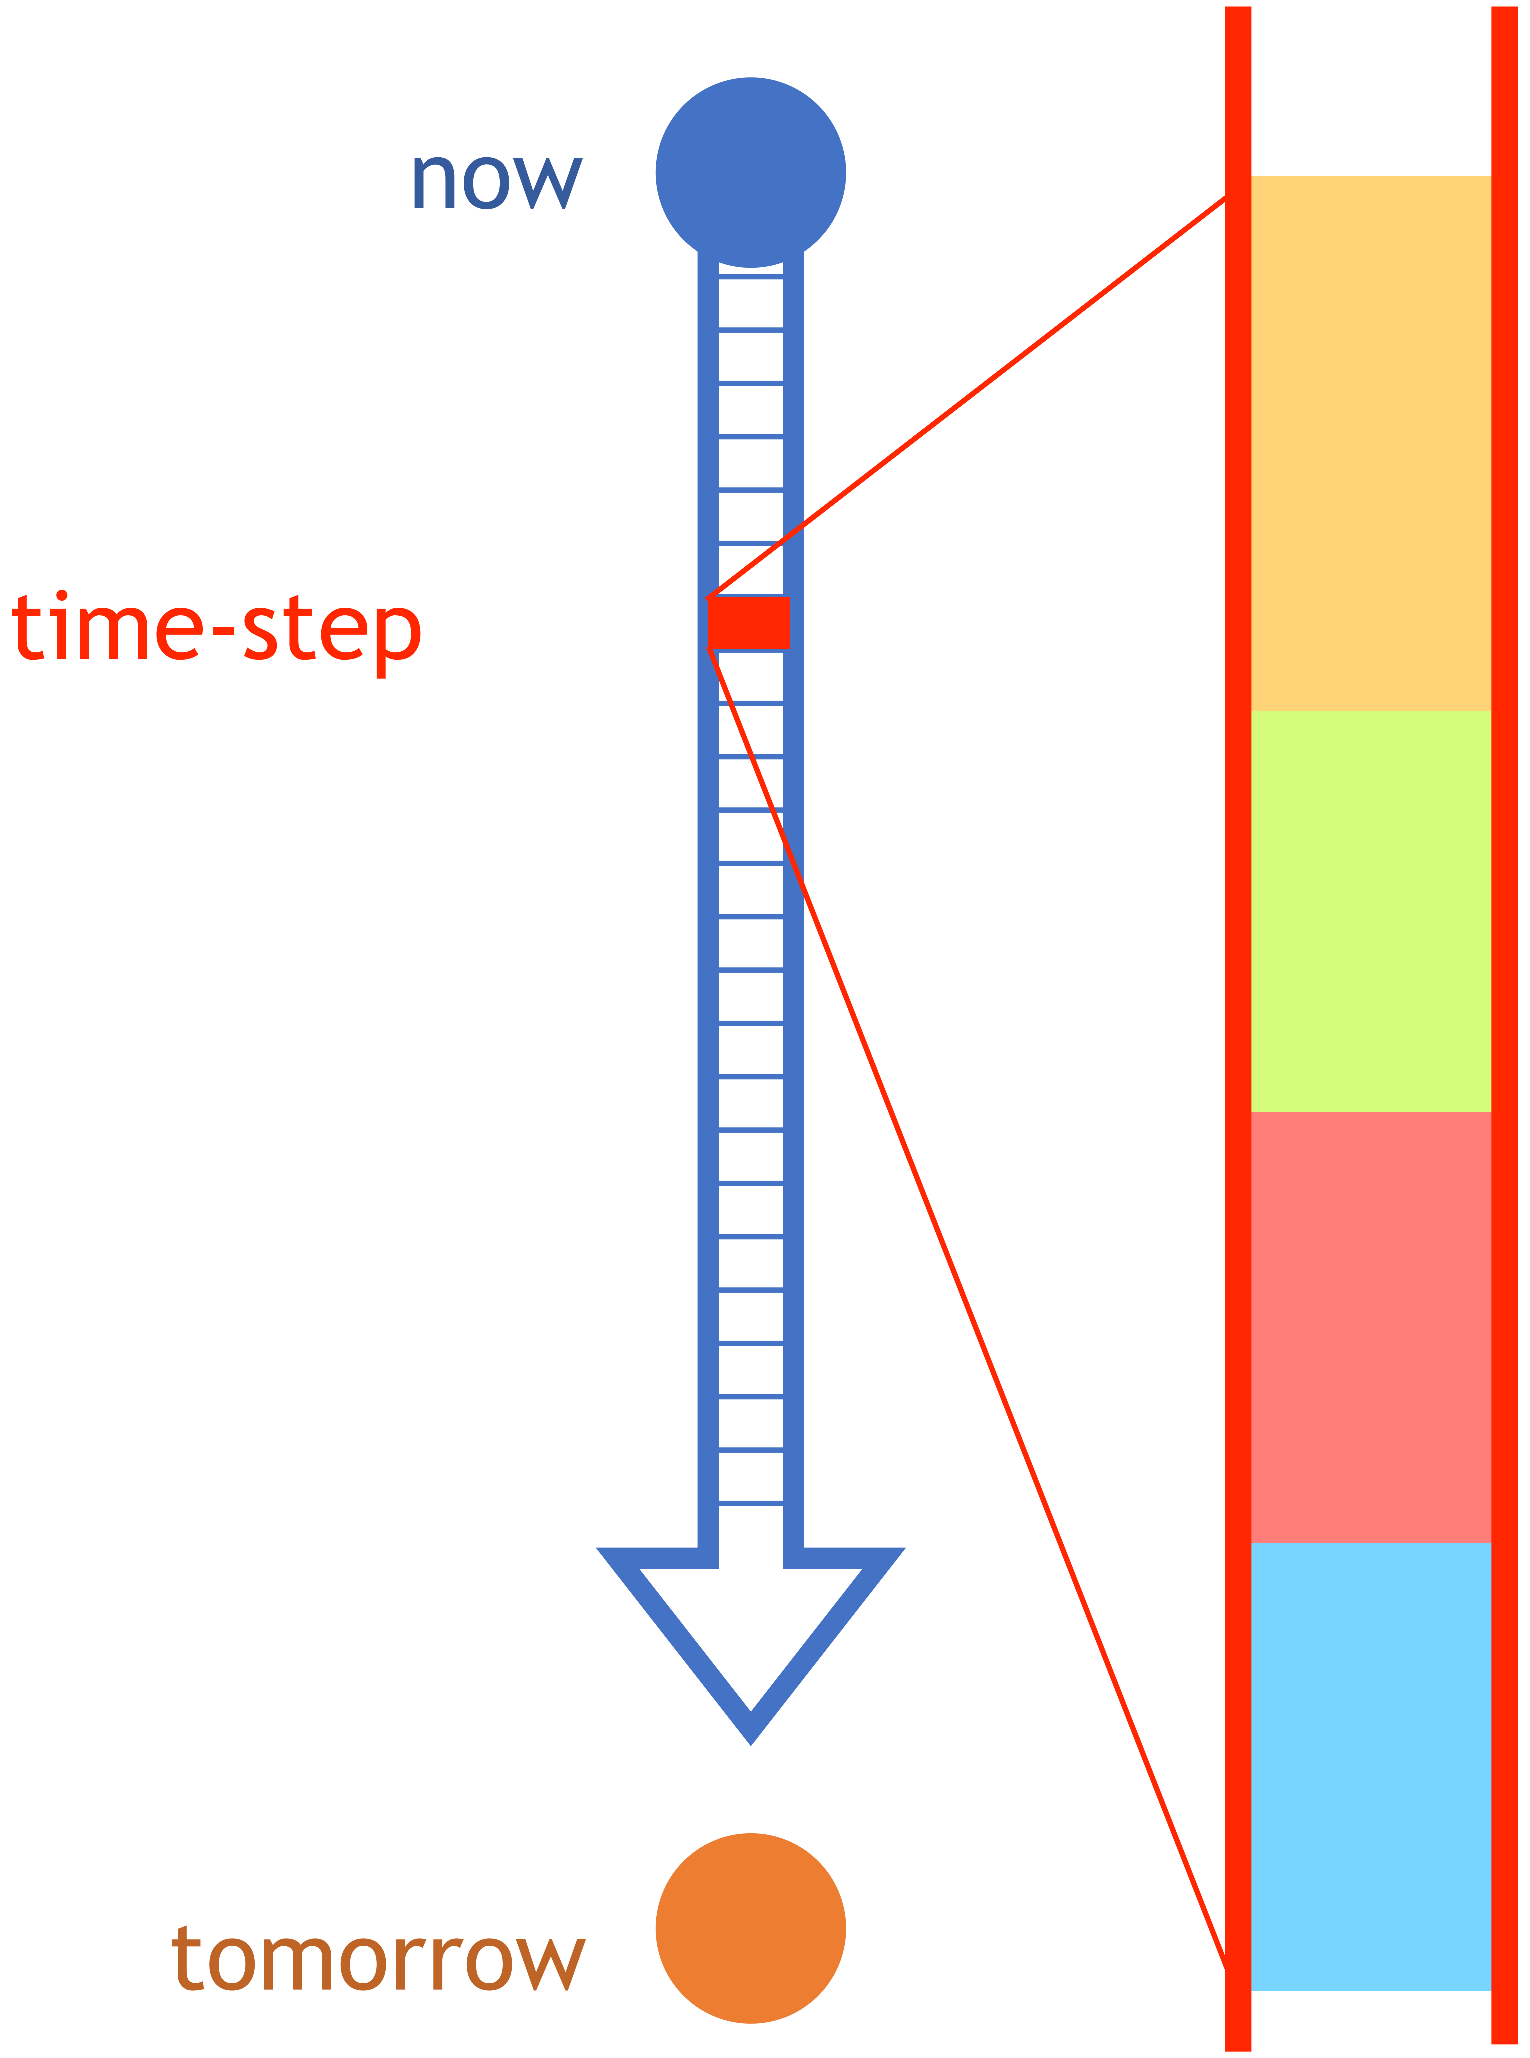

components:	options:	IFS	FVM	ALARO	EULAG	GRASS
time-stepping	spectral transform	✓				
	biFFT			✓		
	elliptic solver, multigrid		✓		✓	
	HEVI					✓
advection	semi-Lagrangian	✓		✓		
	LAITRI	✓		✓		
	MPDATA		✓		✓	
gradient computation	high order finite difference					✓
	finite volume		✓		✓	
physics	ACRANEB2 radiation			✓		
	CLOUDSC microphysics	✓	✓			

ESCAPE ESCAPE 2

Funded by the European Union

Concept behind ESCAPE

components:	options:	IFS	FVM	ALARO	EULAG	GRASS	prototype	Atlas	MPI	OpenMP	OpenACC	DSL	Optalysys
time-stepping	spectral transform	✓					✓	✓	✓	✓	✓		✓
	biFFT			✓			✓		✓	✓	✓		✓
	elliptic solver, multigrid		✓		✓		✓	✓	✓	✓			
	HEVI					✓			✓	✓			
advection	semi-Lagrangian	✓		✓			✓	✓	✓	✓			
	LAITRI	✓		✓			✓			✓	✓		
	MPDATA		✓		✓		✓	✓	✓	✓	✓	✓	
gradient computation	high order finite difference					✓			✓	✓			
	finite volume		✓		✓			✓	✓	✓			
physics	ACRANEB2 radiation			✓			✓		✓	✓	✓		
	CLOUDSC microphysics	✓	✓				✓		✓	✓	✓		

Relevance of our dwarfs

% of the forecast runtime

IFS 9km (ECMWF)

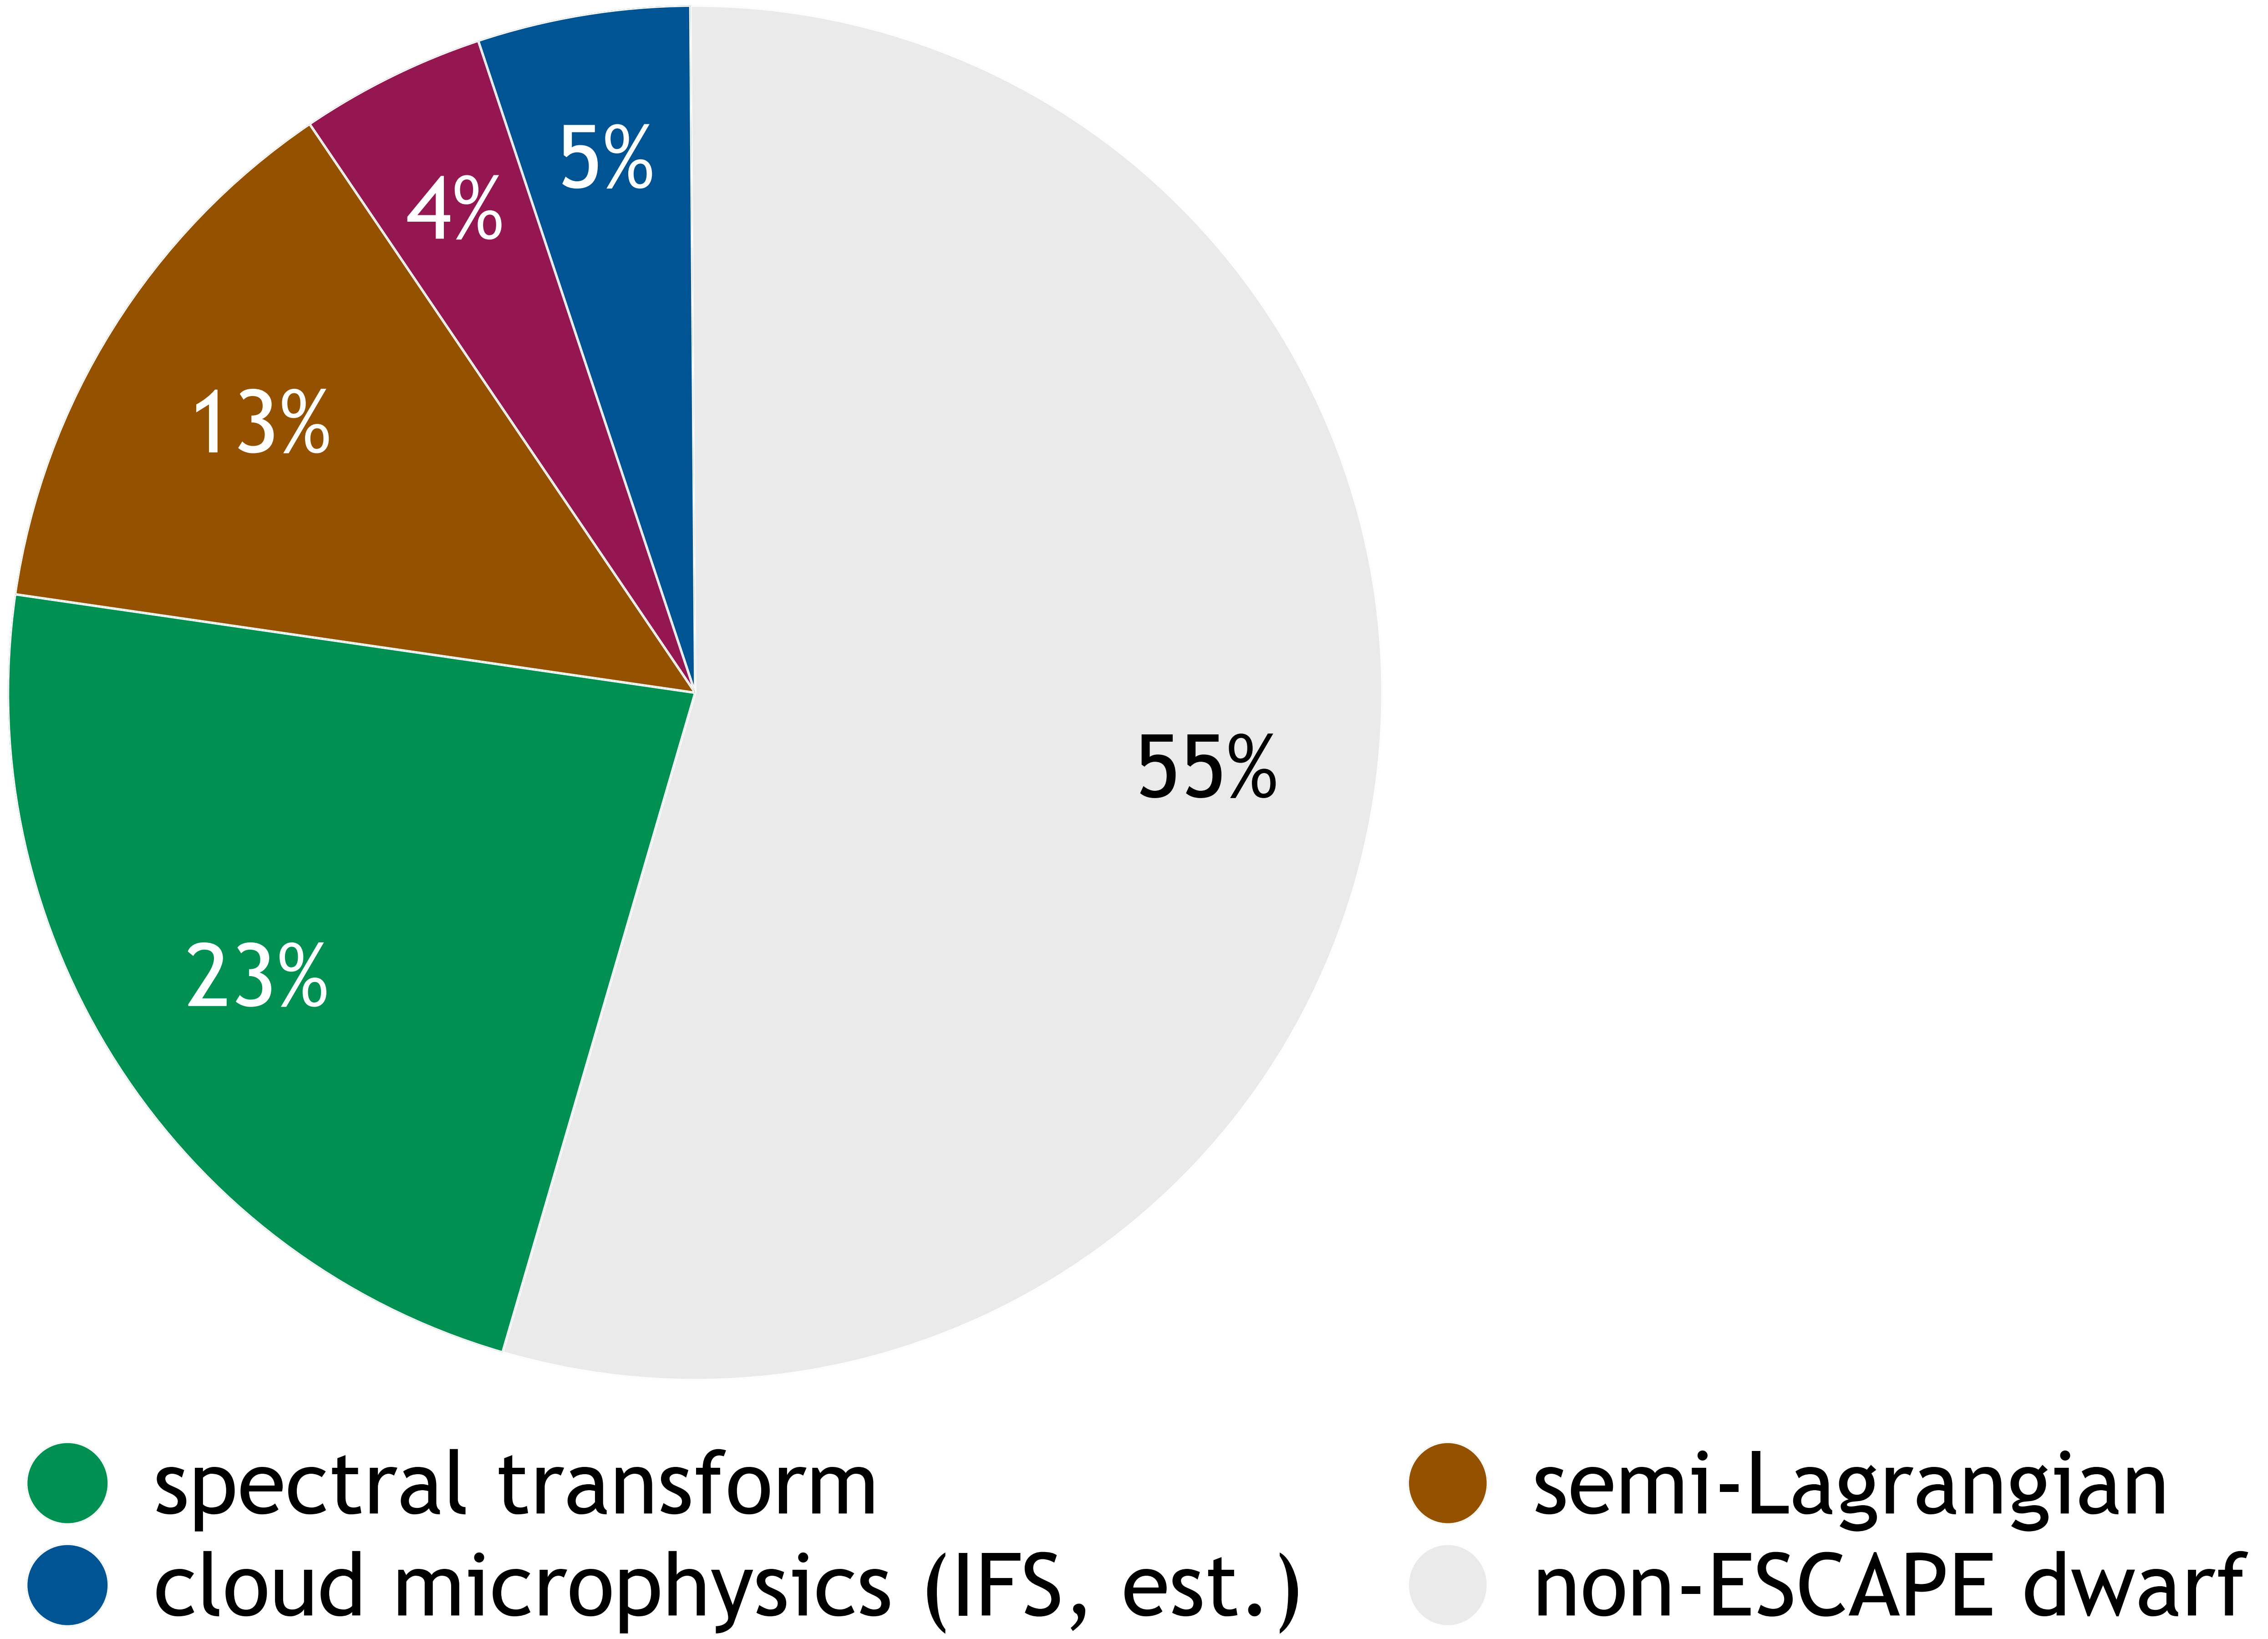

Relevance of our dwarfs

% of the forecast runtime

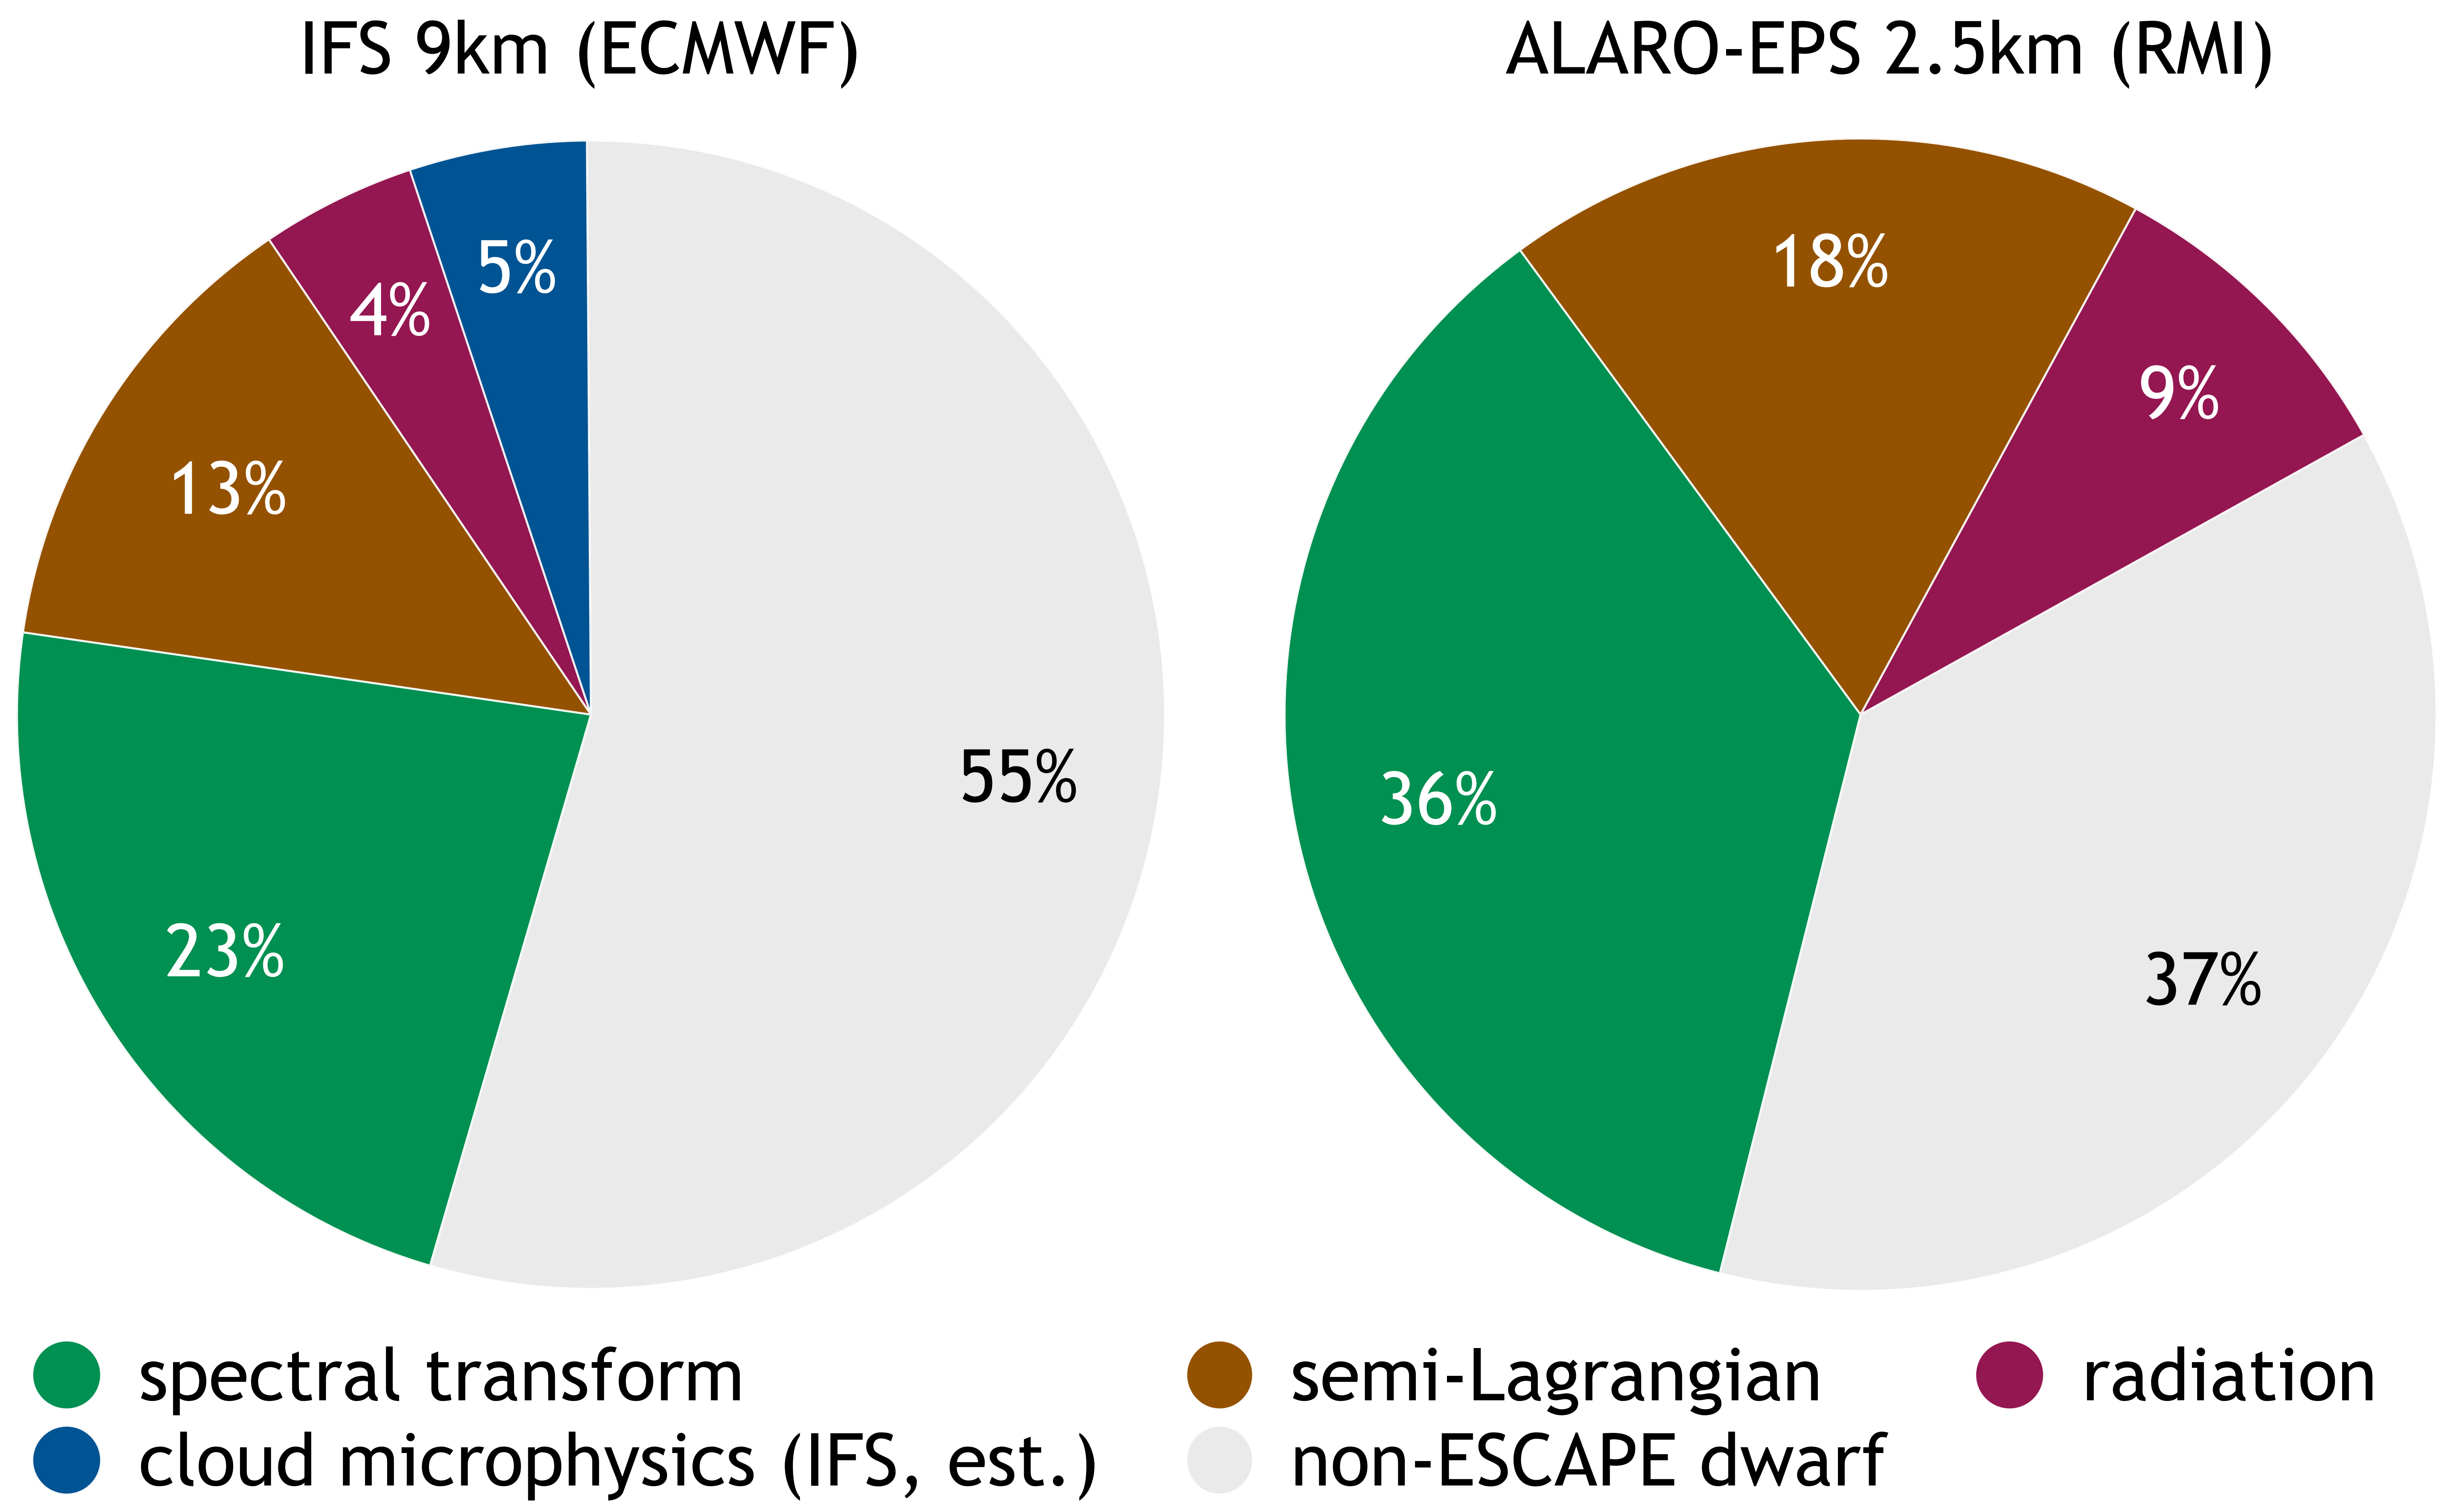

Relevance of our dwarfs

% of the forecast runtime

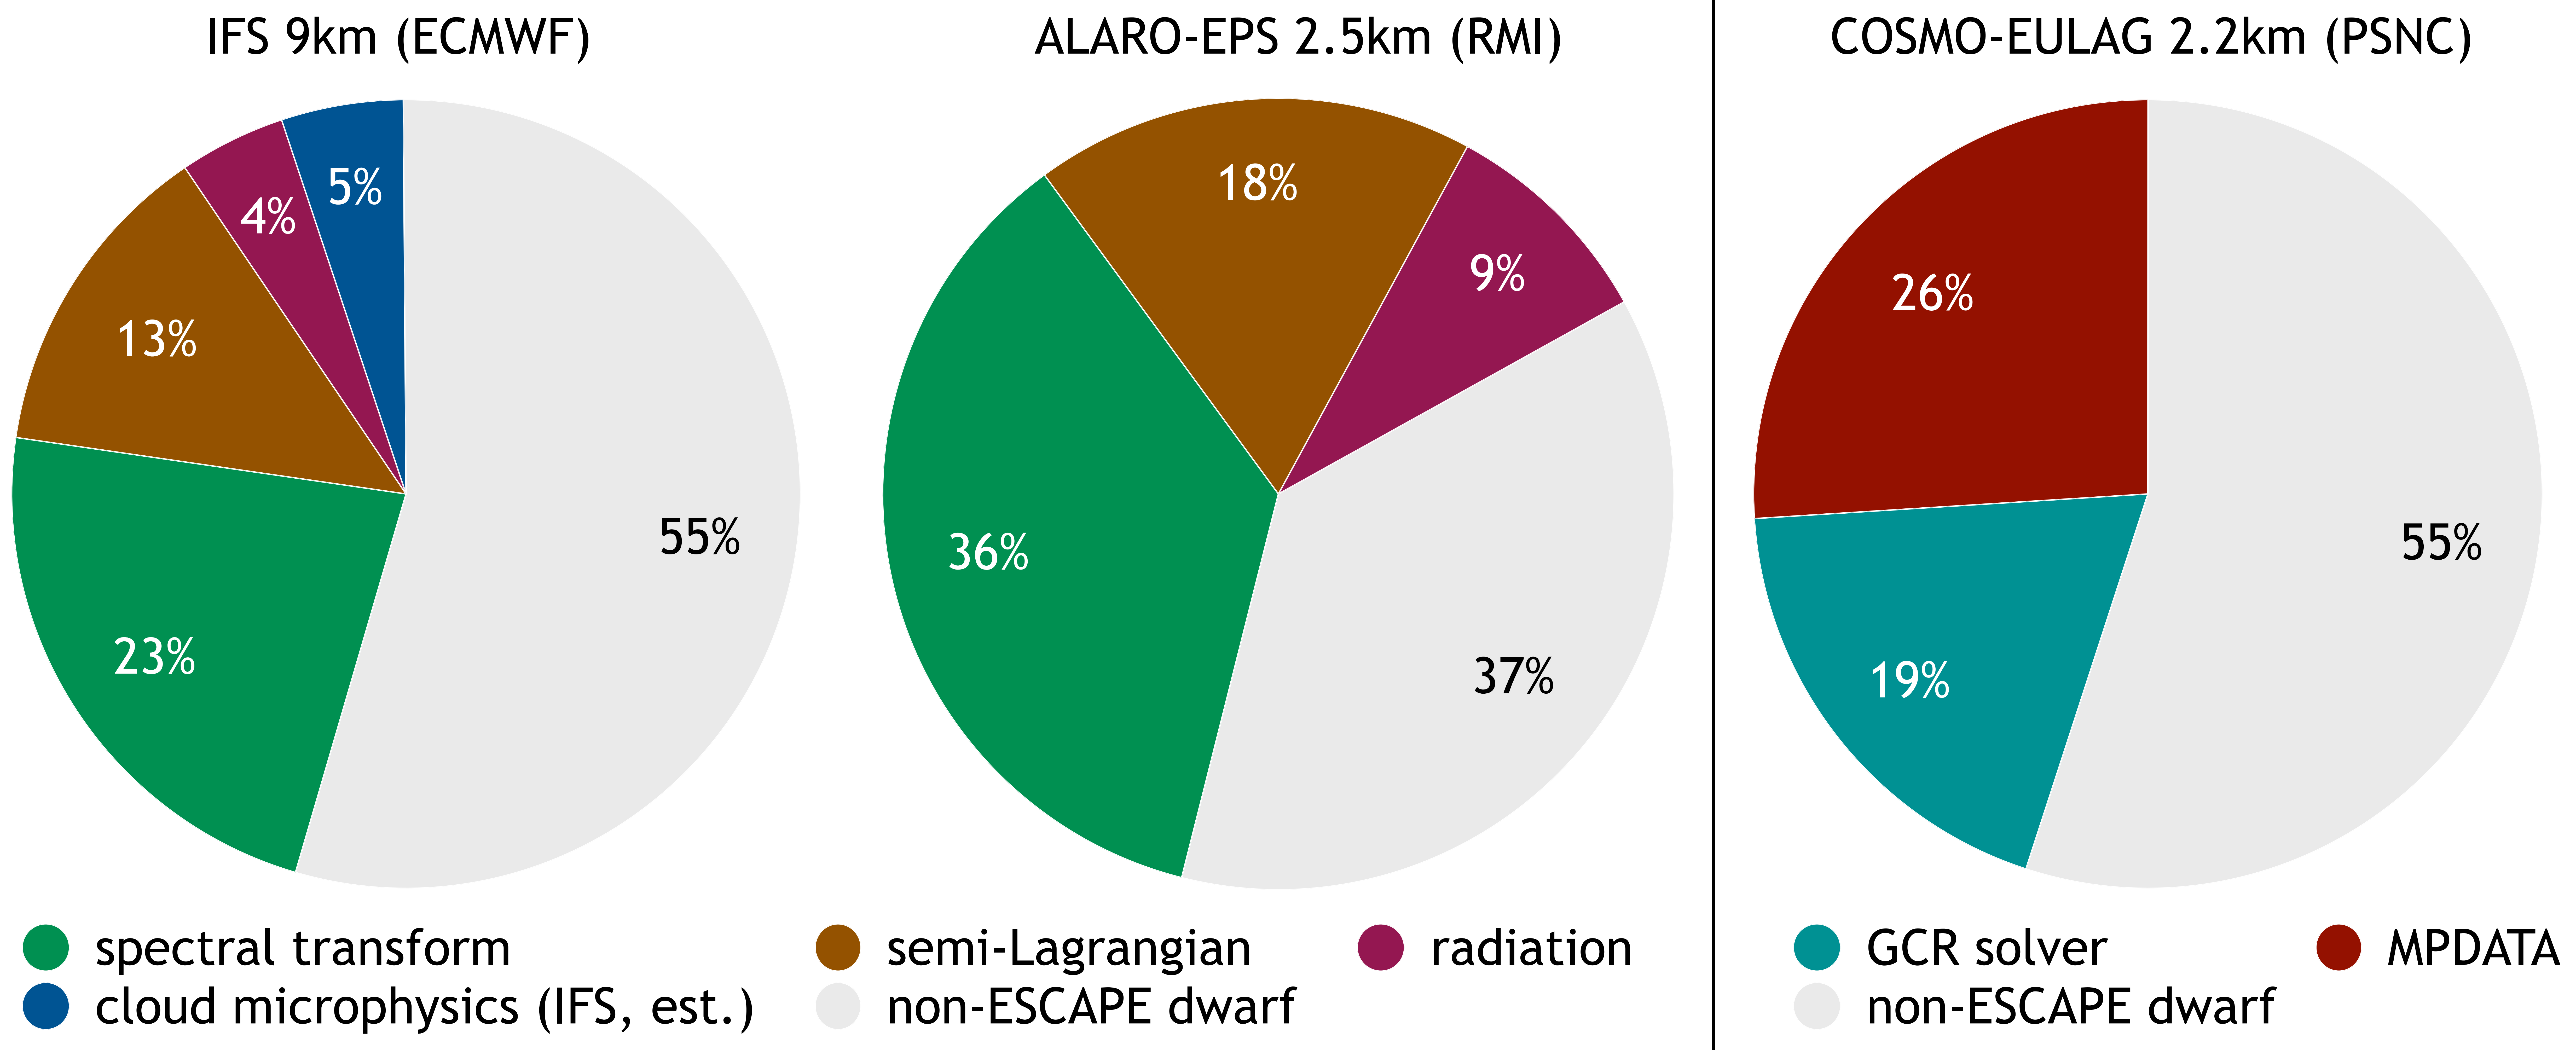

simplified model for MPI communication

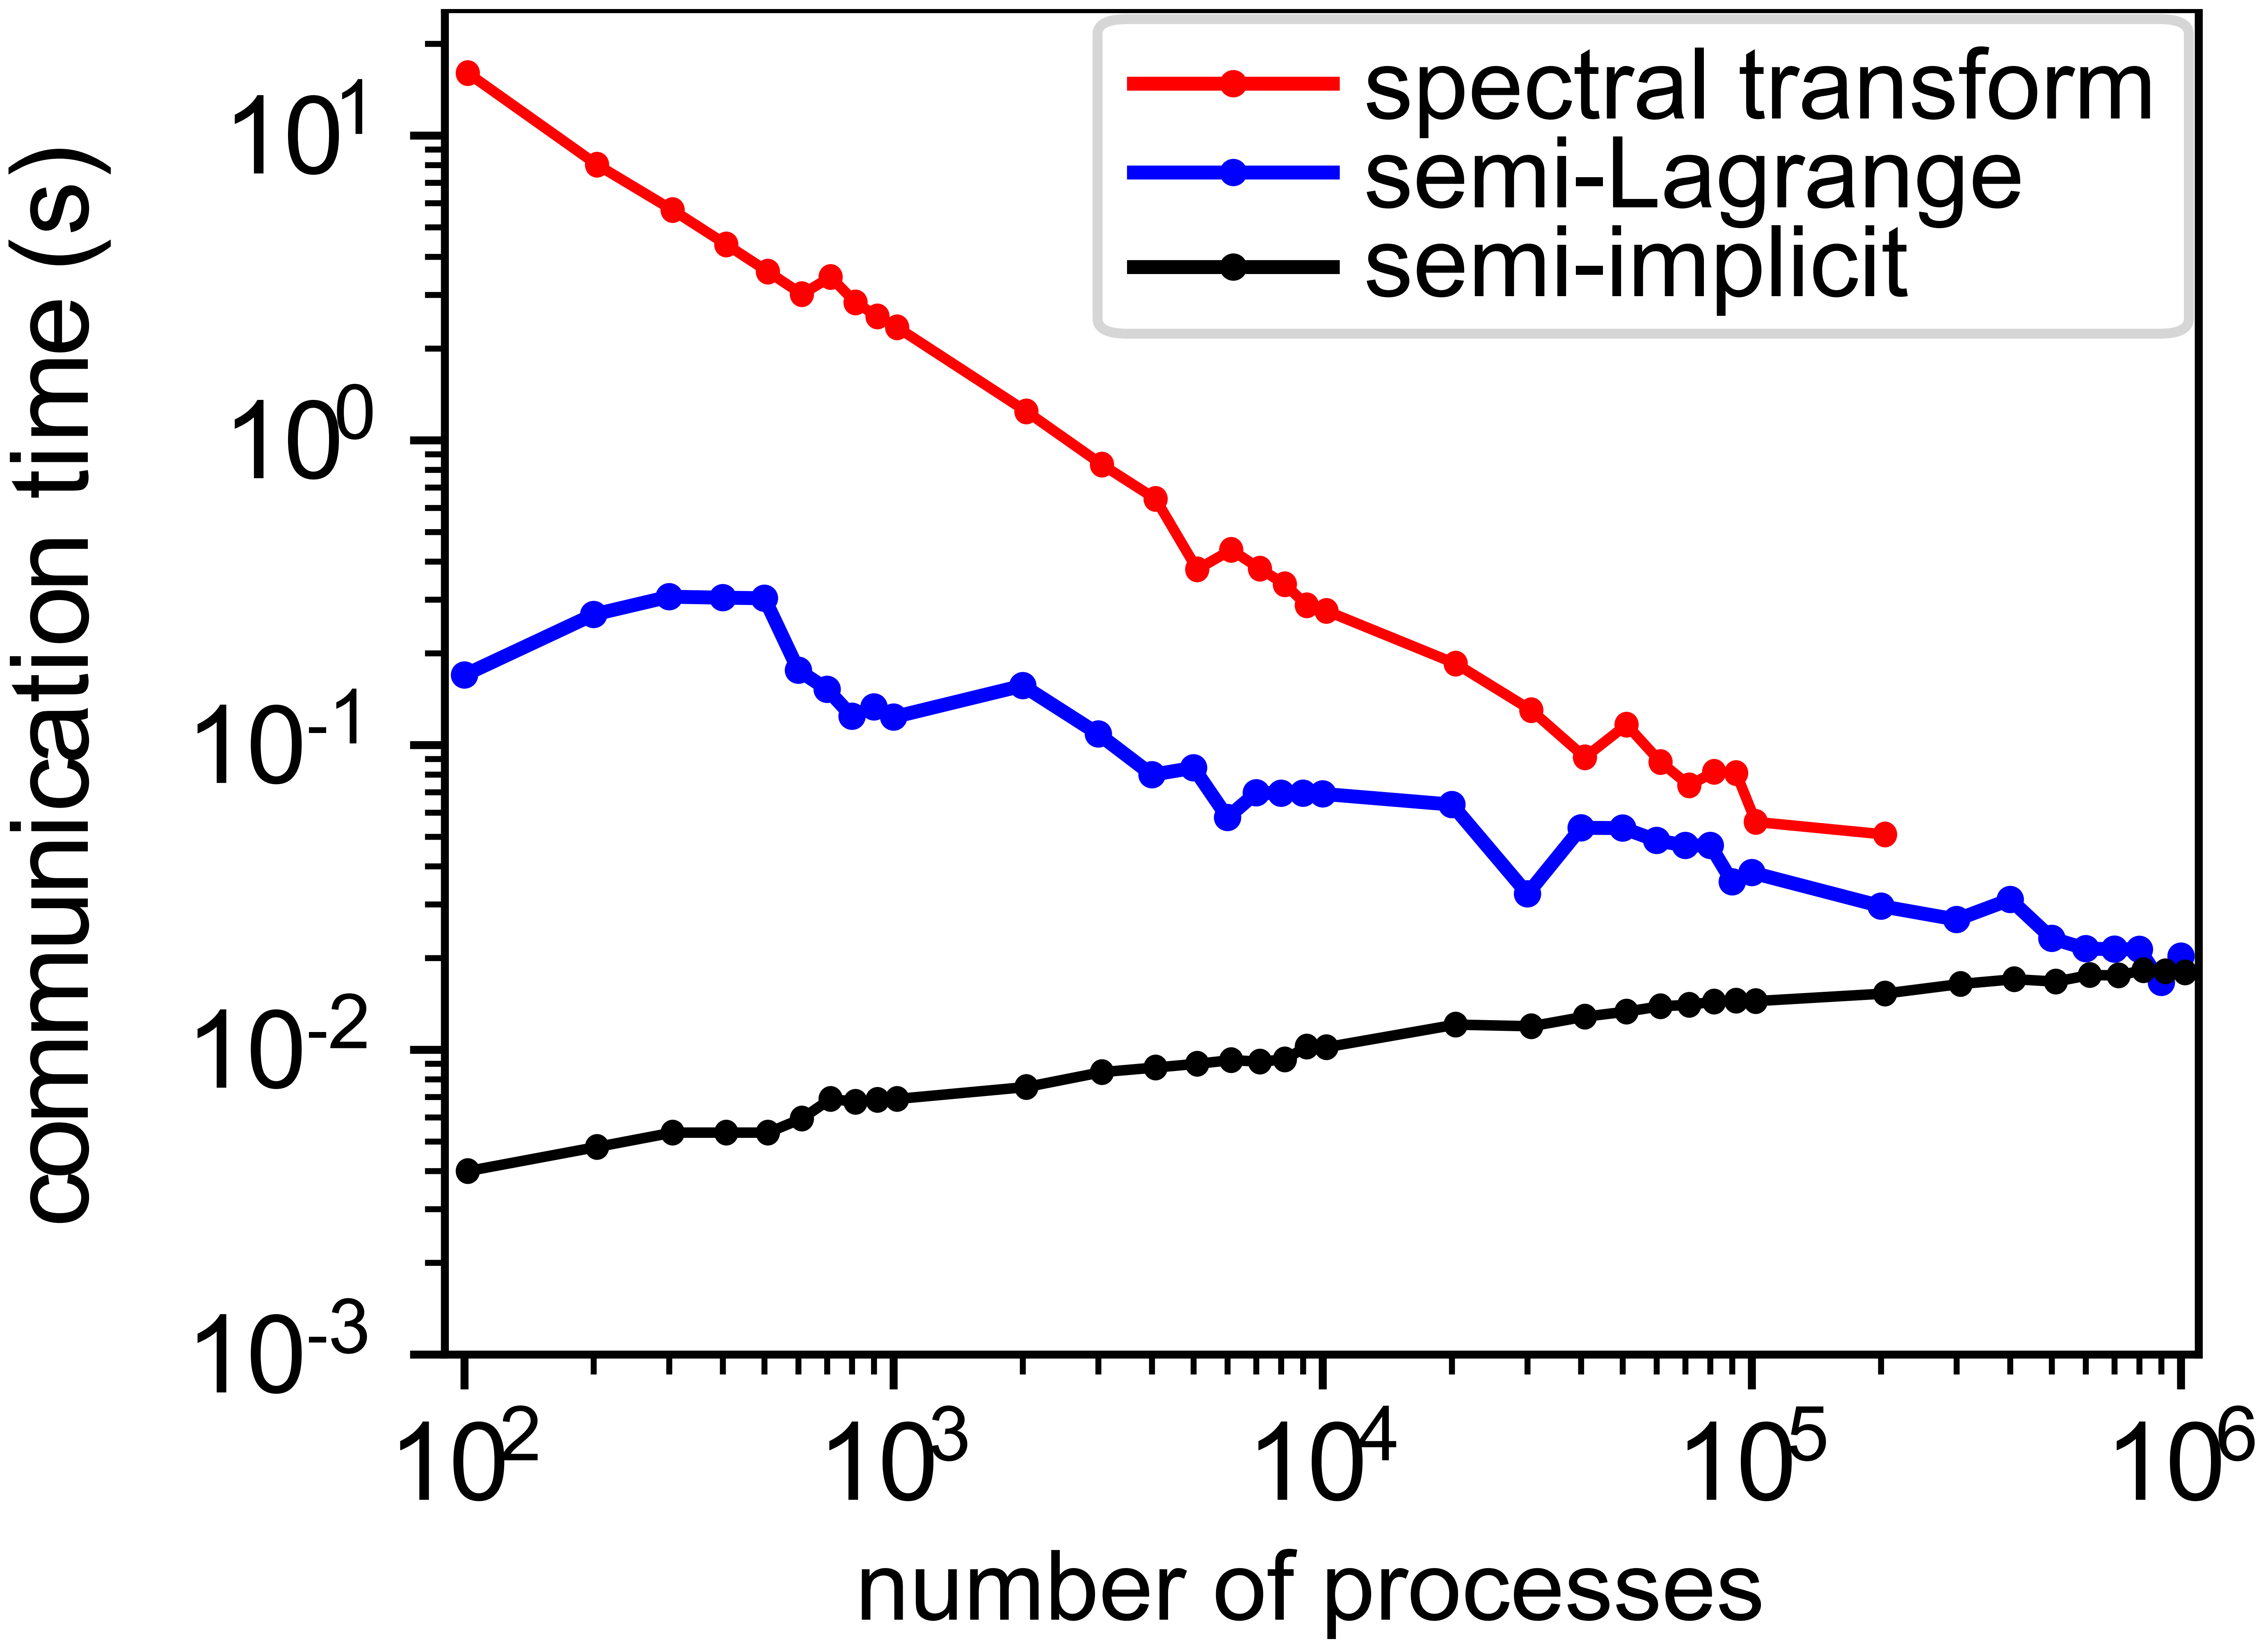

Zheng and Marguinaud (2018)

Funded by the European Union

IFS on the CPUs of PizDaint and Summit

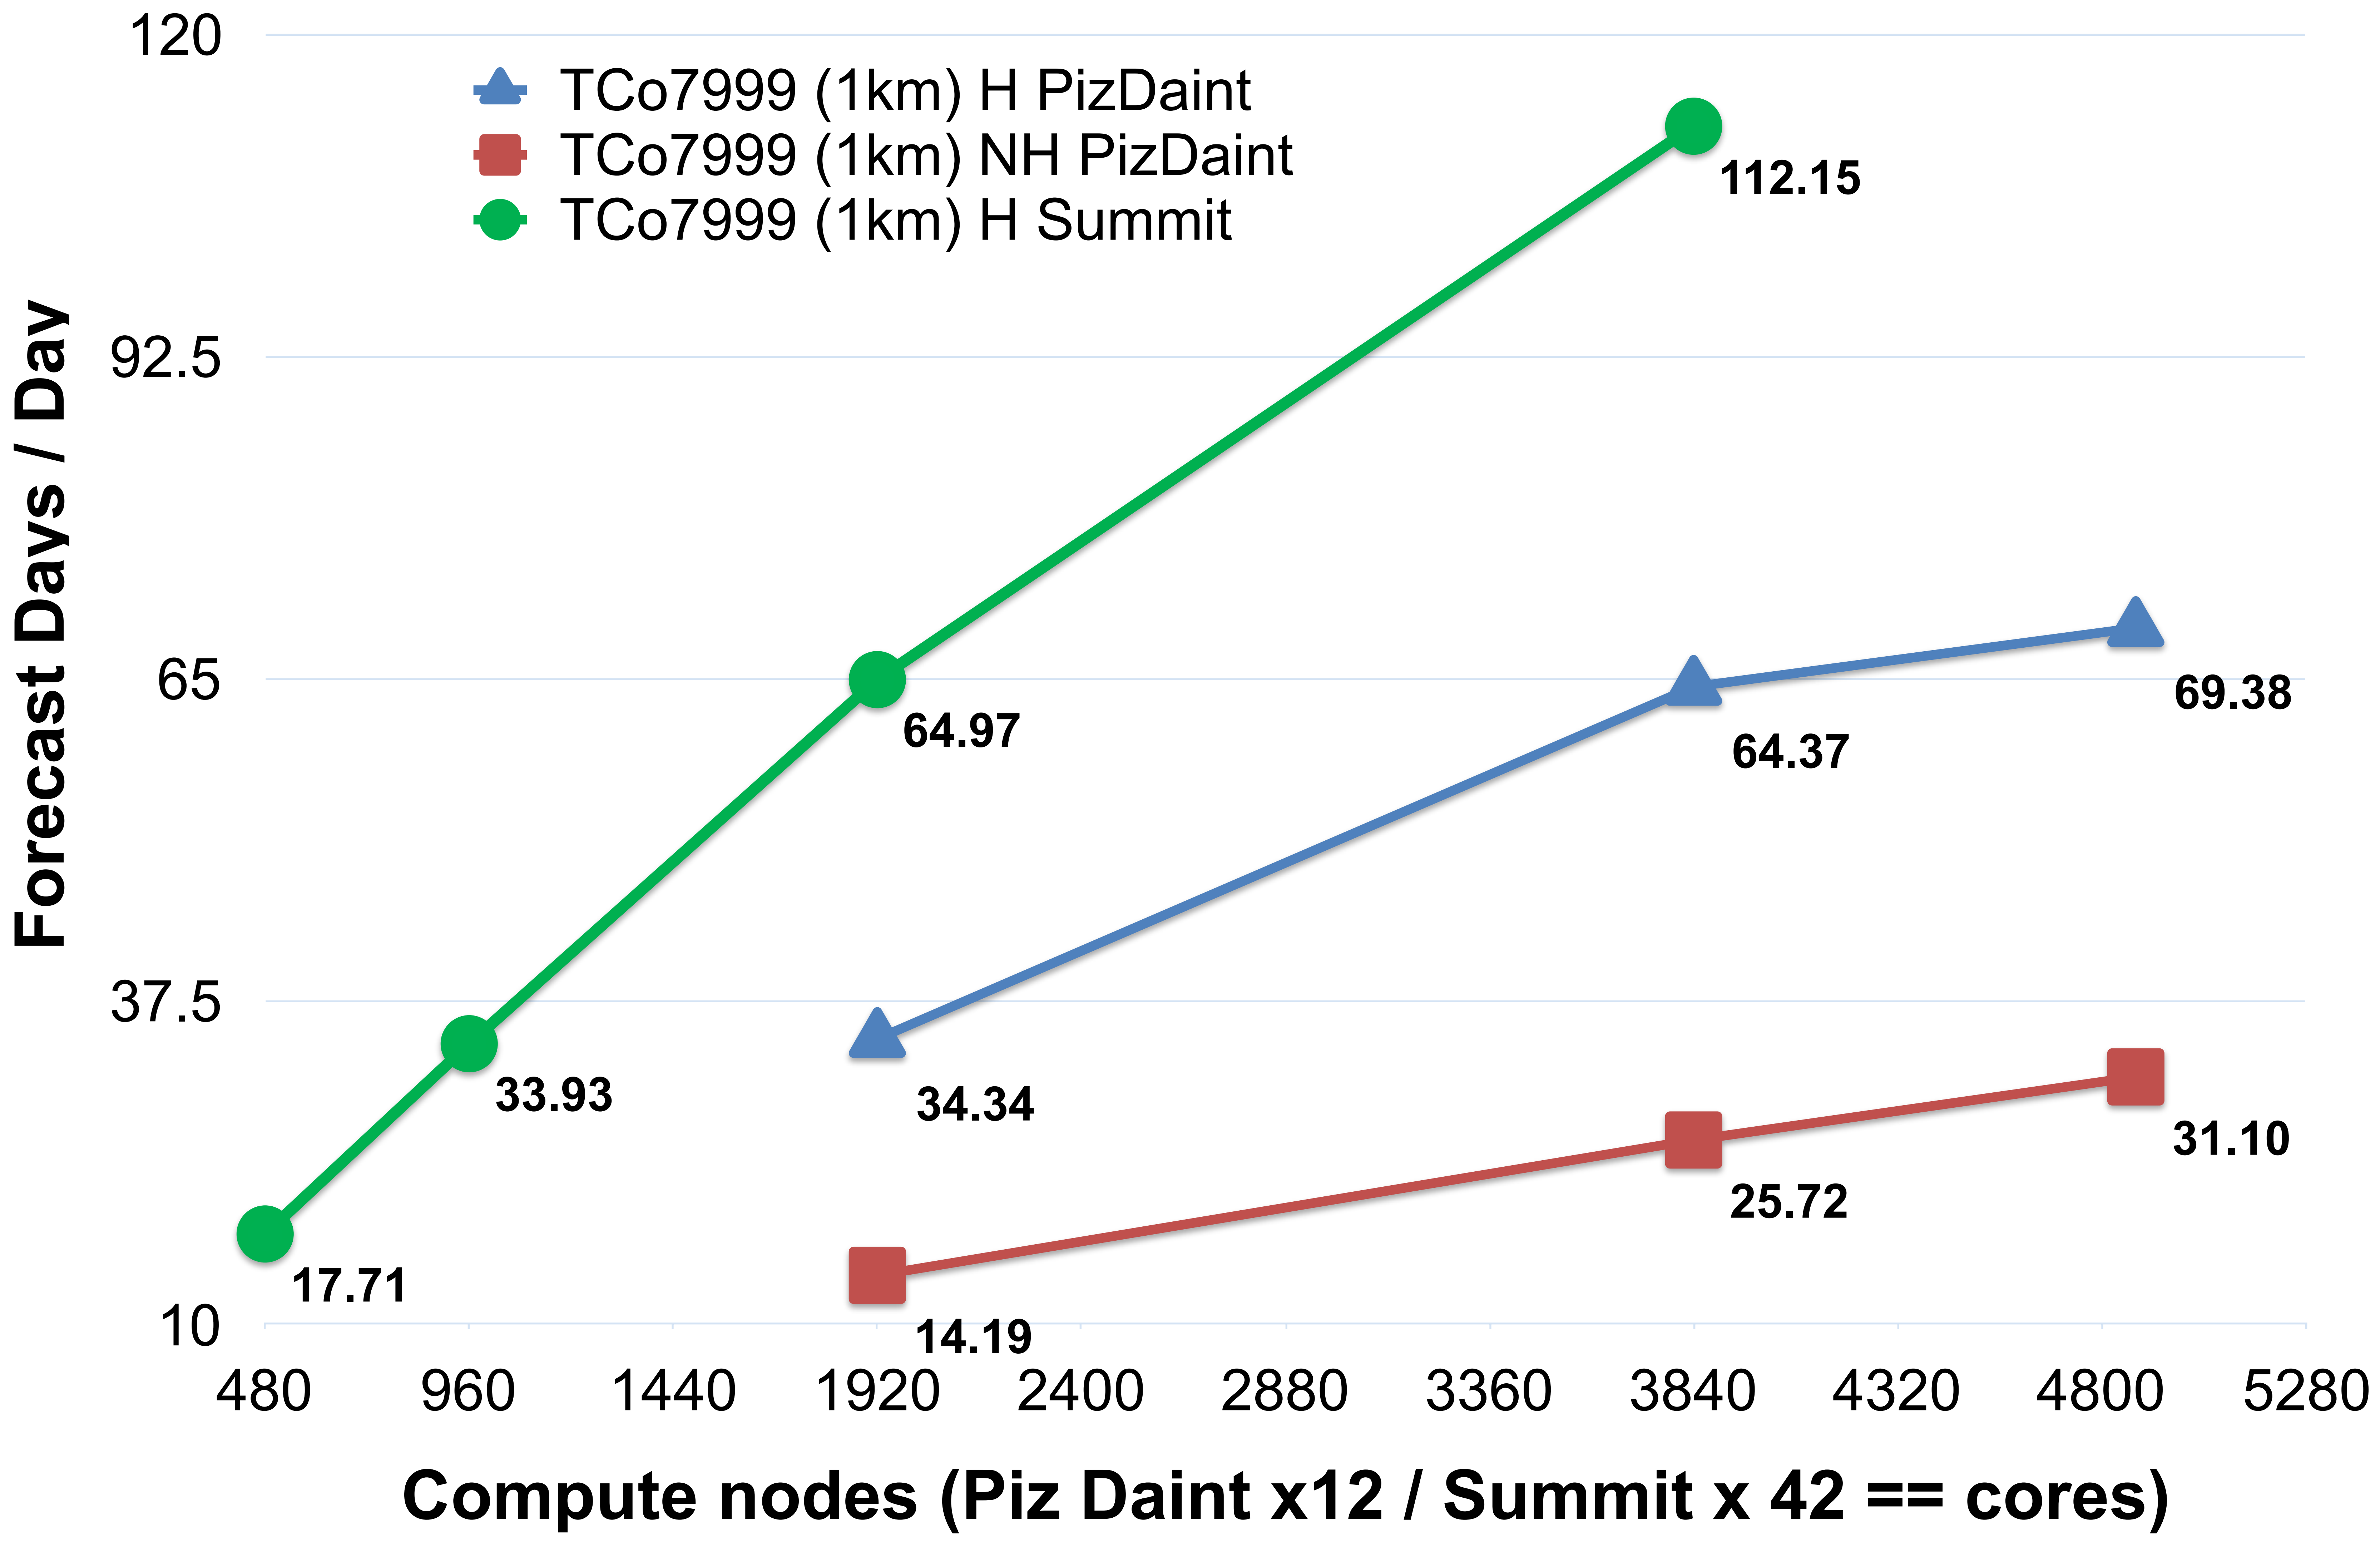

H = hydrostatic
NH = non-hydrostatic

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

Funded by the European Union

optimisations for NVIDIA GPUs

spectral transform at resolution TCo639 (18km)

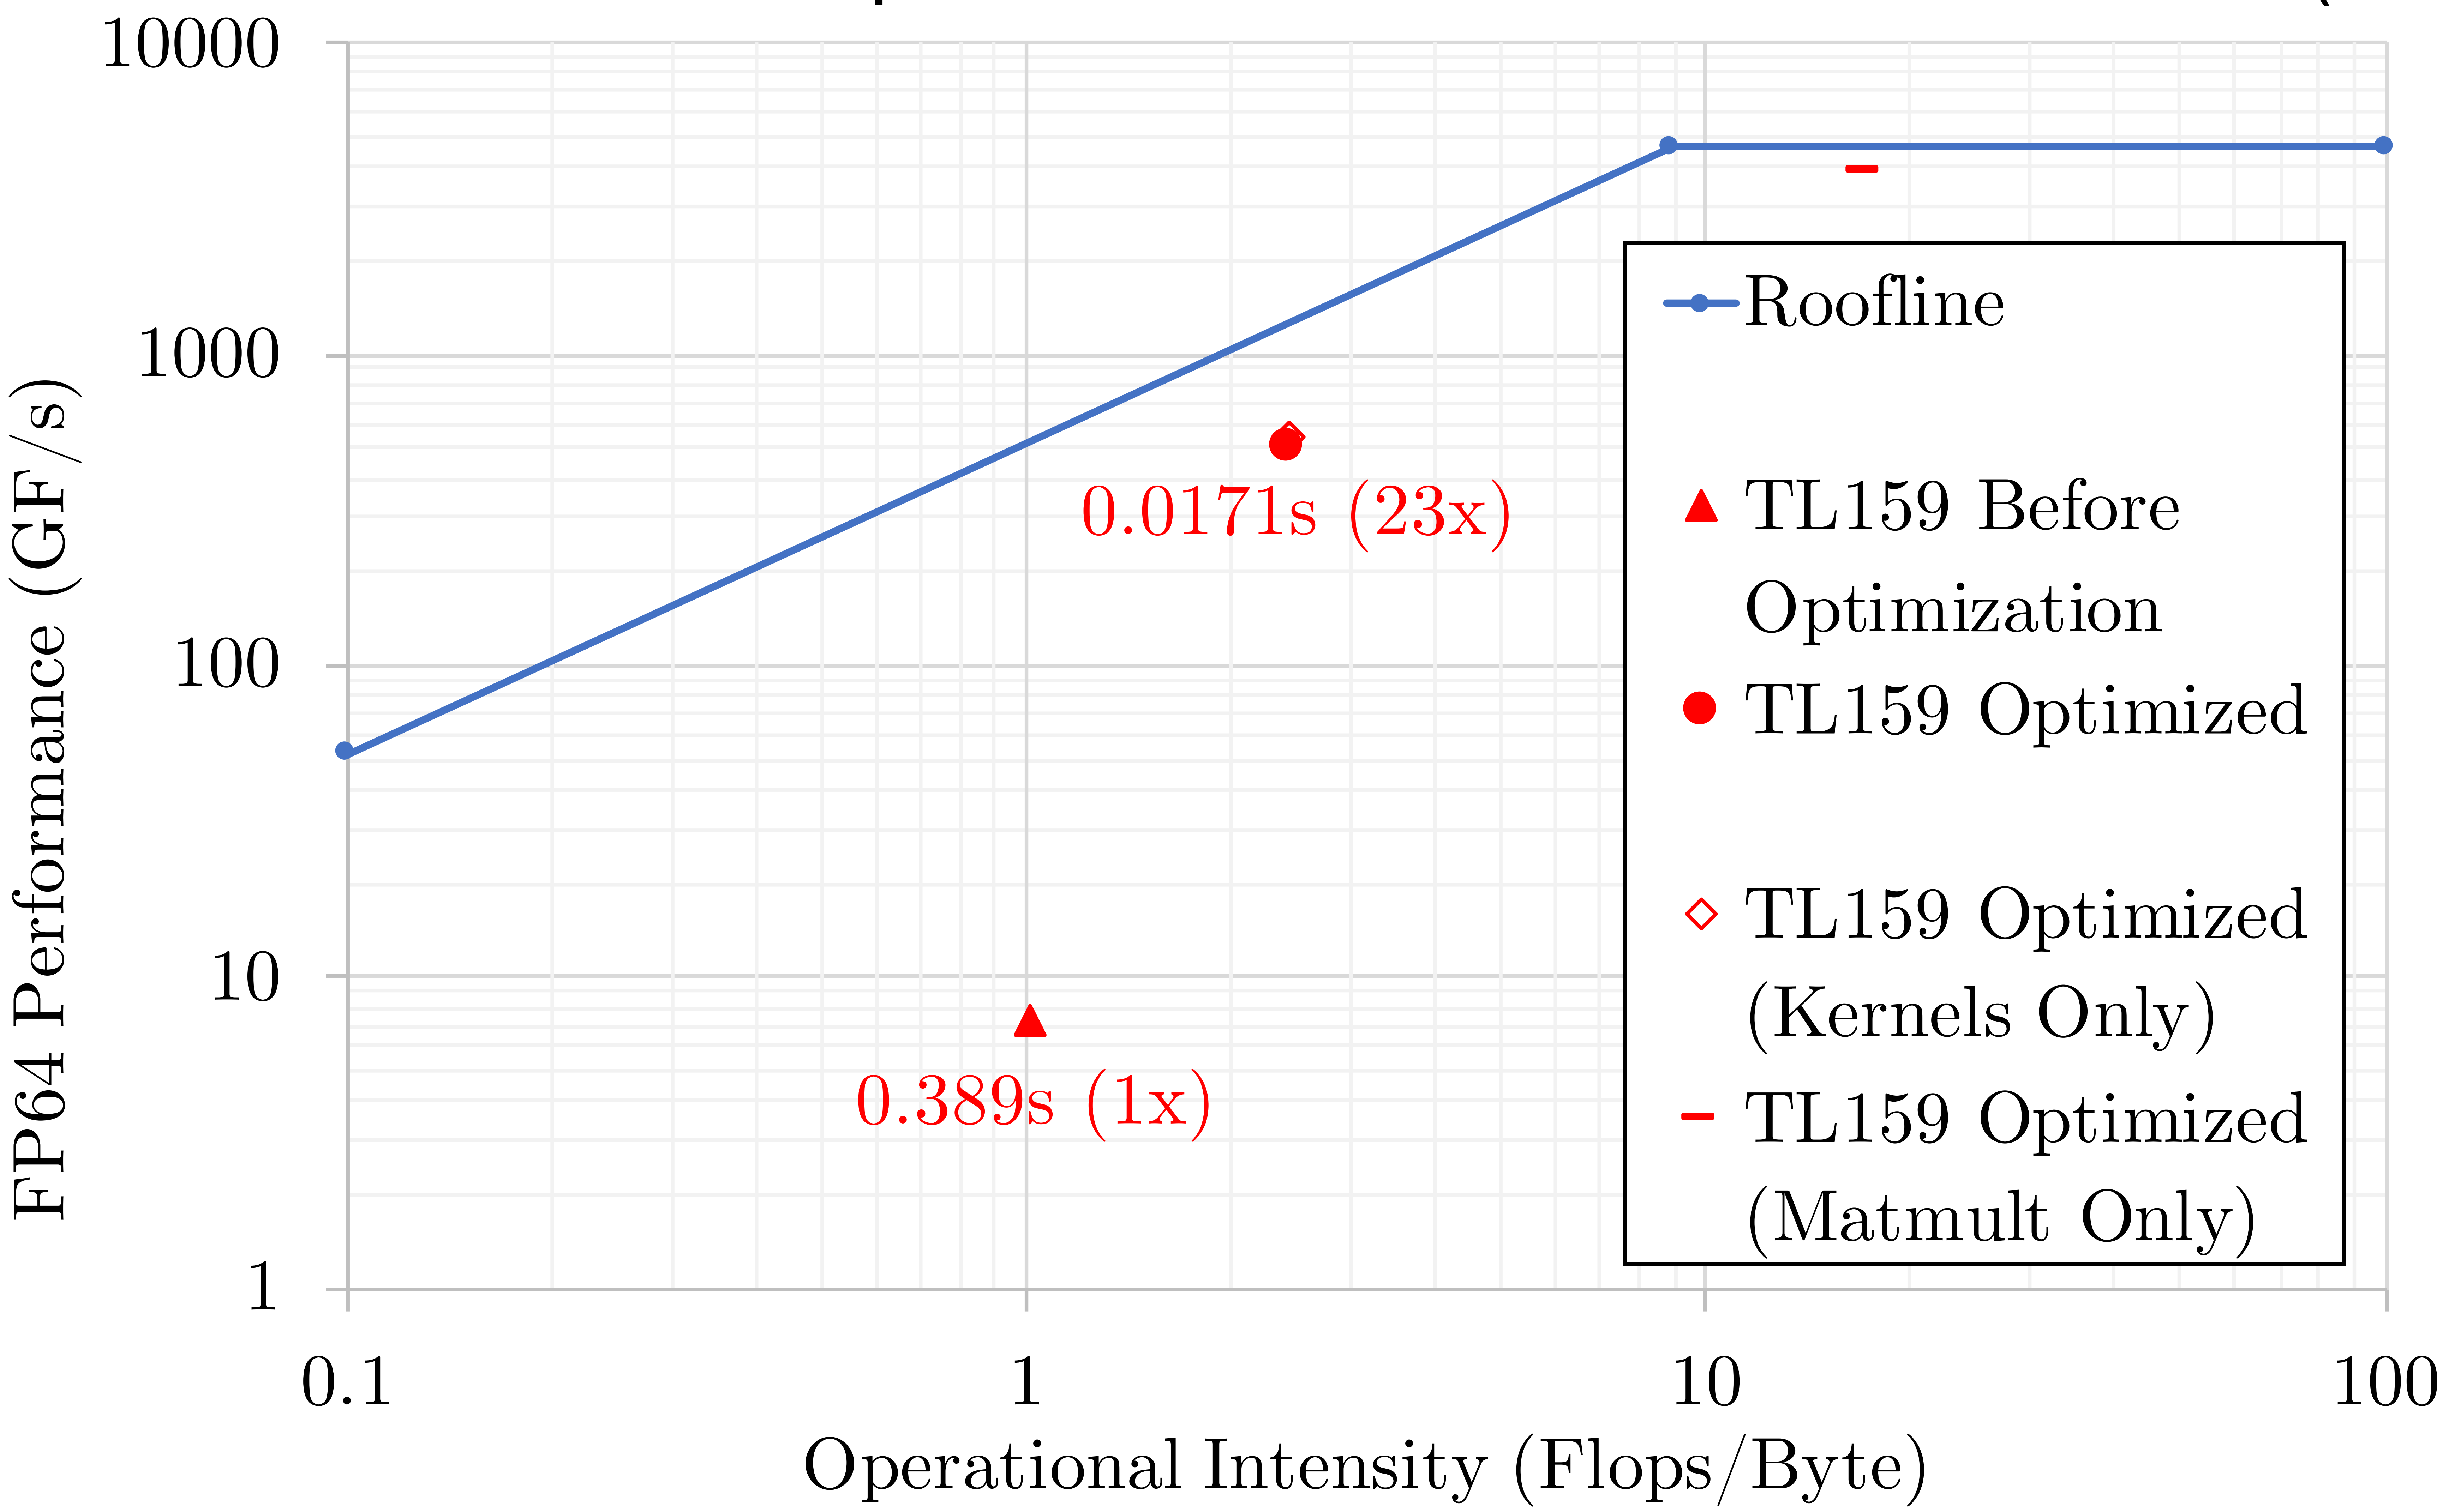

optimised version includes zero operations for padding matrices for batched matrix multiplication

spectral transform
DGX-2 with NVSwitch vs multiple DGX-1V

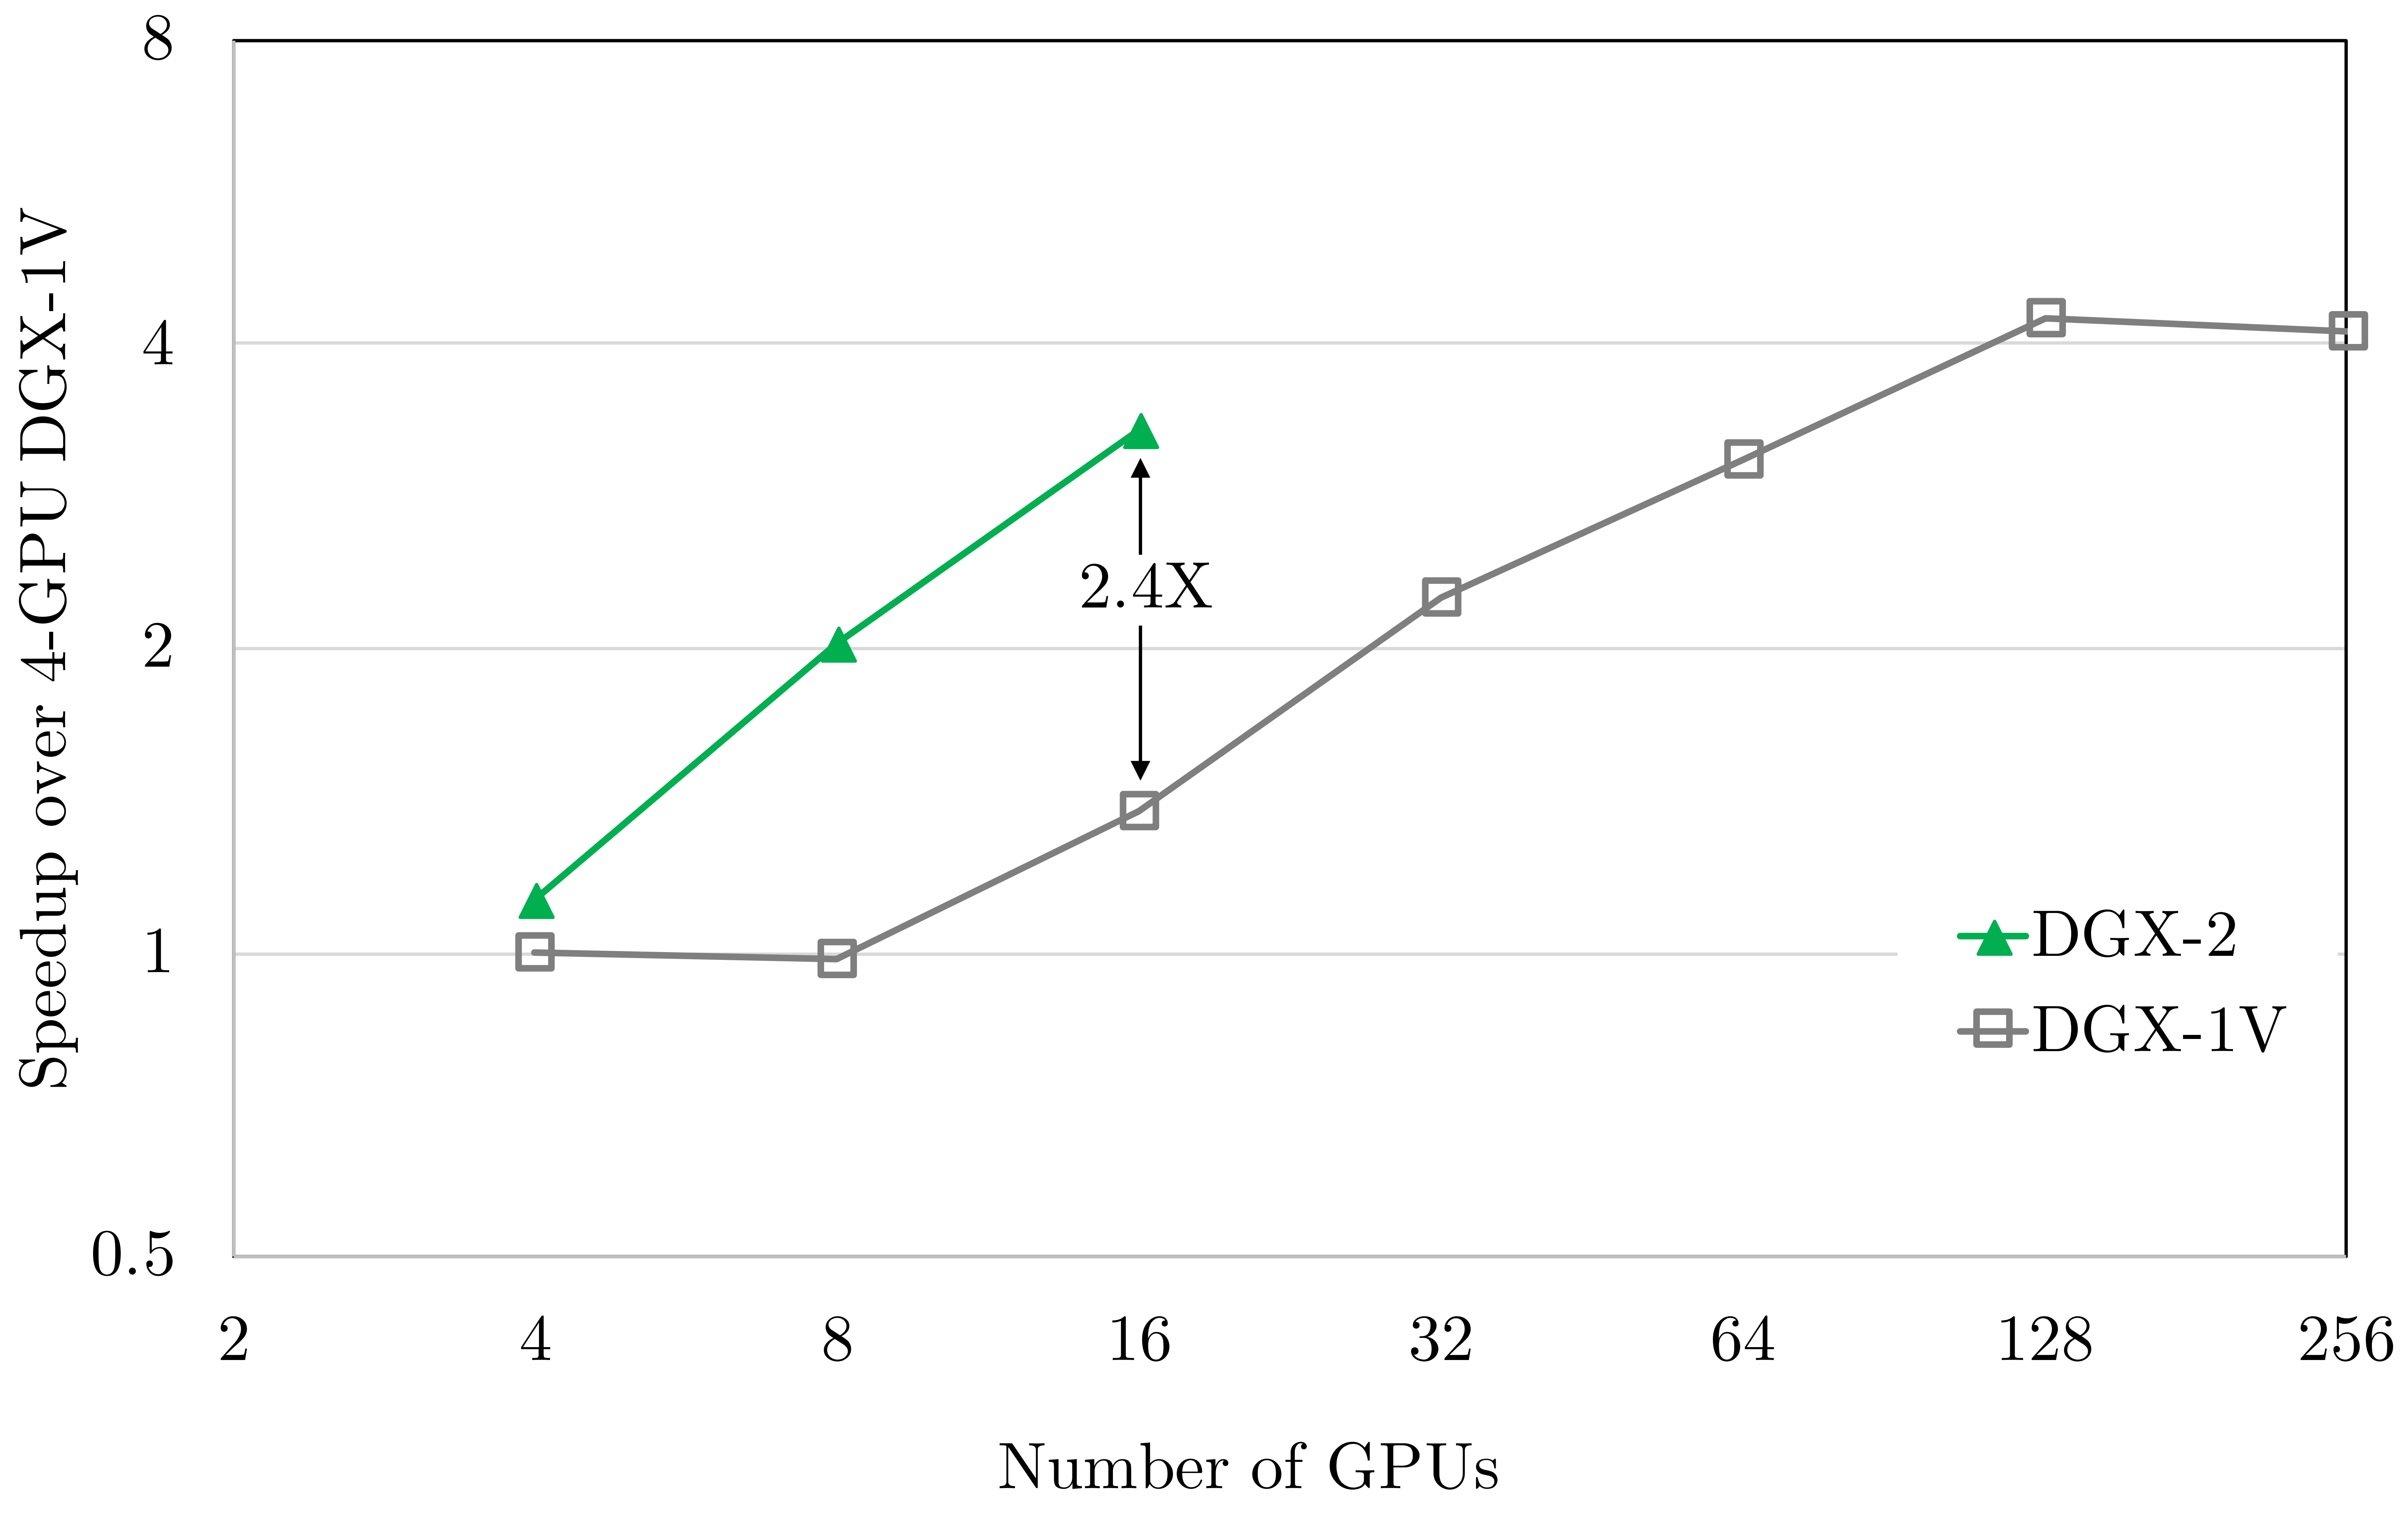

Funded by the European Union

Funded by the European Union

Spectral transform dwarf on Summit

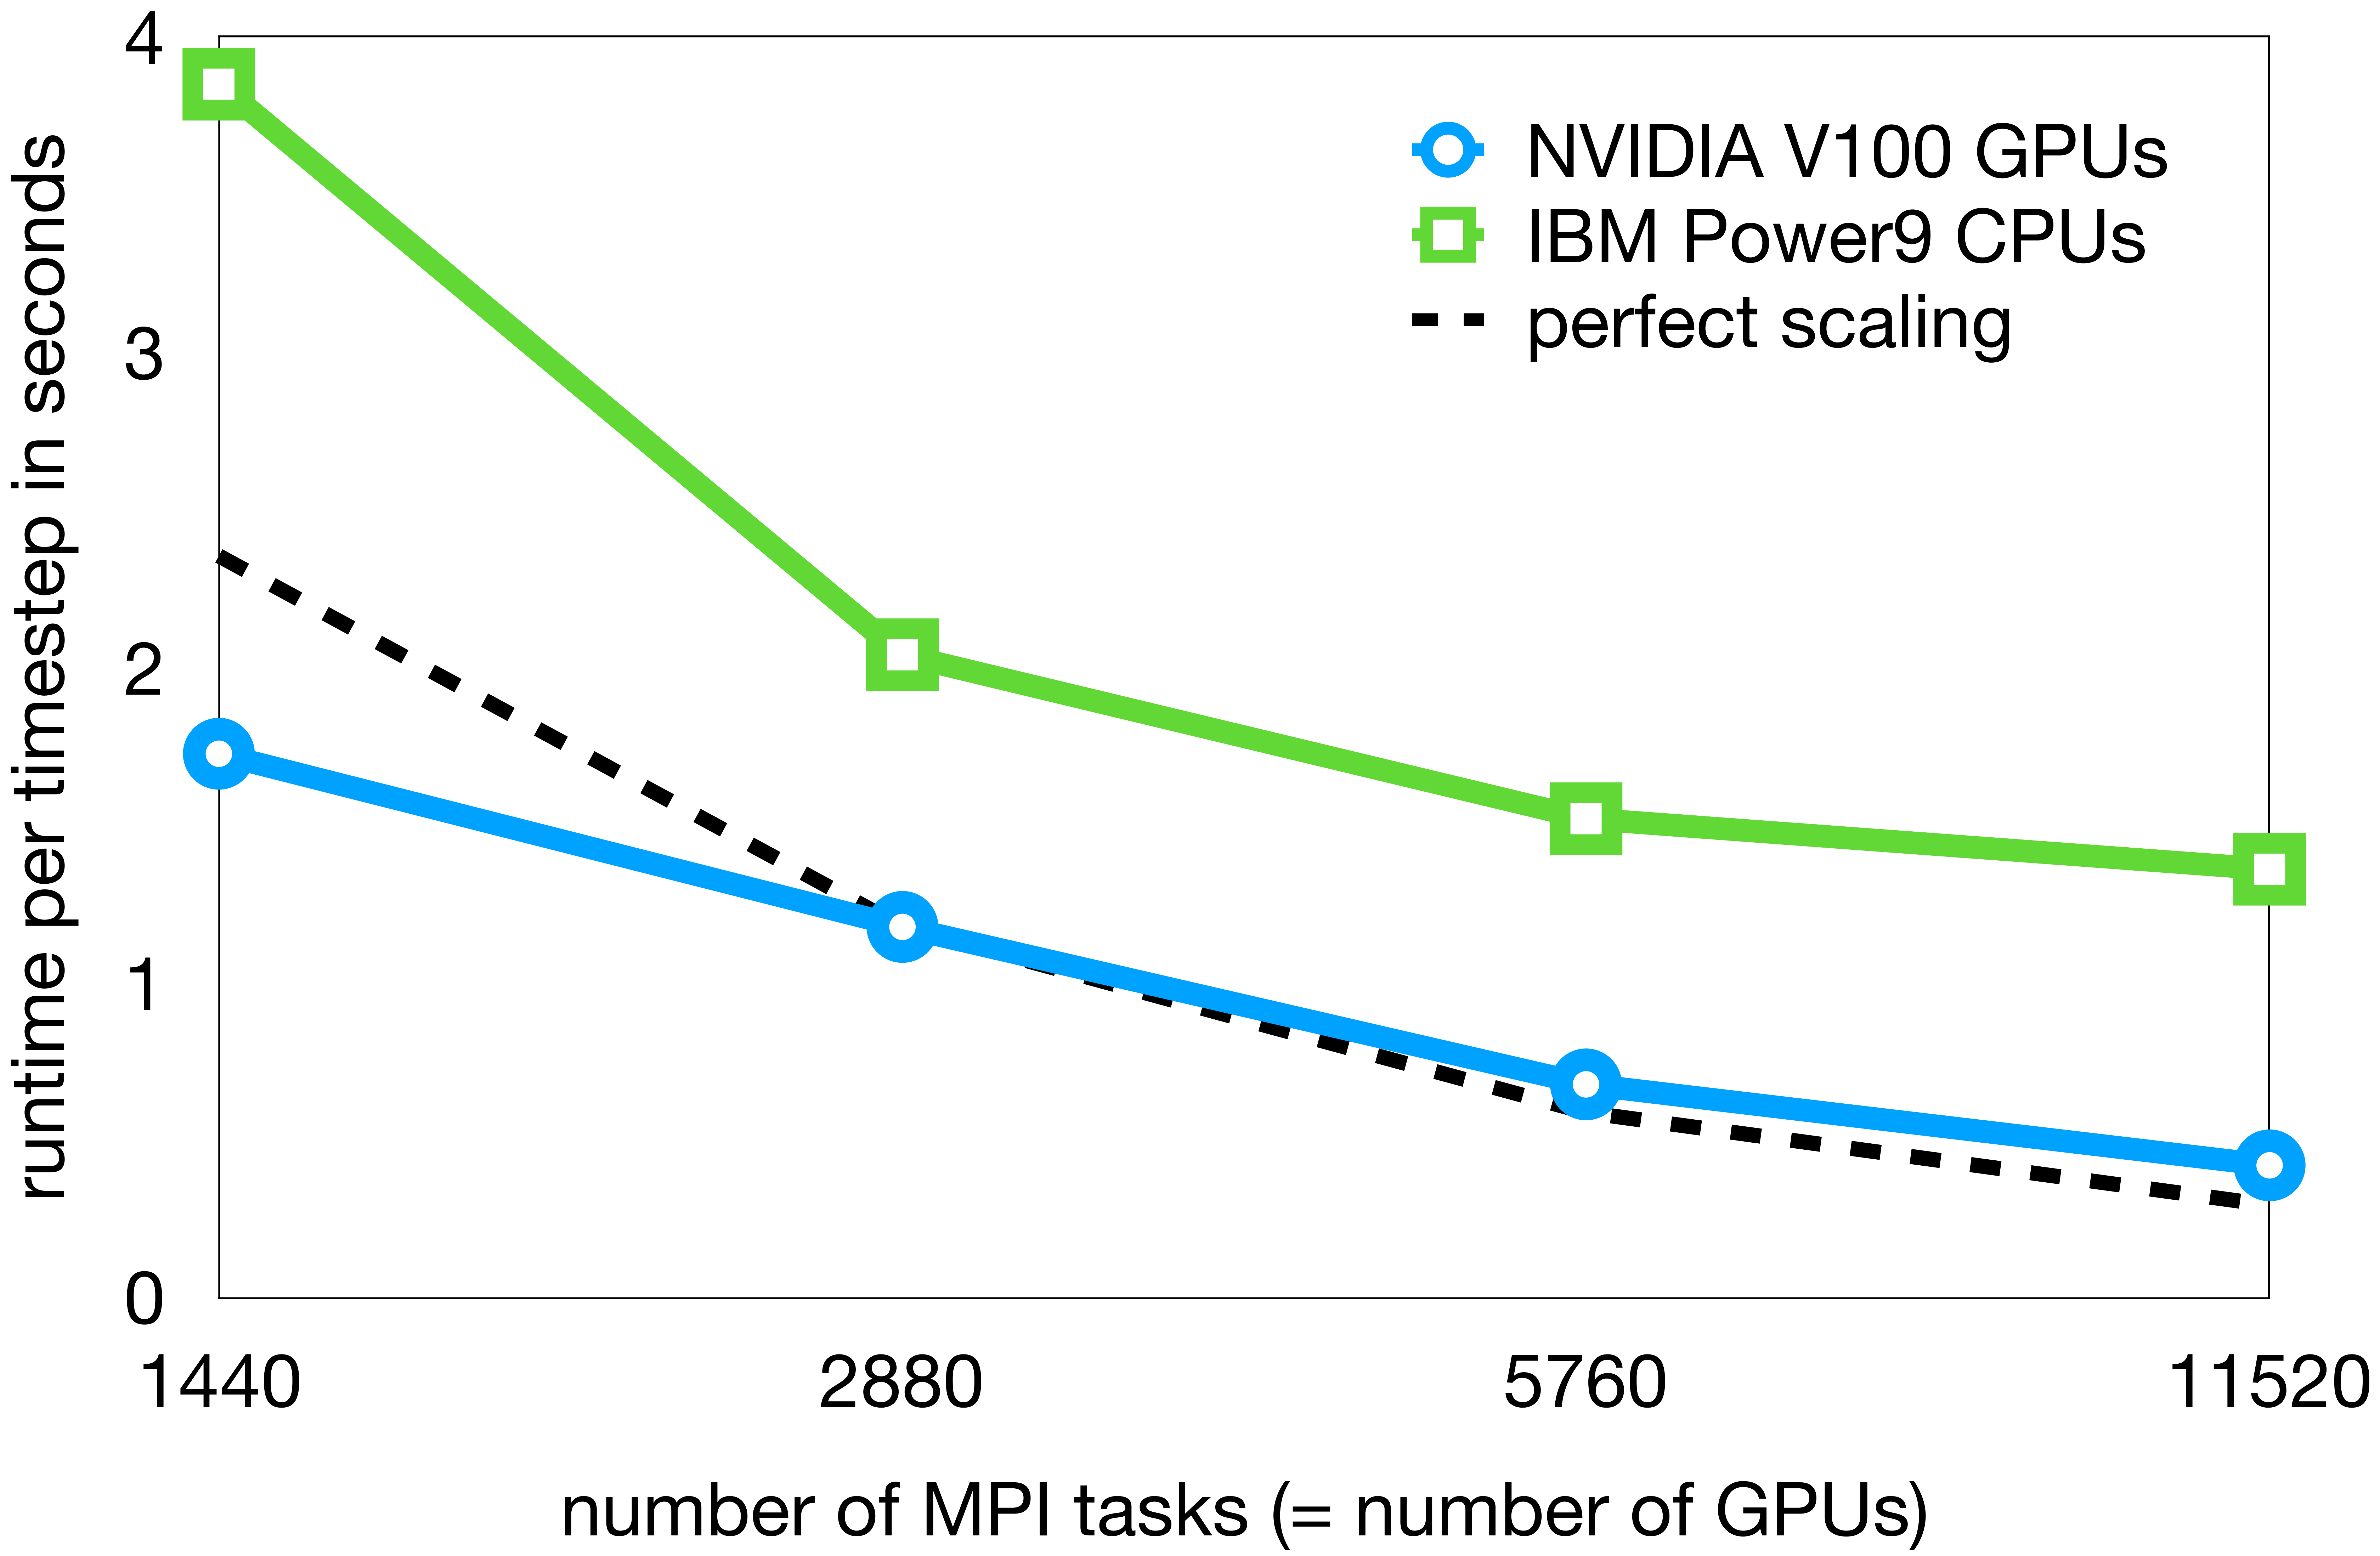

At 2.5km resolution, less than 1s per time-step fits operational needs.

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE office of Science User Facility supported under contract DE-AC05-00OR22725.

CPU optimisation

spectral transform at resolution TCo639 (18km)

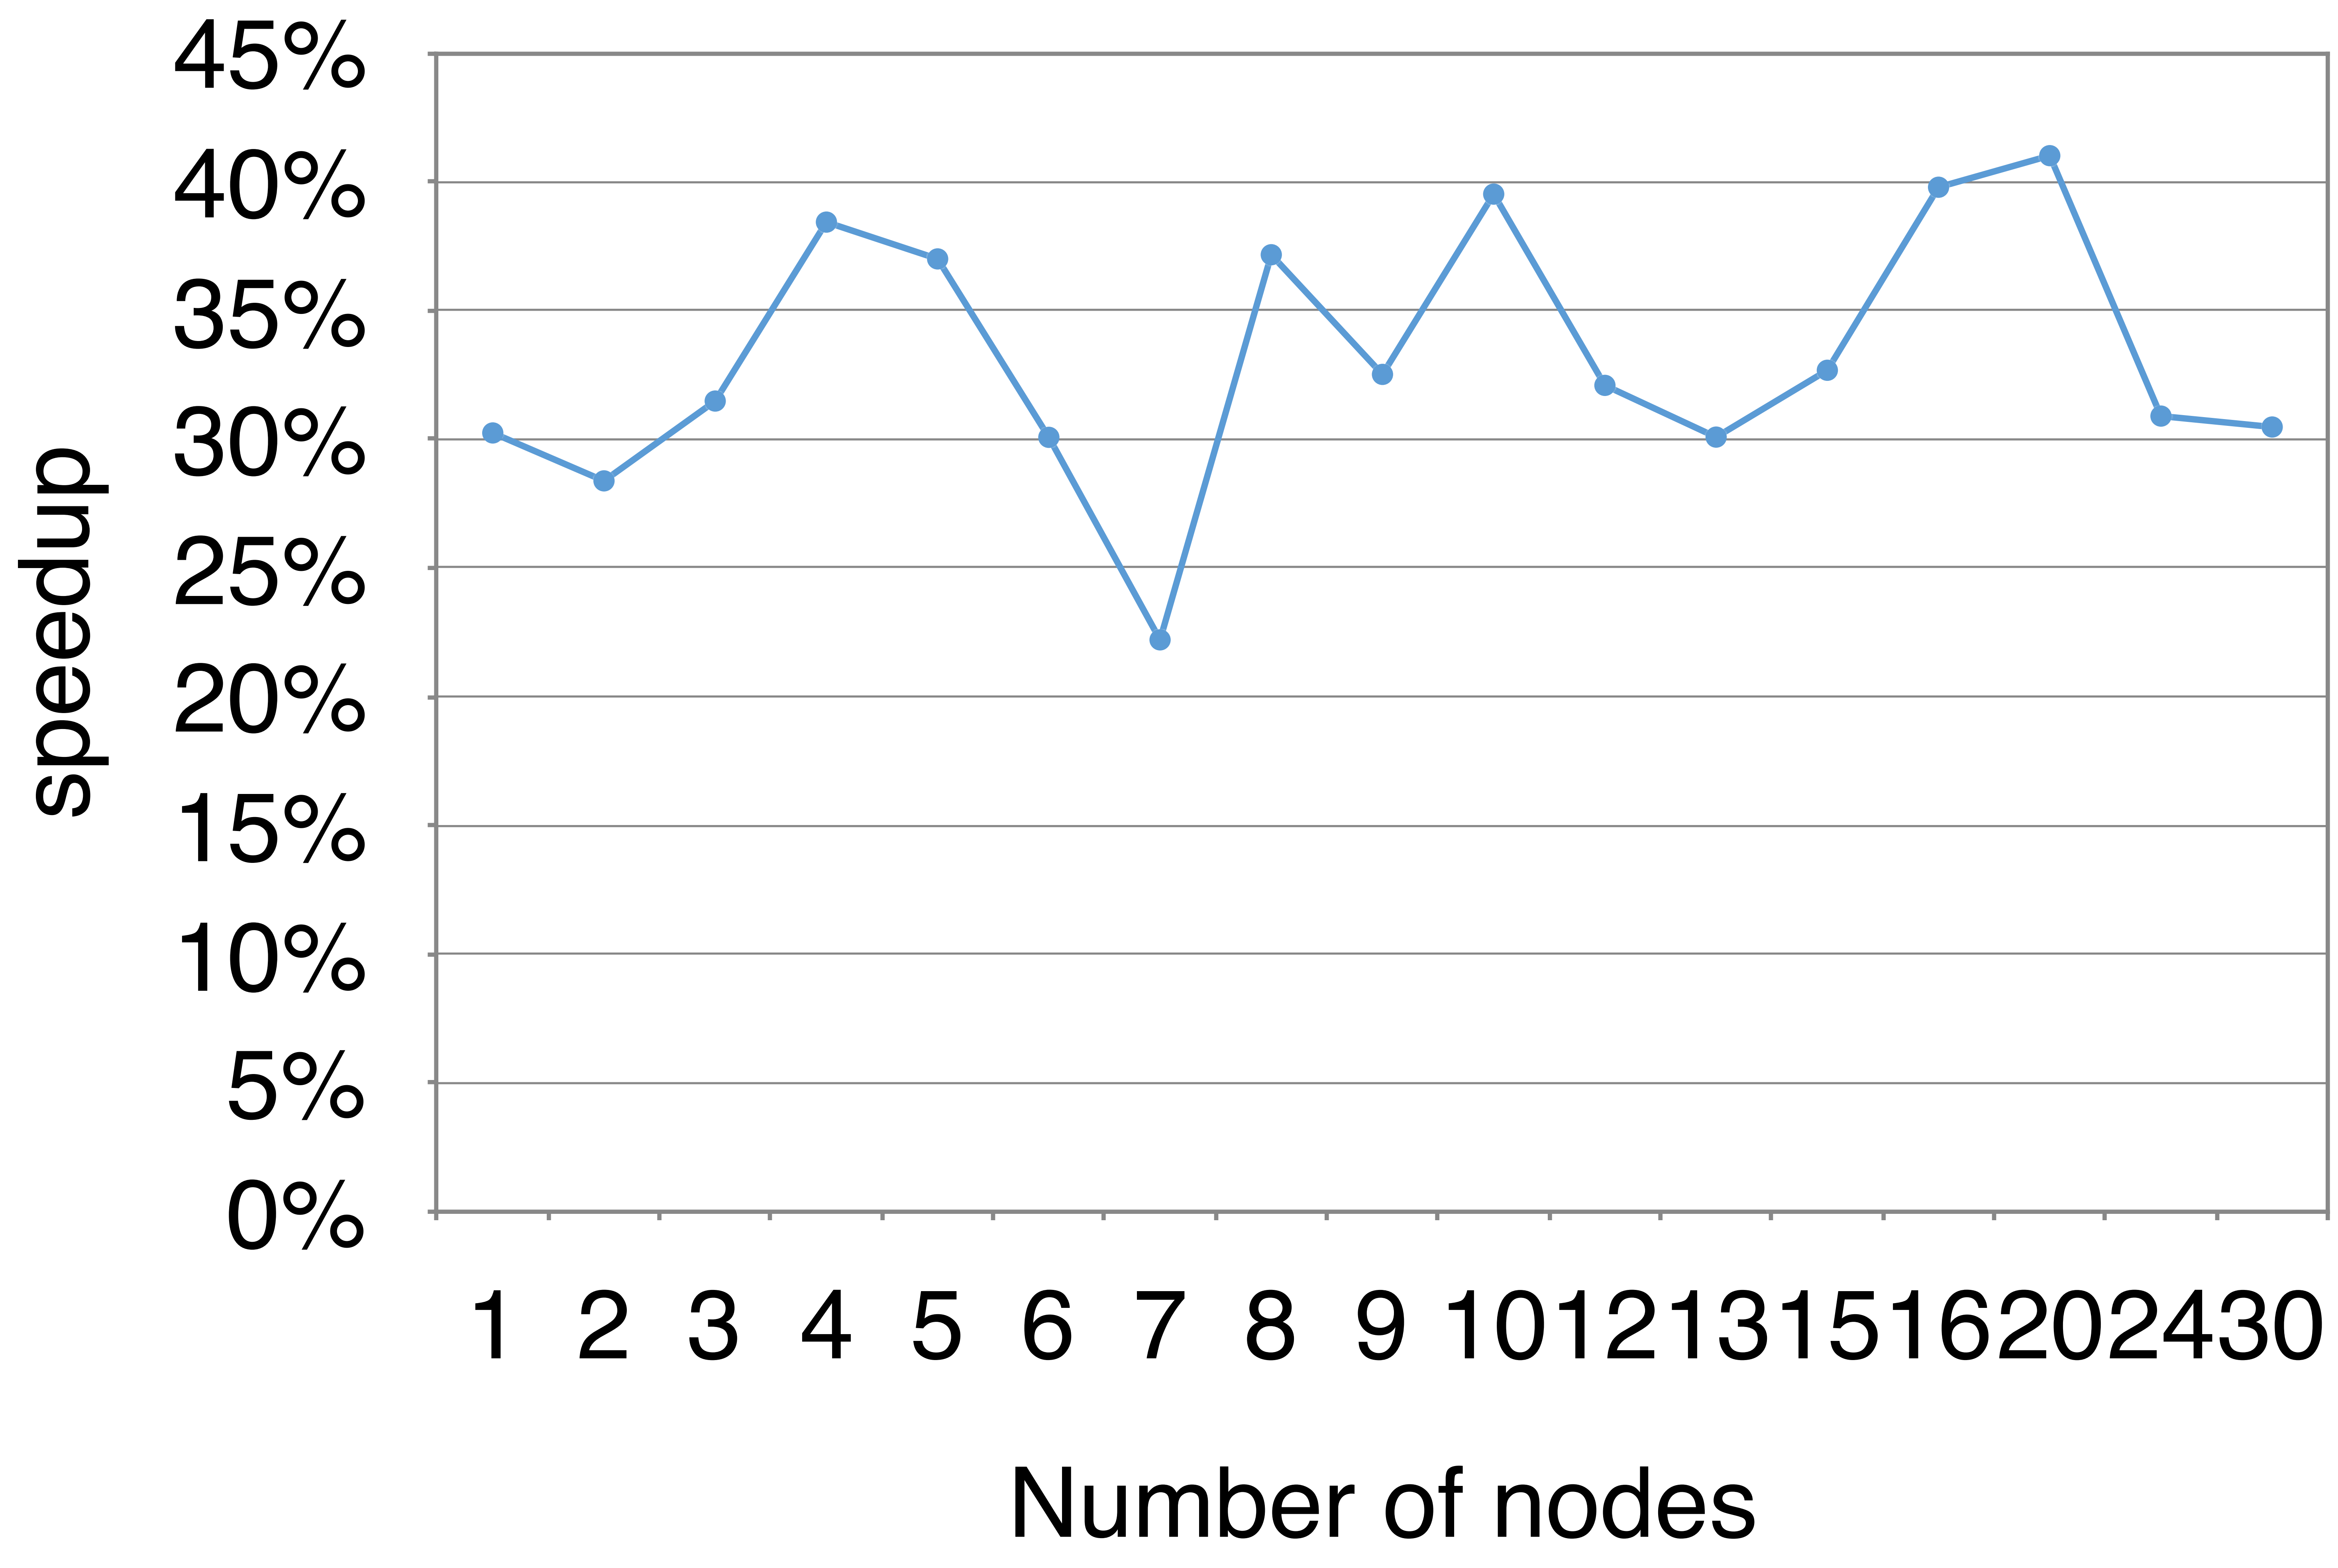

Funded by the European Union

Optimisation of radiation dwarf ACRANEB2

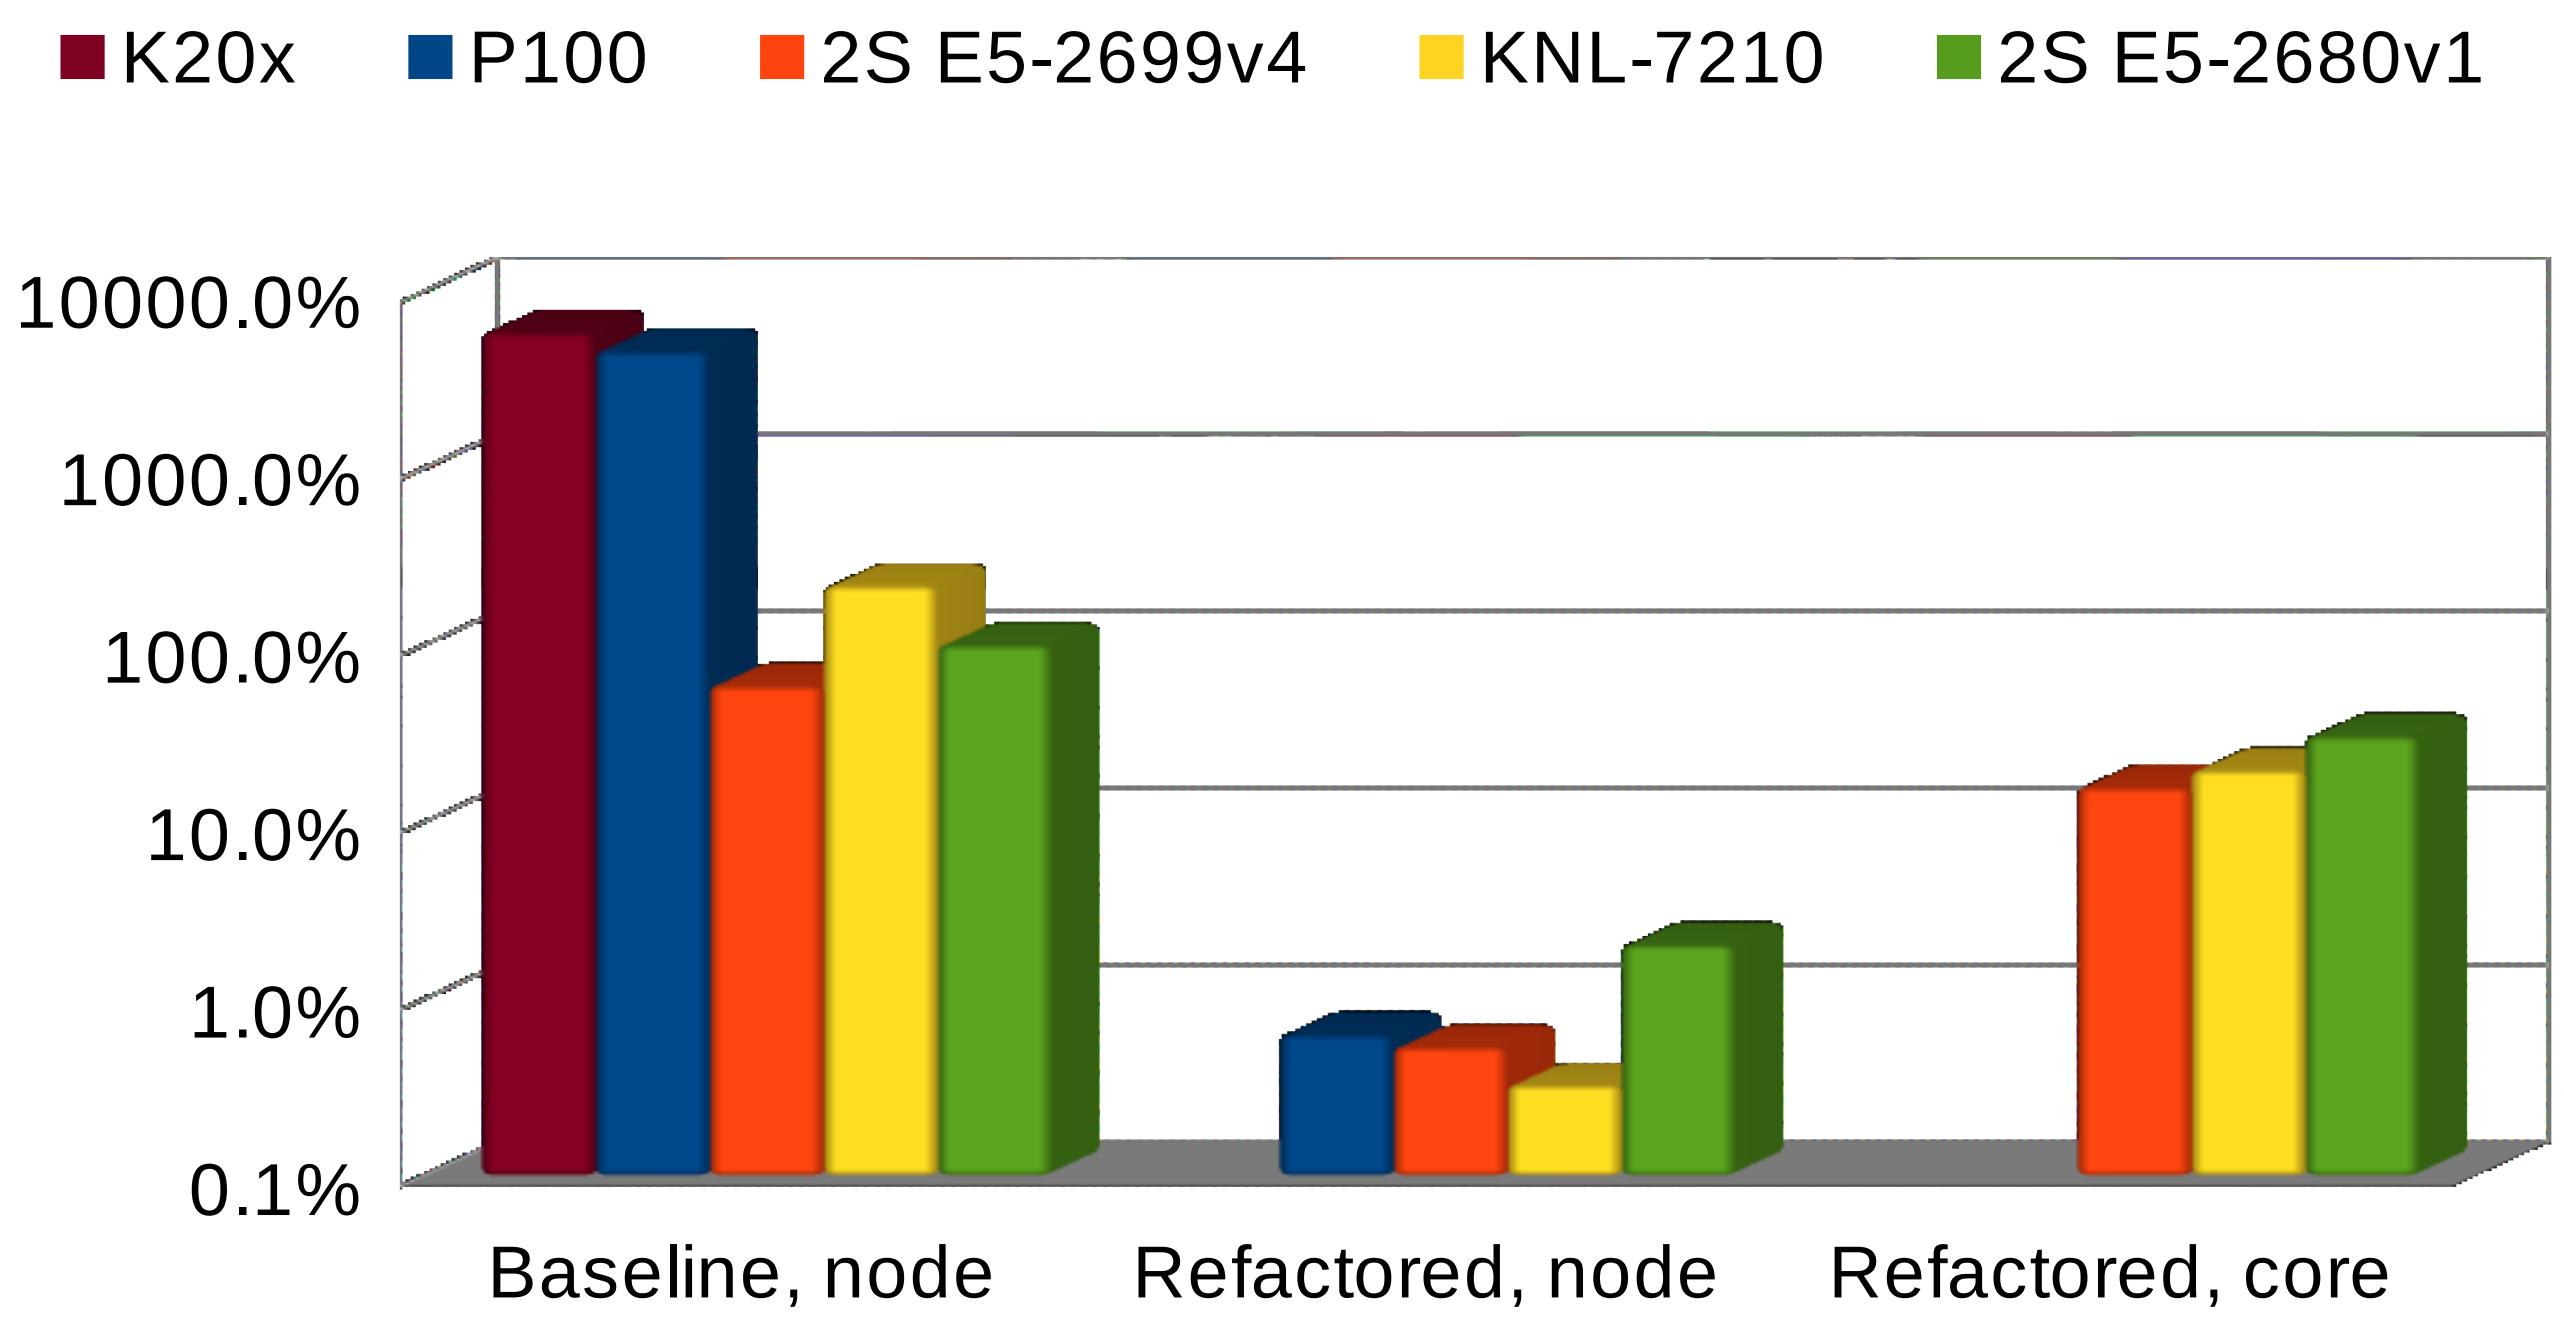

baseline: 1 Intel Sandy Bridge node (green)

Poulsen and Berg (2017)

Optimisation of radiation dwarf ACRANEB2

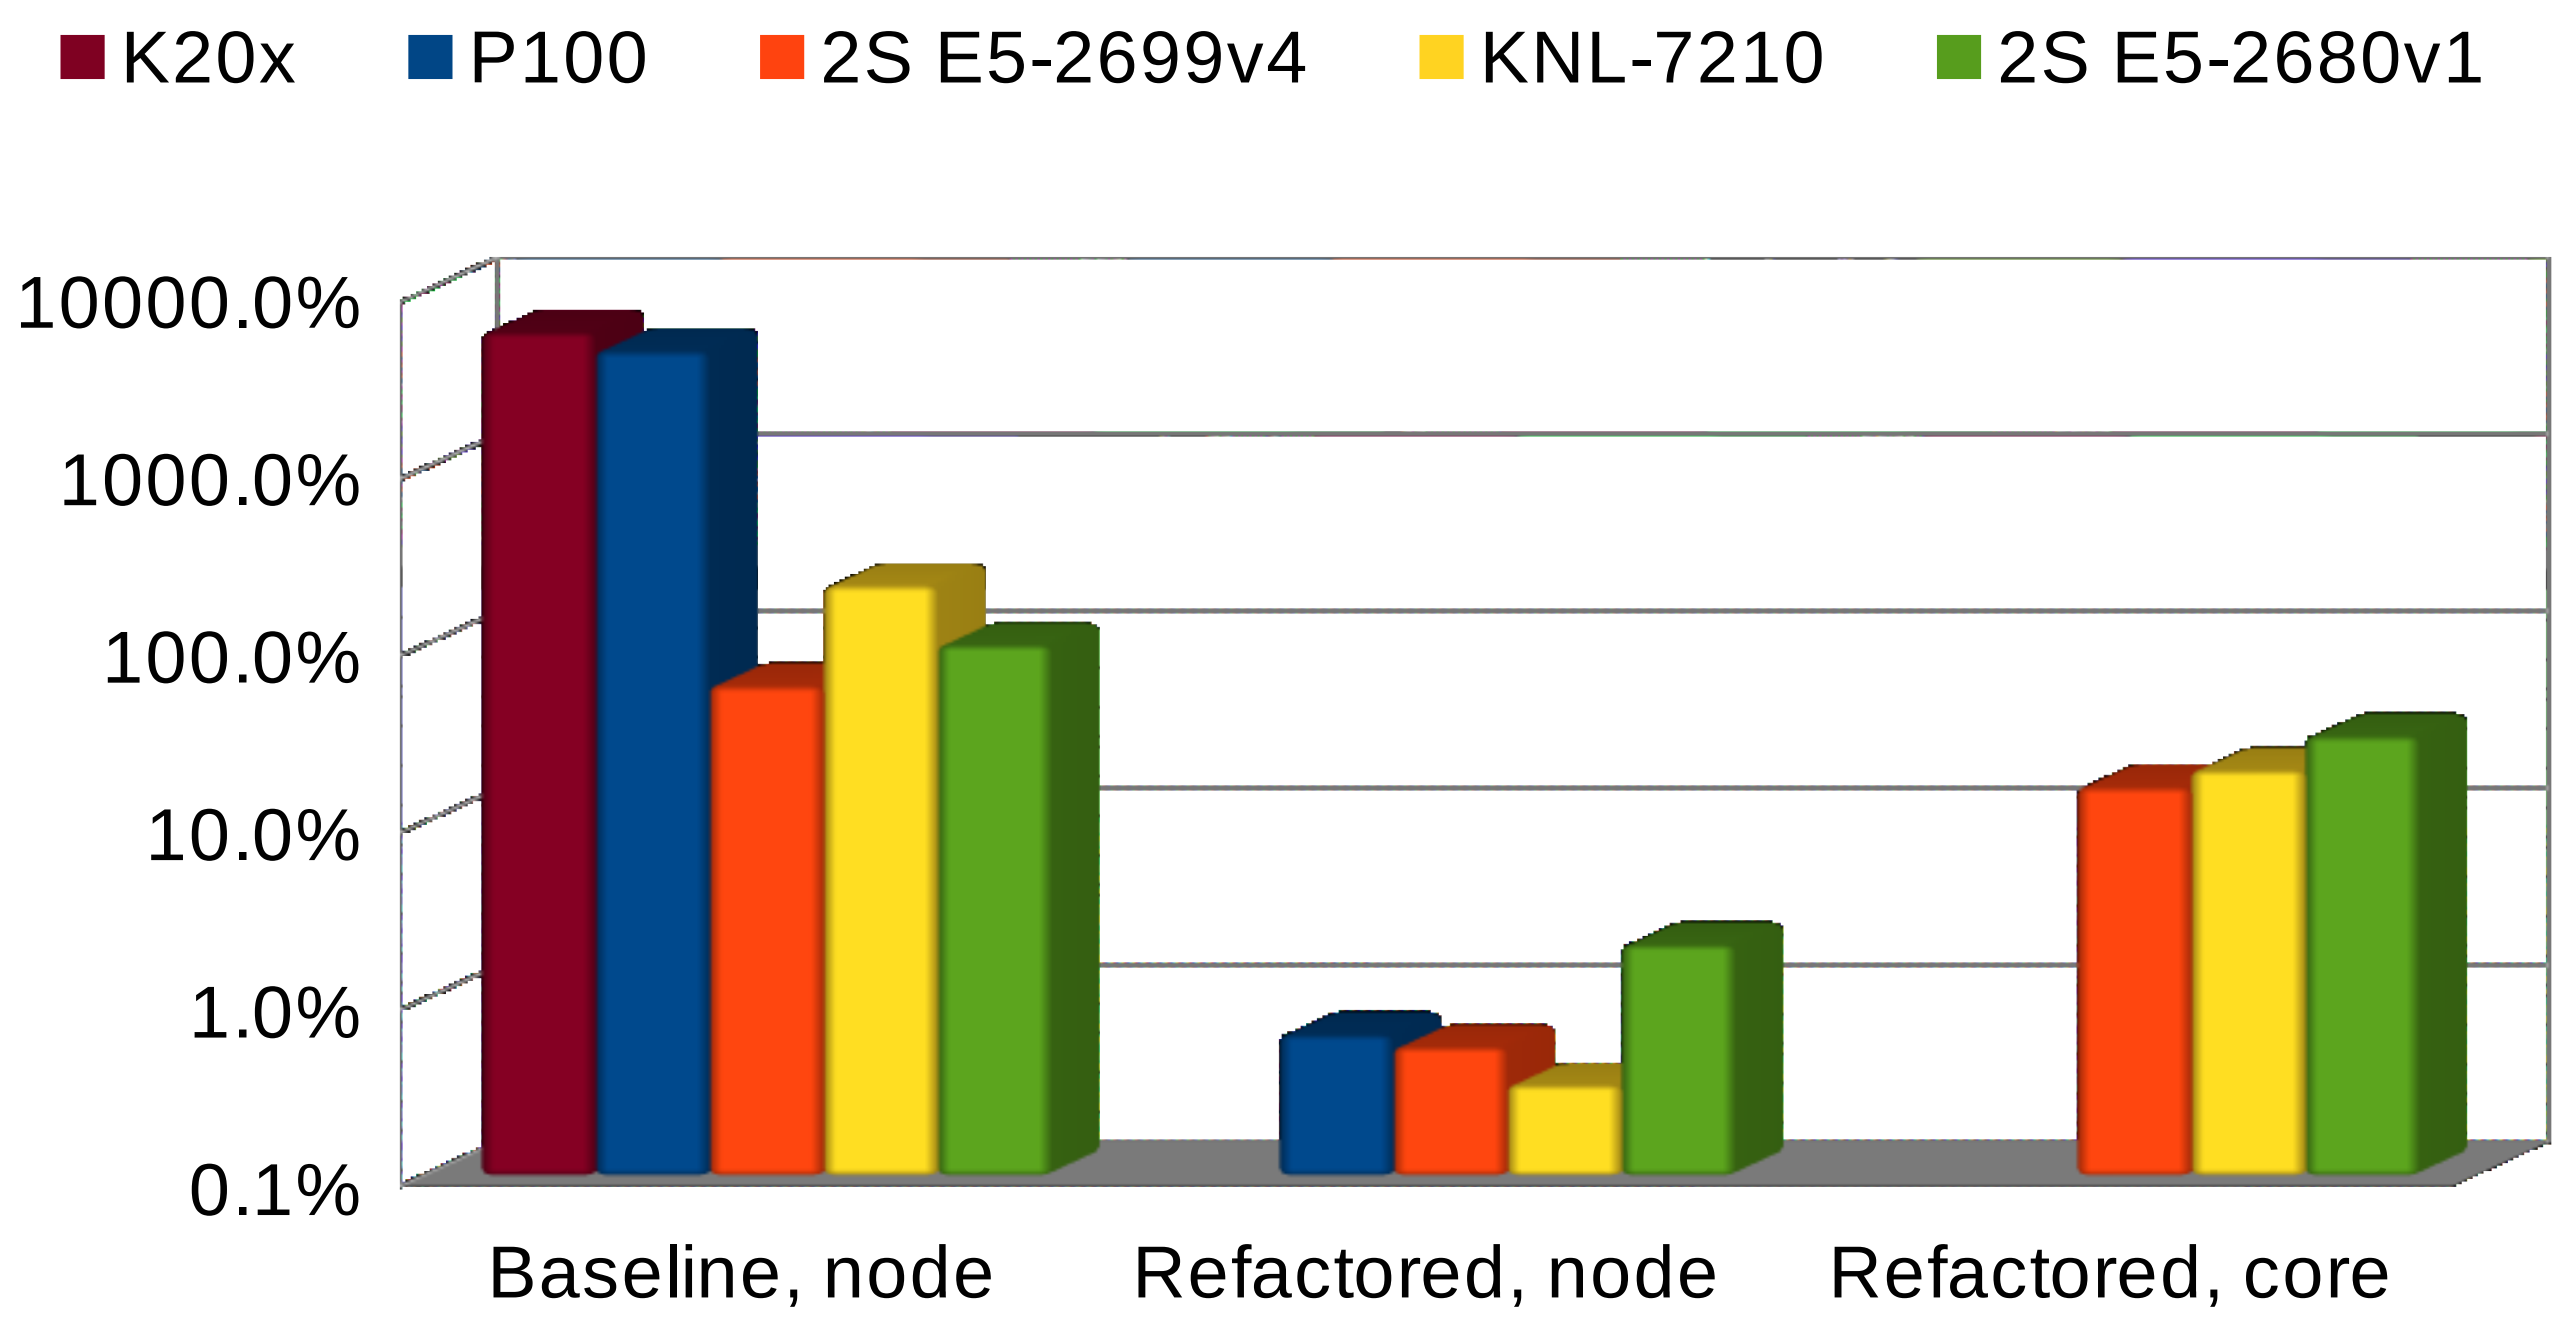

baseline: 1 Intel Sandy Bridge node (green)

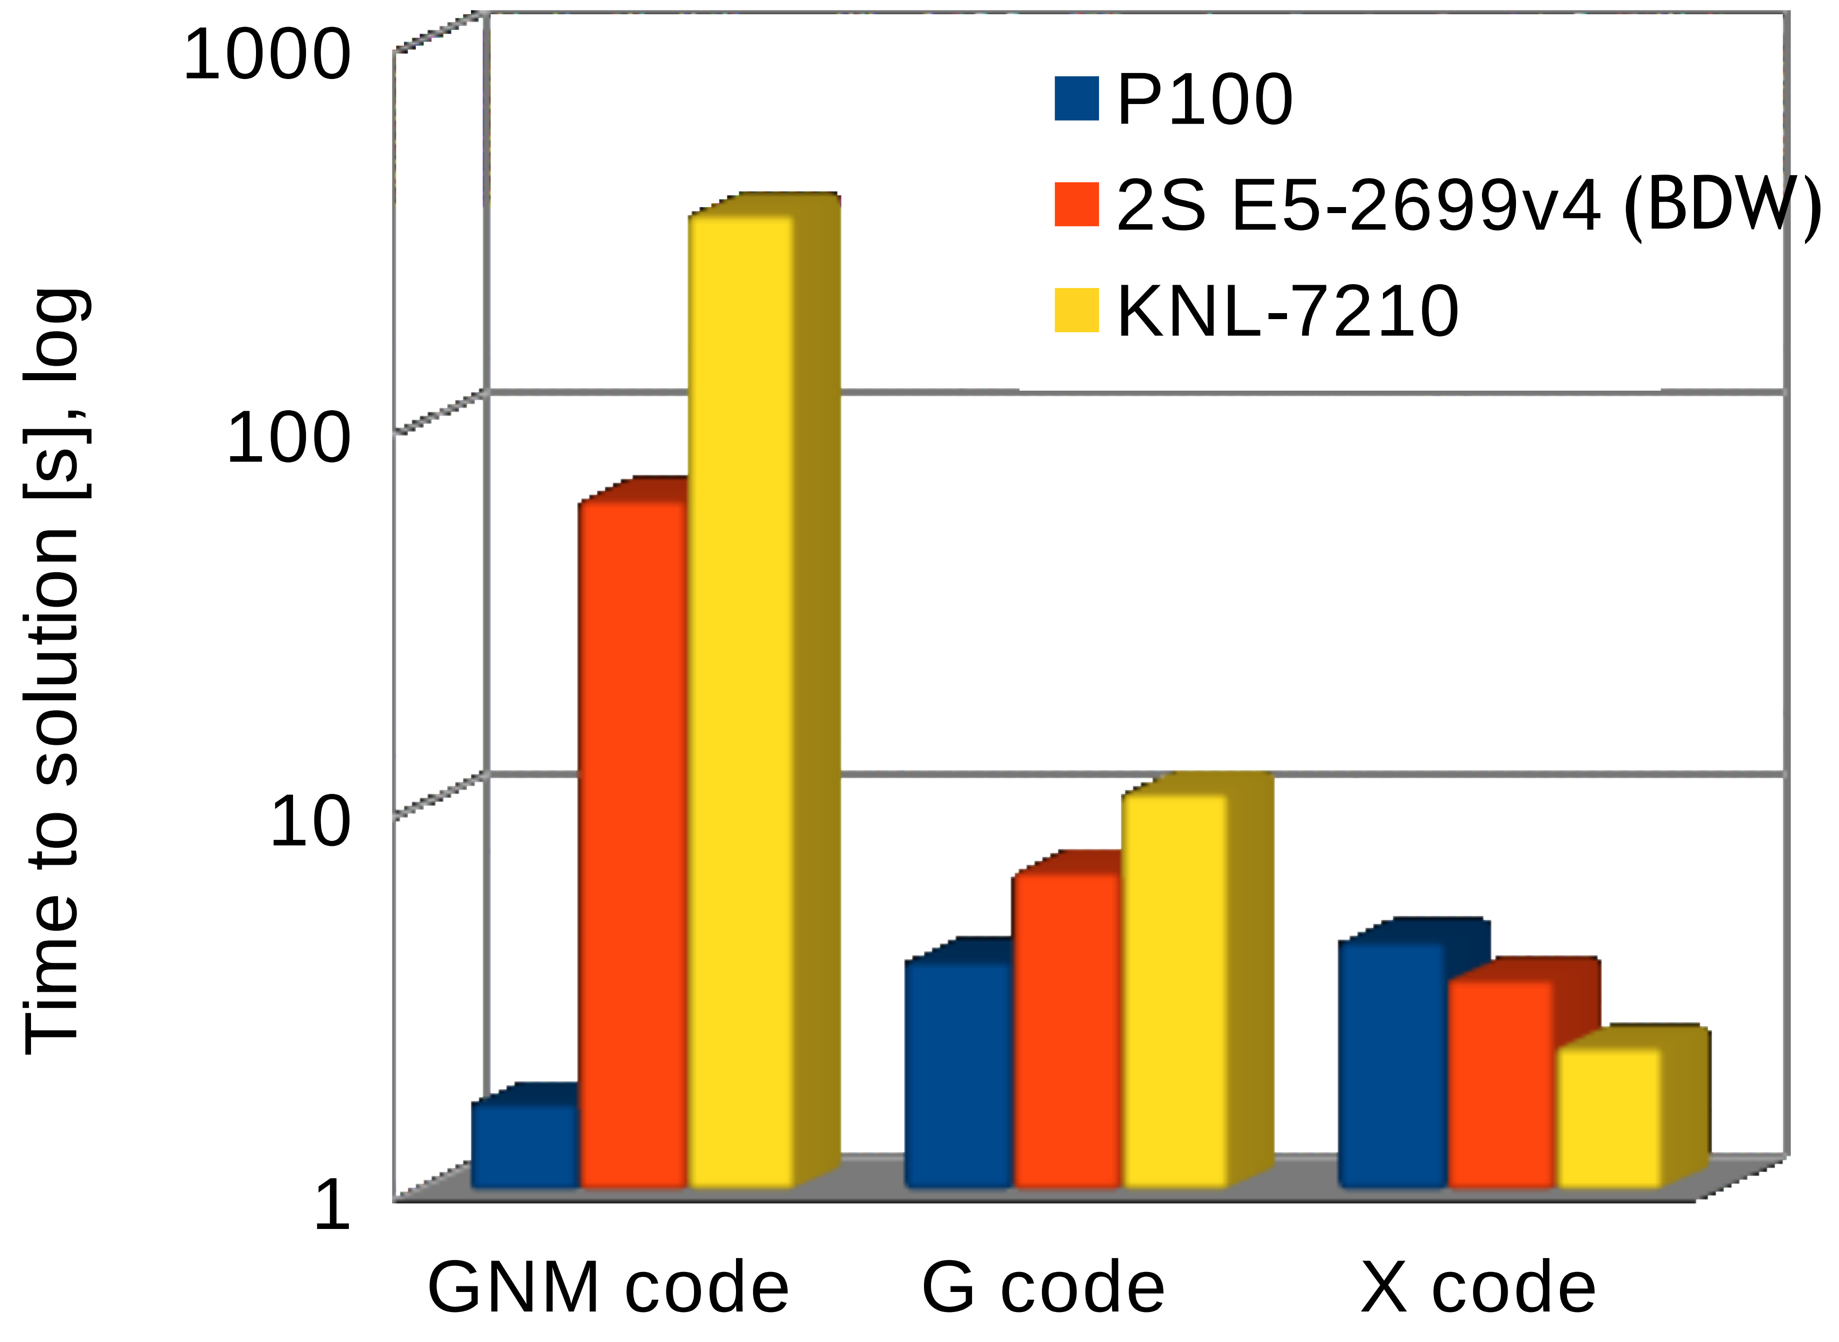

X: optimised for CPU
G: optimised for GPU, data structures from X
GNM: GPU with transposed data structures

Poulsen and Berg (2017)

Funded by the European Union

MPDATA: DSL version with Gridtools

```fortran
subroutine compute_upwind_flux(this,pflux,pD,pVn)
type(MPDATA_type), intent(inout) :: this
real(wp), intent(out) :: pflux(:,:)
real(wp), intent(in)  :: pVn(:,:), pD(:,:)
real(wp) :: zpos, zneg
integer :: nb_edges
integer :: nb_levels
integer  :: jedge, jlev, ip1, ip2

call atlas_log%debug('compute_upwind_flux')

nb_edges  = this%dimensions%nb_edges
nb_levels = this%dimensions%nb_levels

!$OMP PARALLEL DO SCHEDULE(STATIC) PRIVATE(jedge,jlev,ip1,ip2,zpos,zneg)
do jedge = 1,nb_edges
  ip1 = iedge2node(1,jedge)
  ip2 = iedge2node(2,jedge)
  do jlev = 1,nb_levels
    zpos                = max(0._wp,pVn(jlev,jedge))
    zneg                = min(0._wp,pVn(jlev,jedge))
    pflux(jlev,jedge)   = pD(jlev,ip1)*zpos+pD(jlev,ip2)*zneg
  enddo
enddo
!$OMP END PARALLEL DO
end subroutine compute_upwind_flux
```

DSL version:
- single code base, readable
- performance portable
- domain knowledge adds flexibility

```cpp
template <uint_t Color> struct upwind_flux {
  using flux = accessor<0, enumtype::inout, icosahedral_topology_t::edges>;
  using pD =
      in_accessor<1, icosahedral_topology_t::vertices, extent<0, 1, 0, 1>>;
  using vn = in_accessor<2, icosahedral_topology_t::edges, extent<0, 1, 0, 1>>;

  typedef boost::mpl::vector<flux, pD, vn> arg_list;

  template <typename Evaluation> static void Do(Evaluation &eval, k_full) {

    constexpr auto neighbors_offsets =
        connectivity<edges, vertices, Color>::offsets();
    constexpr auto ip0 = neighbors_offsets[0];
    constexpr auto ip1 = neighbors_offsets[1];

    float_type pos = math::max(eval(vn()), (float_type)0.);
    float_type neg = math::min(eval(vn()), (float_type)0.);

    eval(flux()) = eval(pos * pD(ip0) + neg * pD(ip1));
  }
};
```

Complementary skills of GridTools (MeteoSwiss) and Atlas (ECMWF)

bringing optimisations into operational workflow

postprocessing of spectral data at 9km resolution

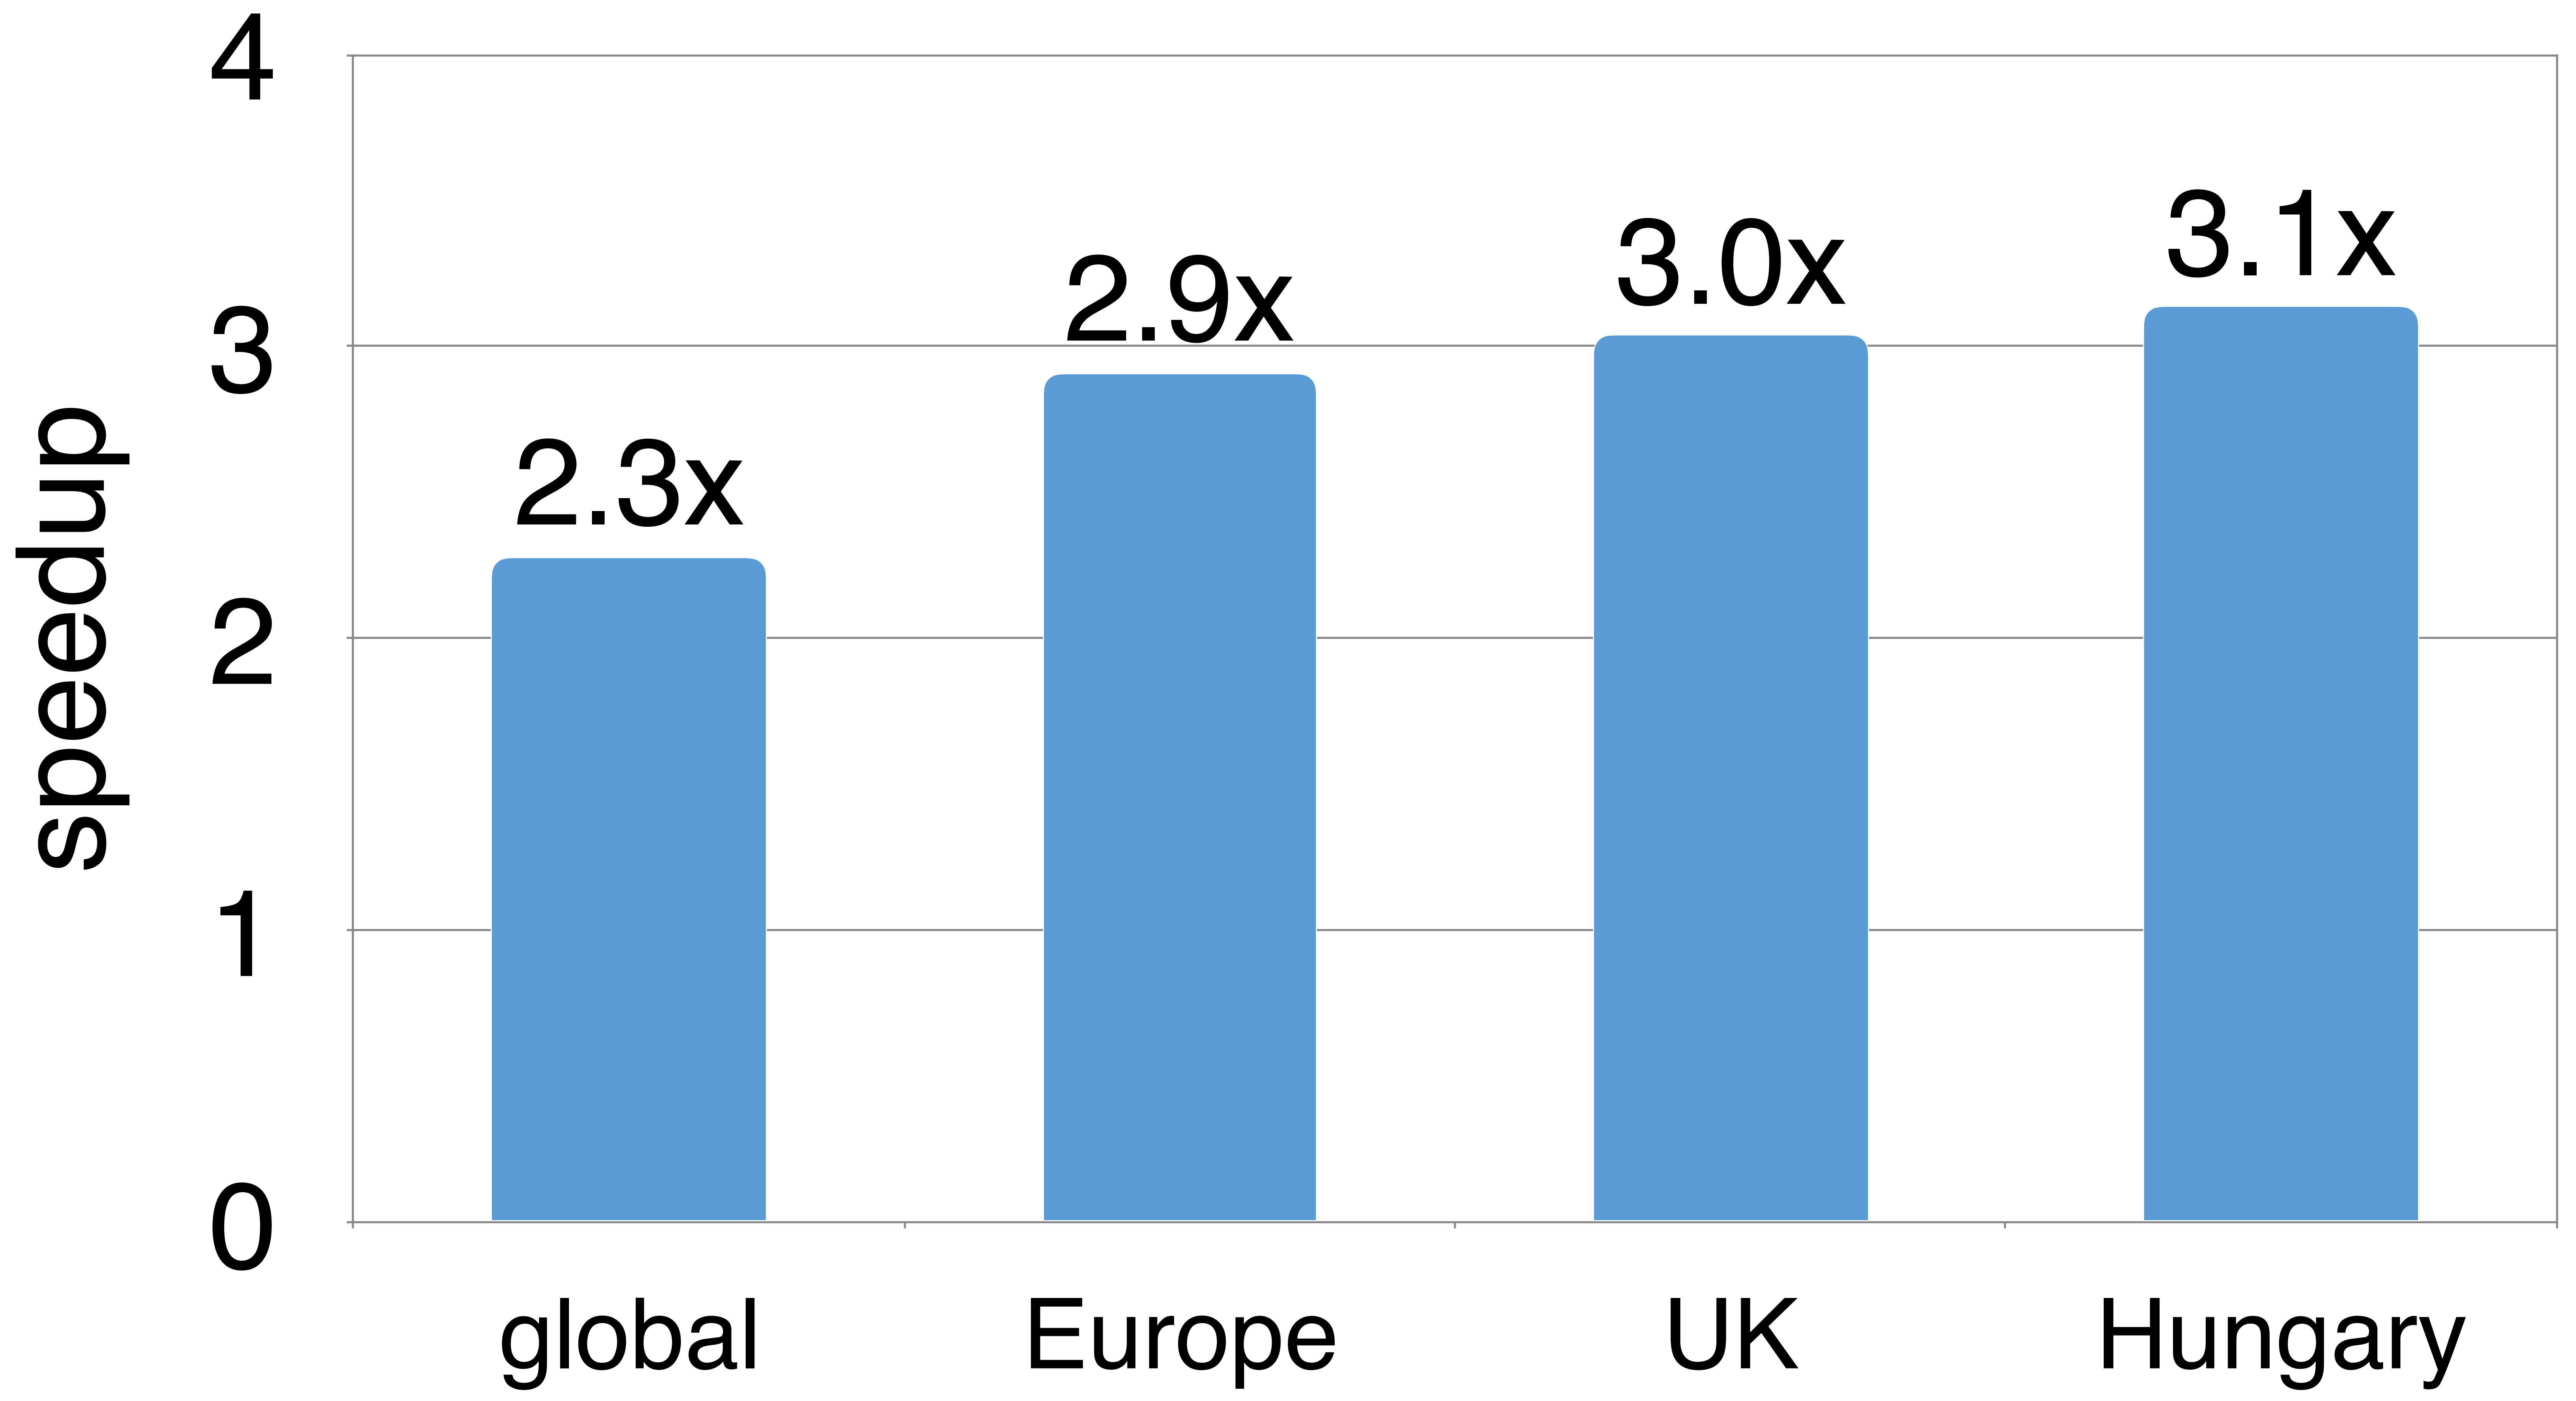

speedup compared to current operational transform used for postprocessing

Energy vs Time-to-solution

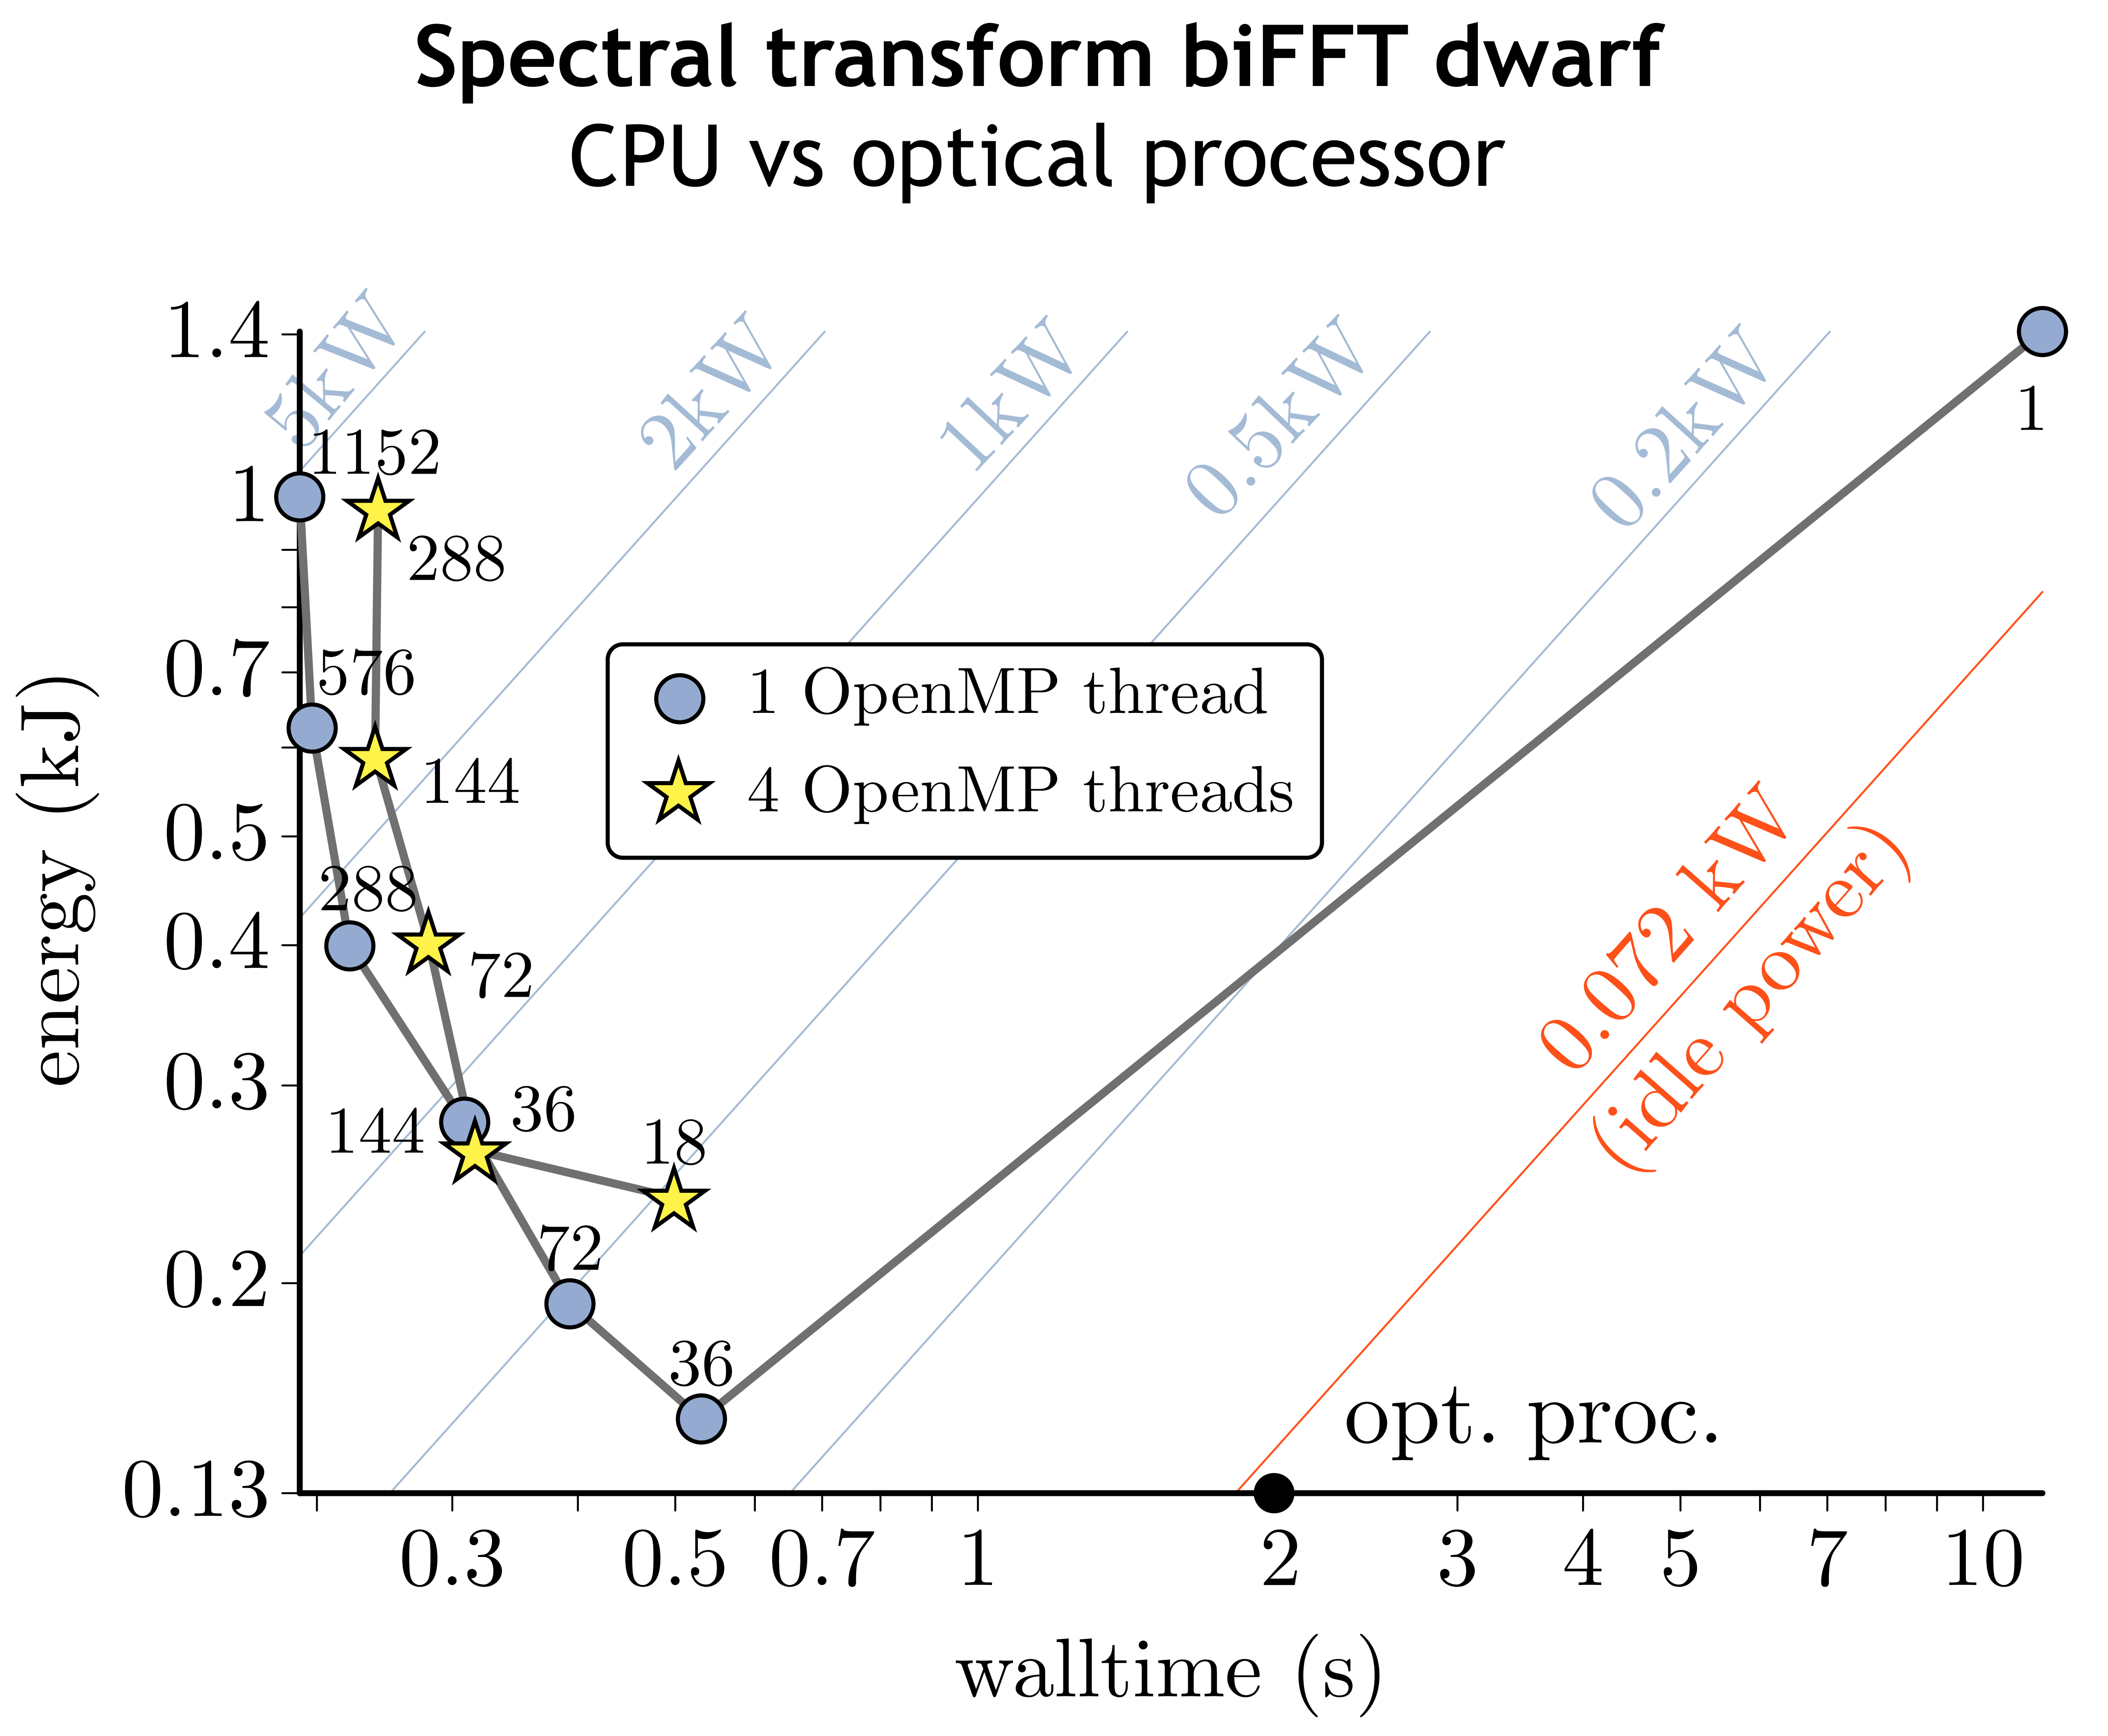

Energy vs Time-to-solution

Spectral transform biFFT dwarf
CPU vs optical processor

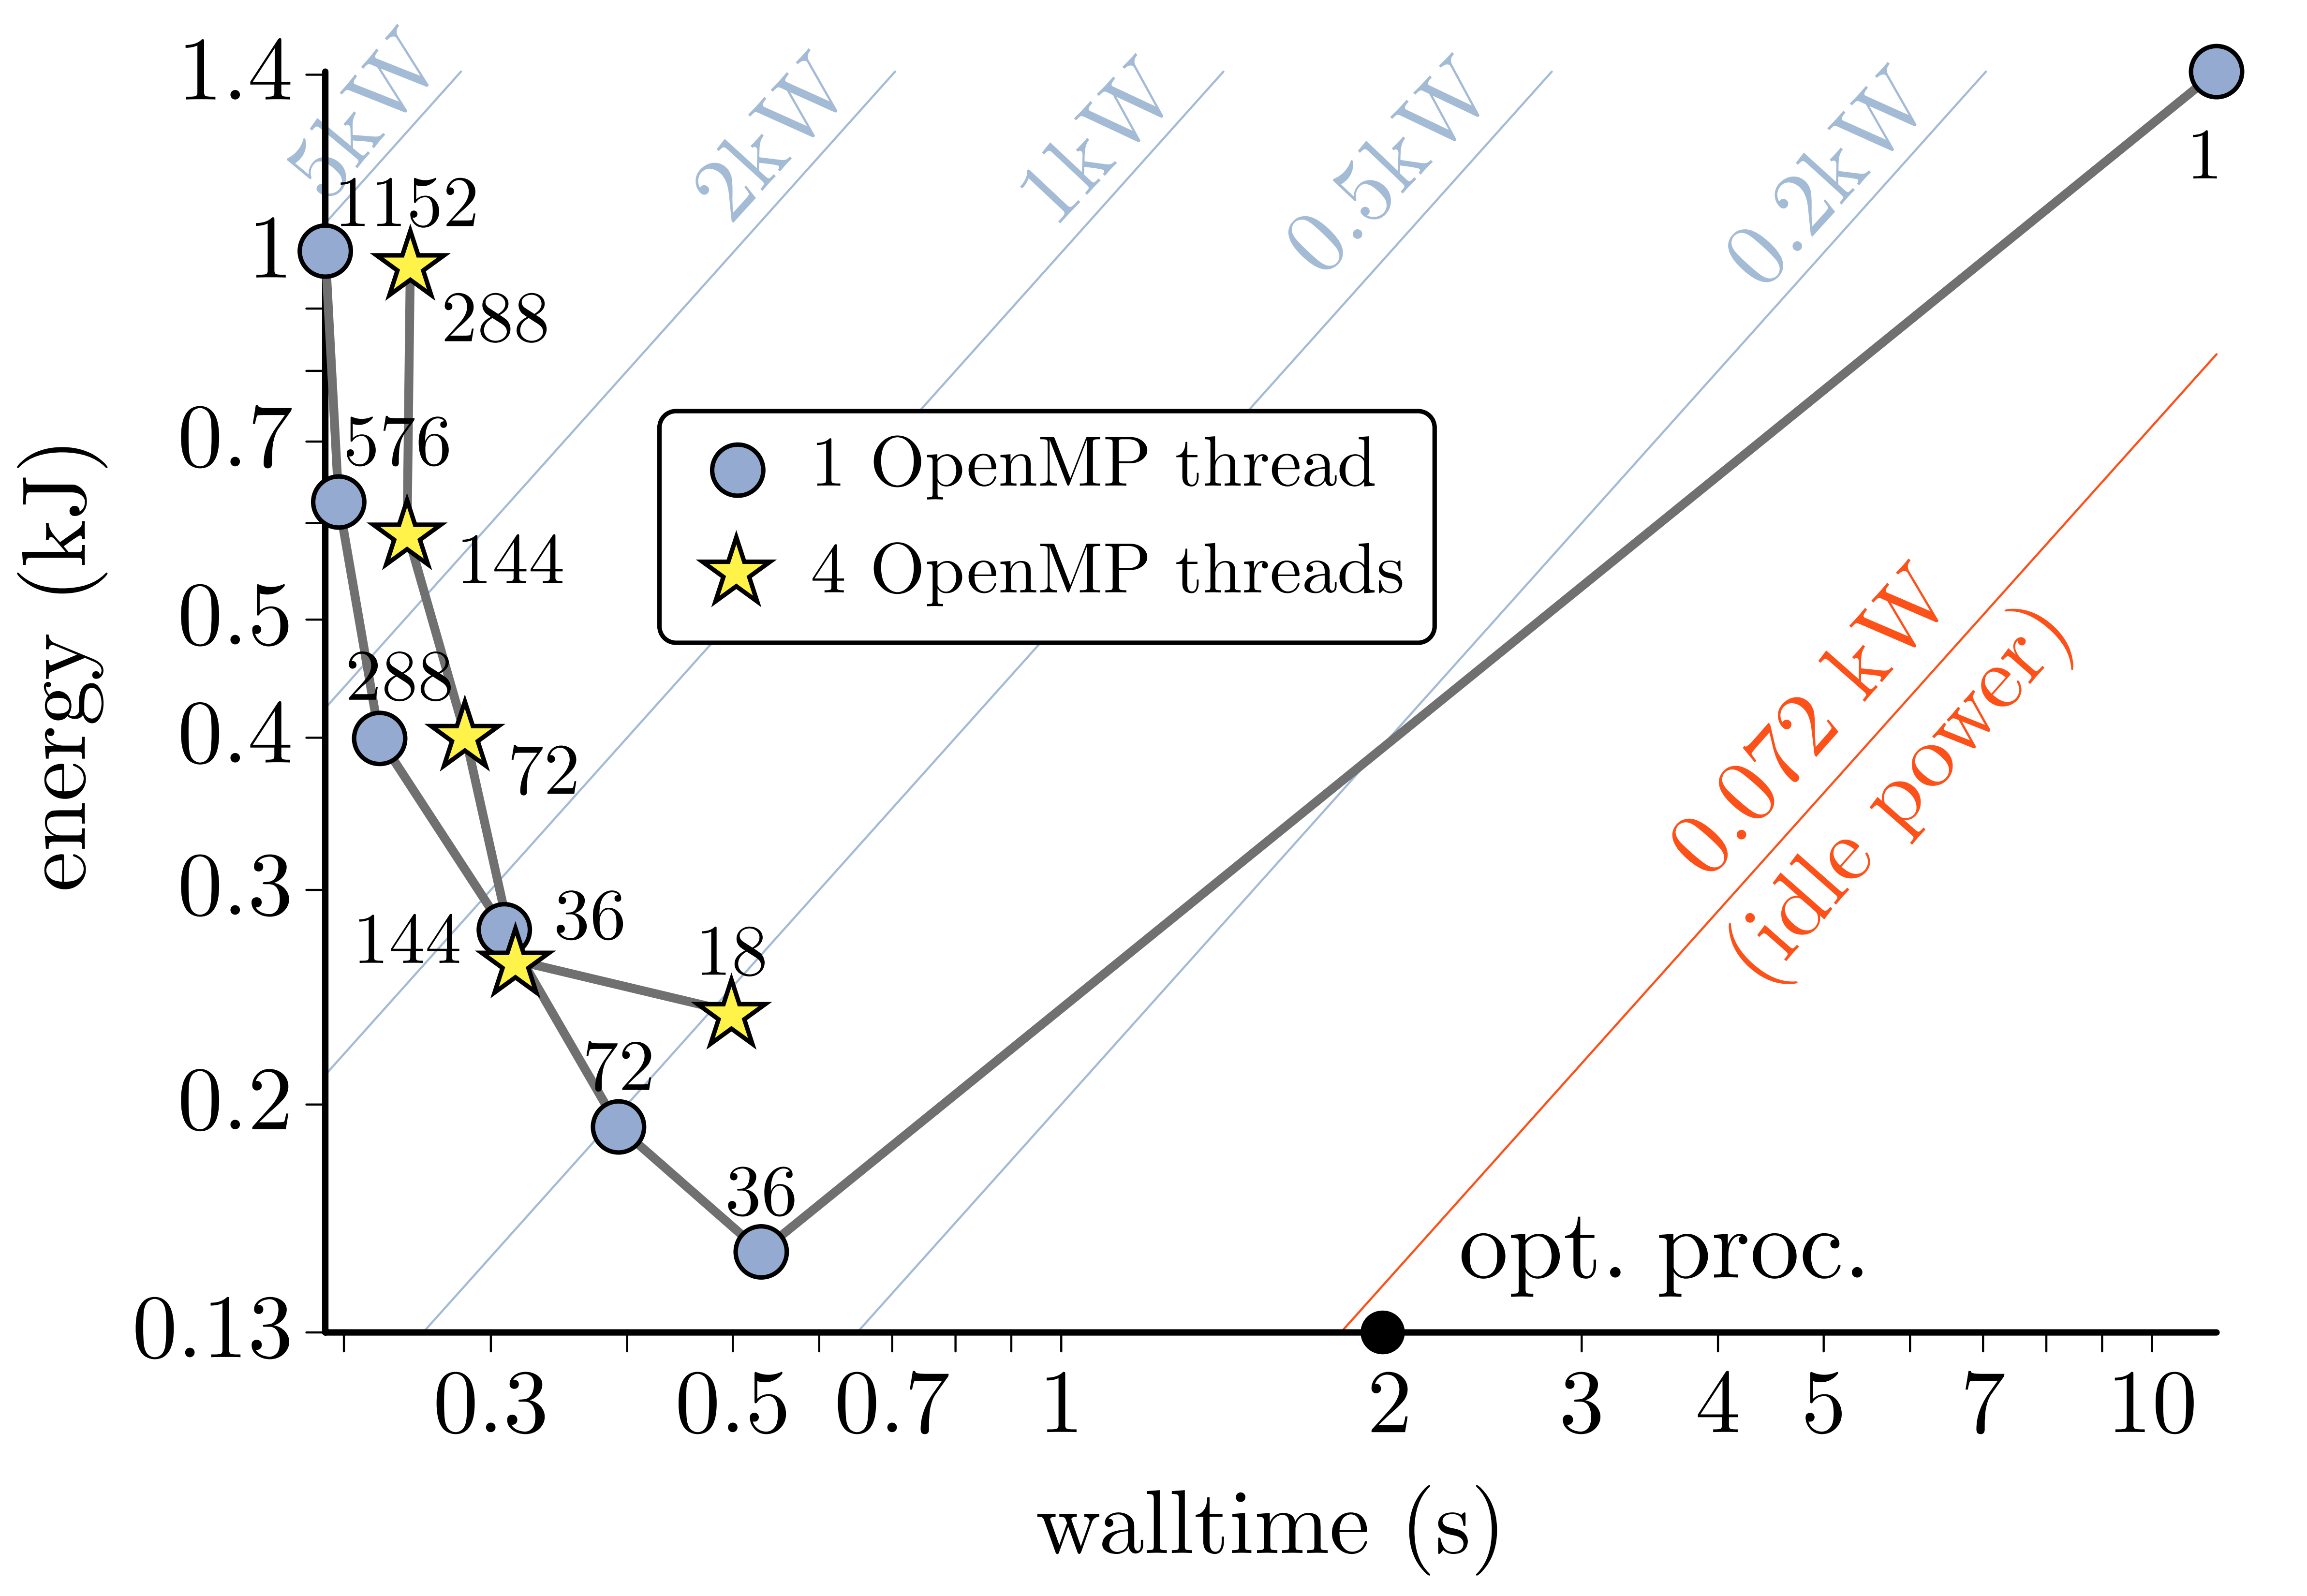

Full forecast model
(ALARO: 2.5km, COSMO-EULAG: 2.2km)

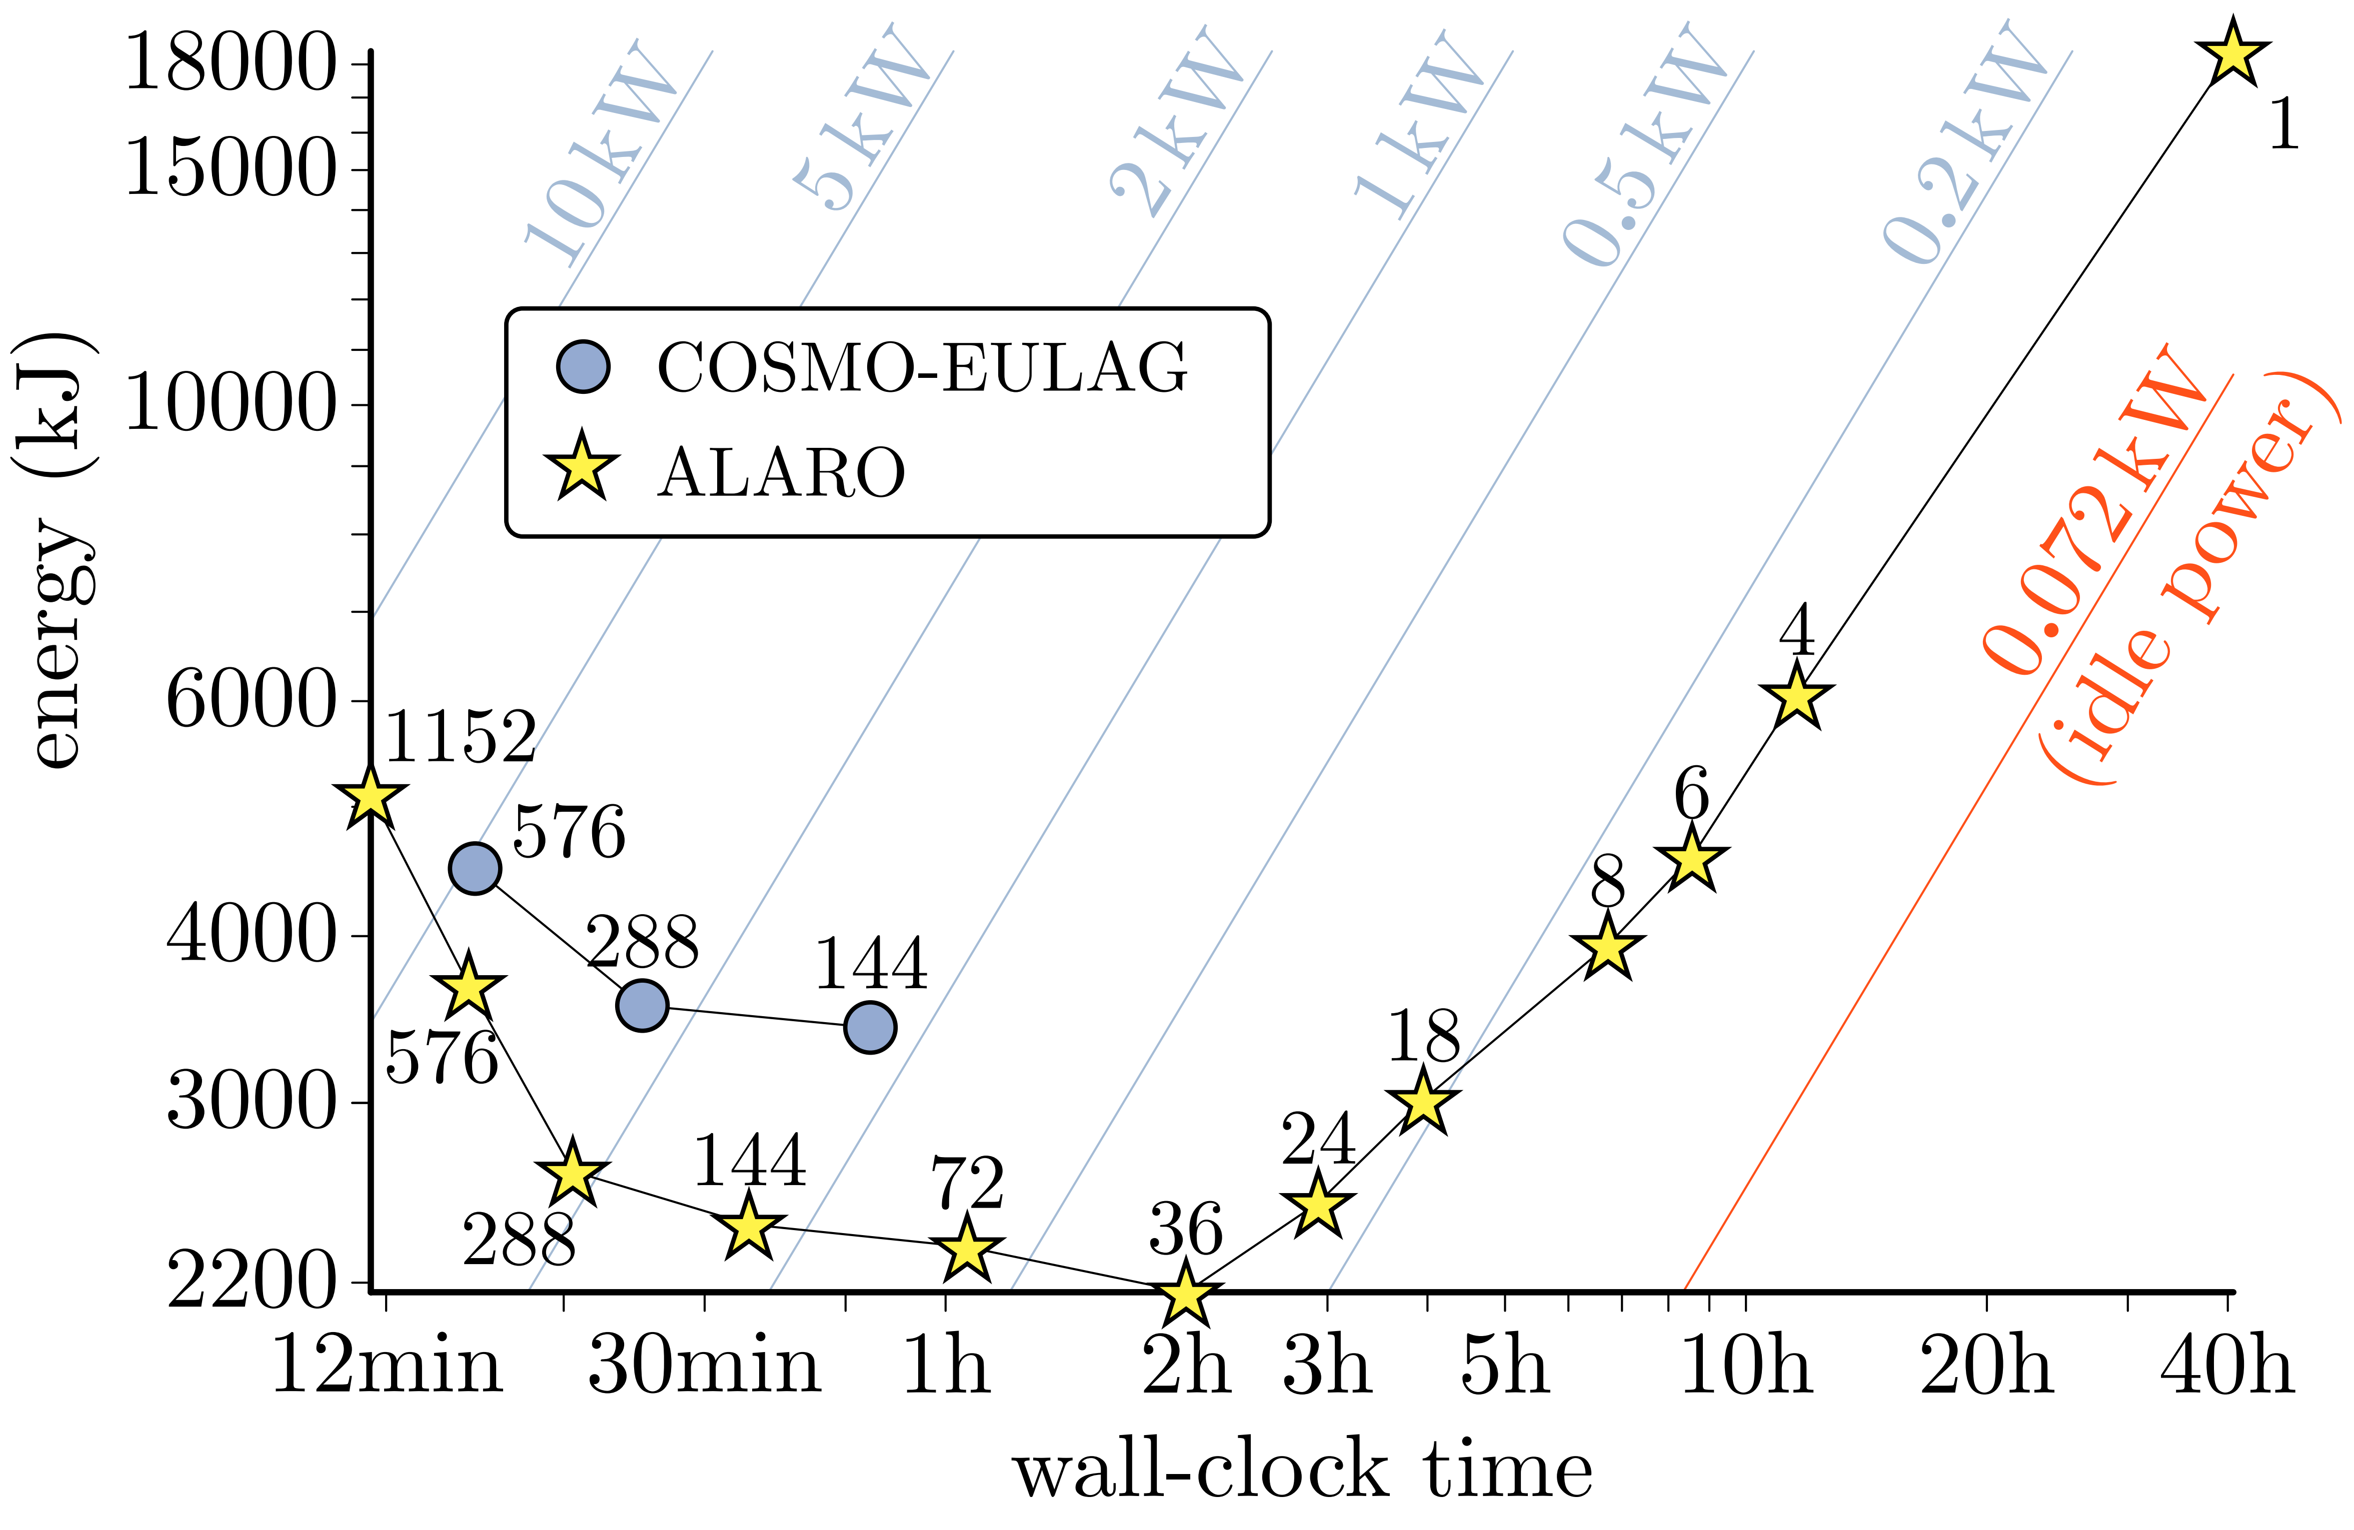

ESCAPE ESCAPE 2

Funded by the European Union

ESCAPE2 algorithms and methods

Mathematical methods and algorithms
- Semi-implicit, semi-Lagrangian CG/DG
- Hierarchical multigrid tools
- Fault resilient solver
- Artificial neural networks

and:
- DSL toolchain
- Ensemble based URANIE
- ● State-of-the-art

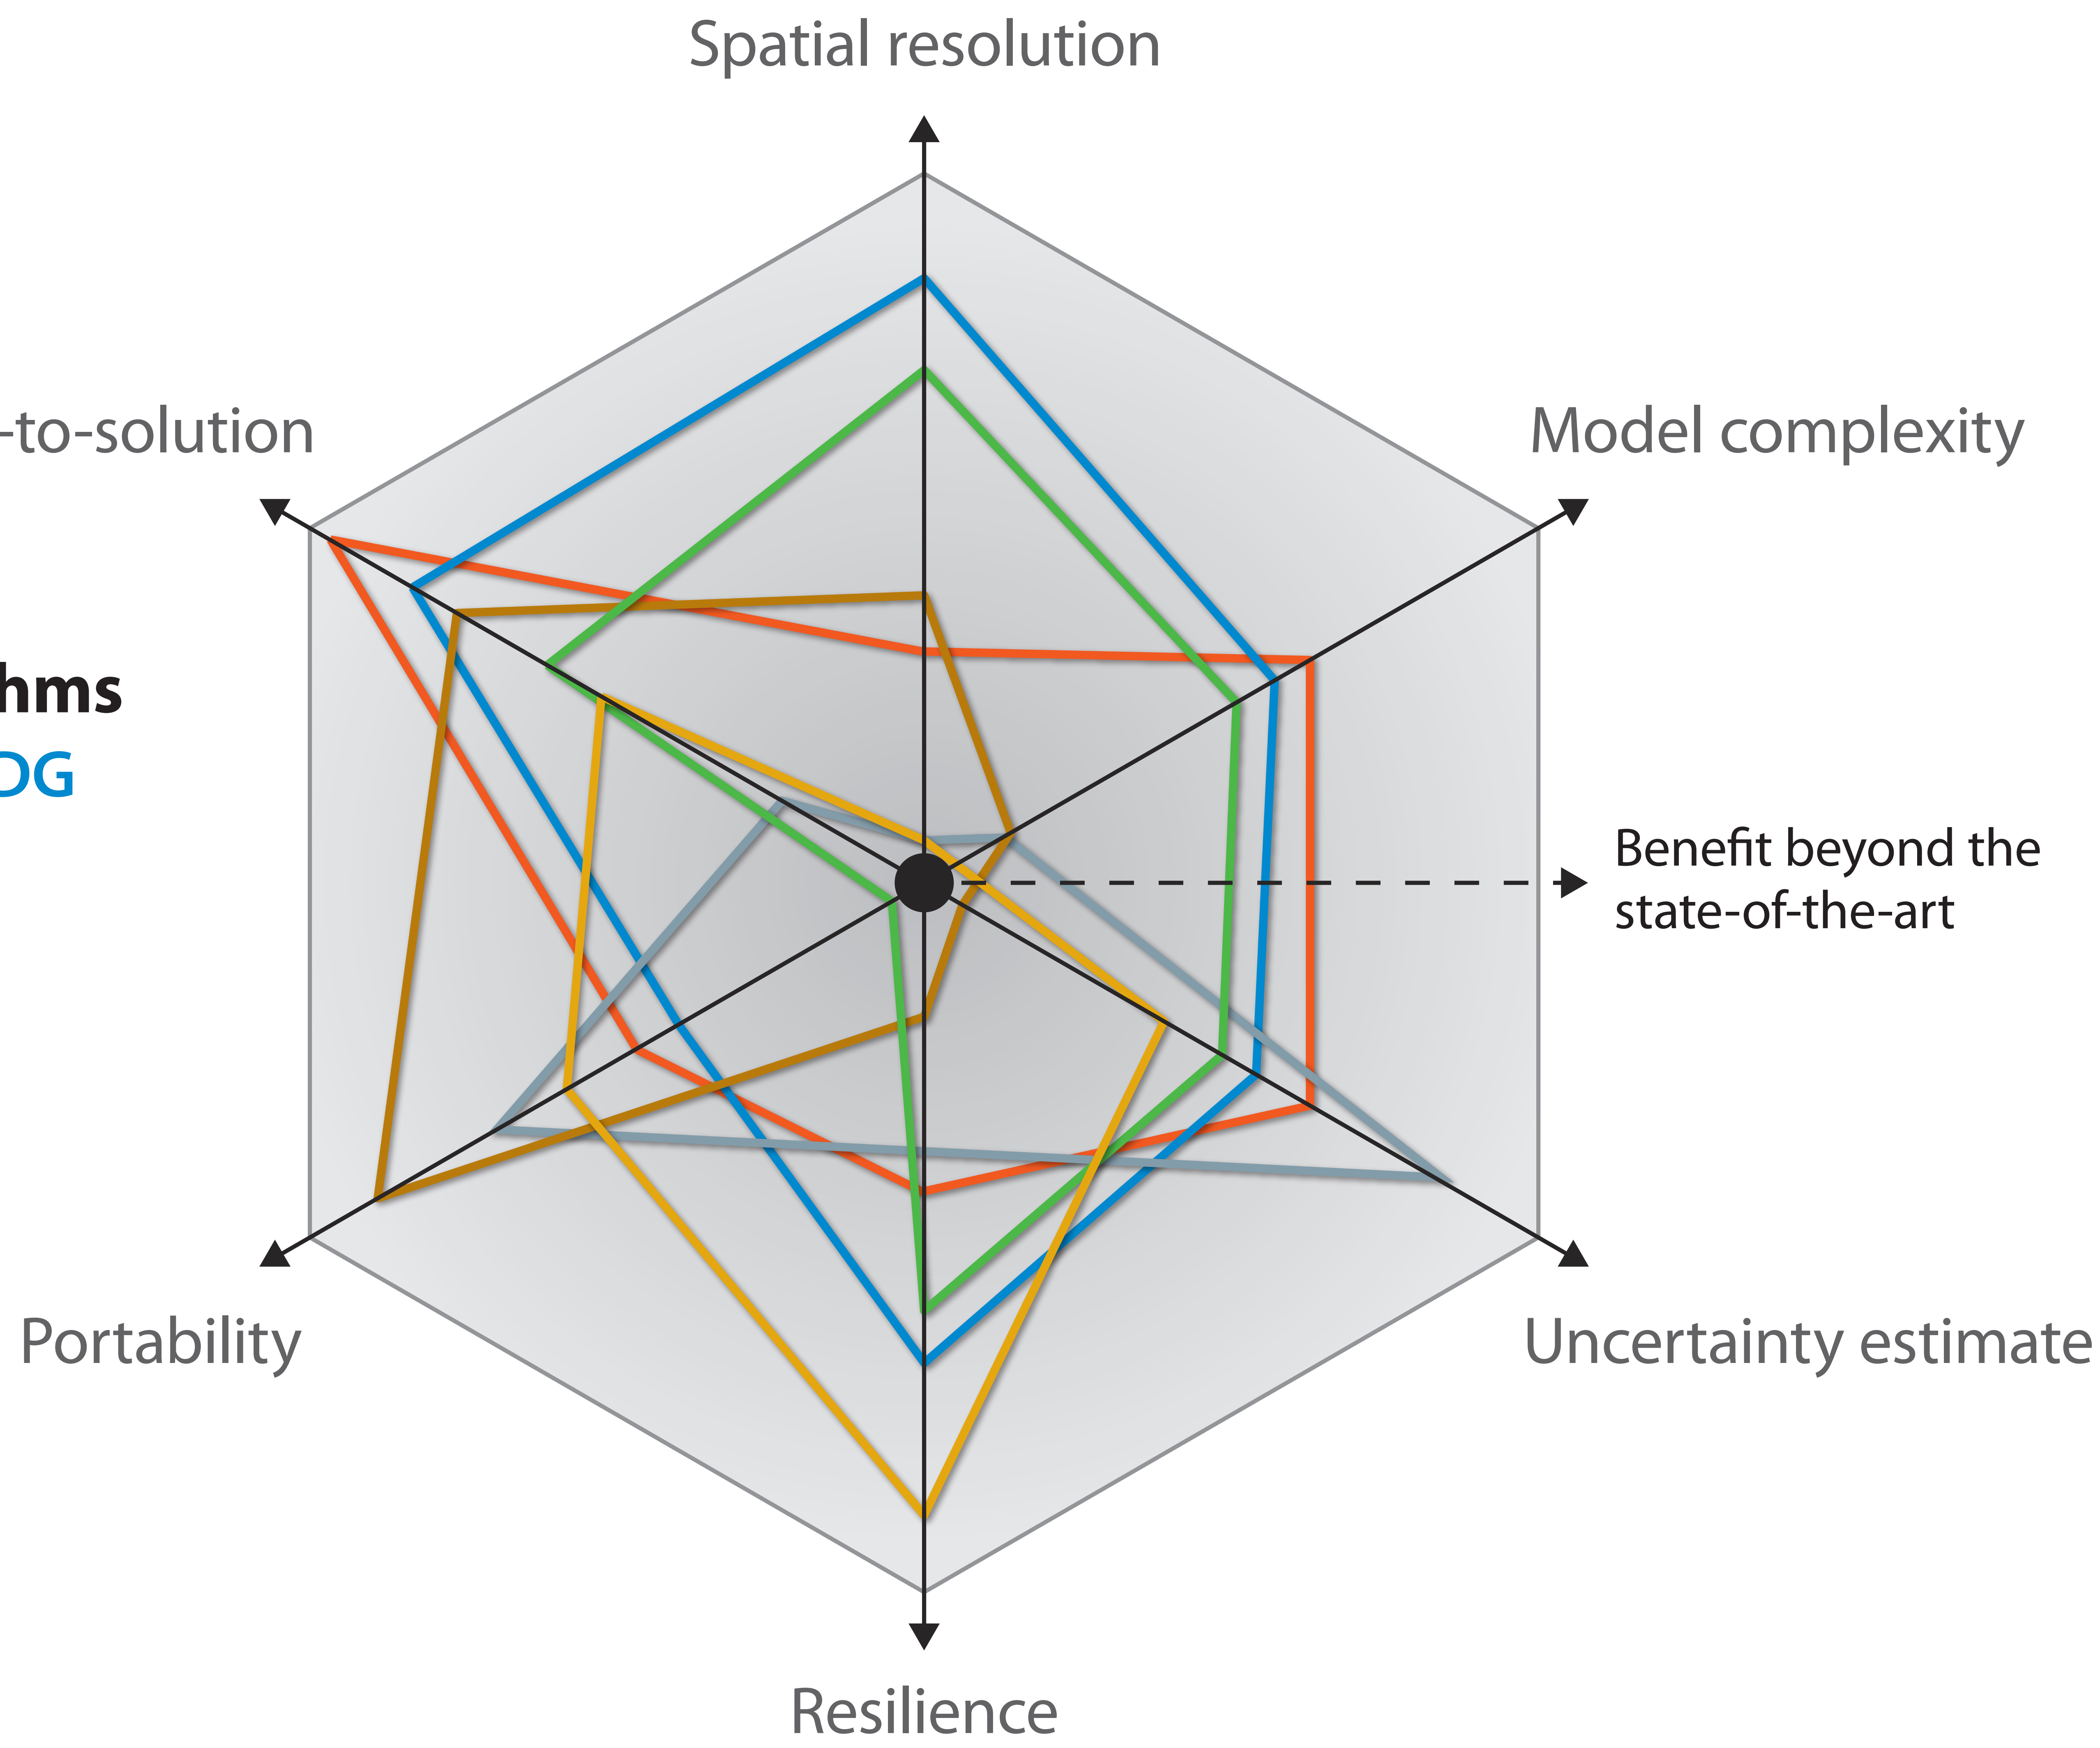

PANTHER: semi-Lagrangian DG with p-adaptivity colliding bubble test case from Robert (1993)

(PANTHER: P-Adaptive Numerical Tool for High order Efficient discRetizations, developed by Giovanni Tumolo)

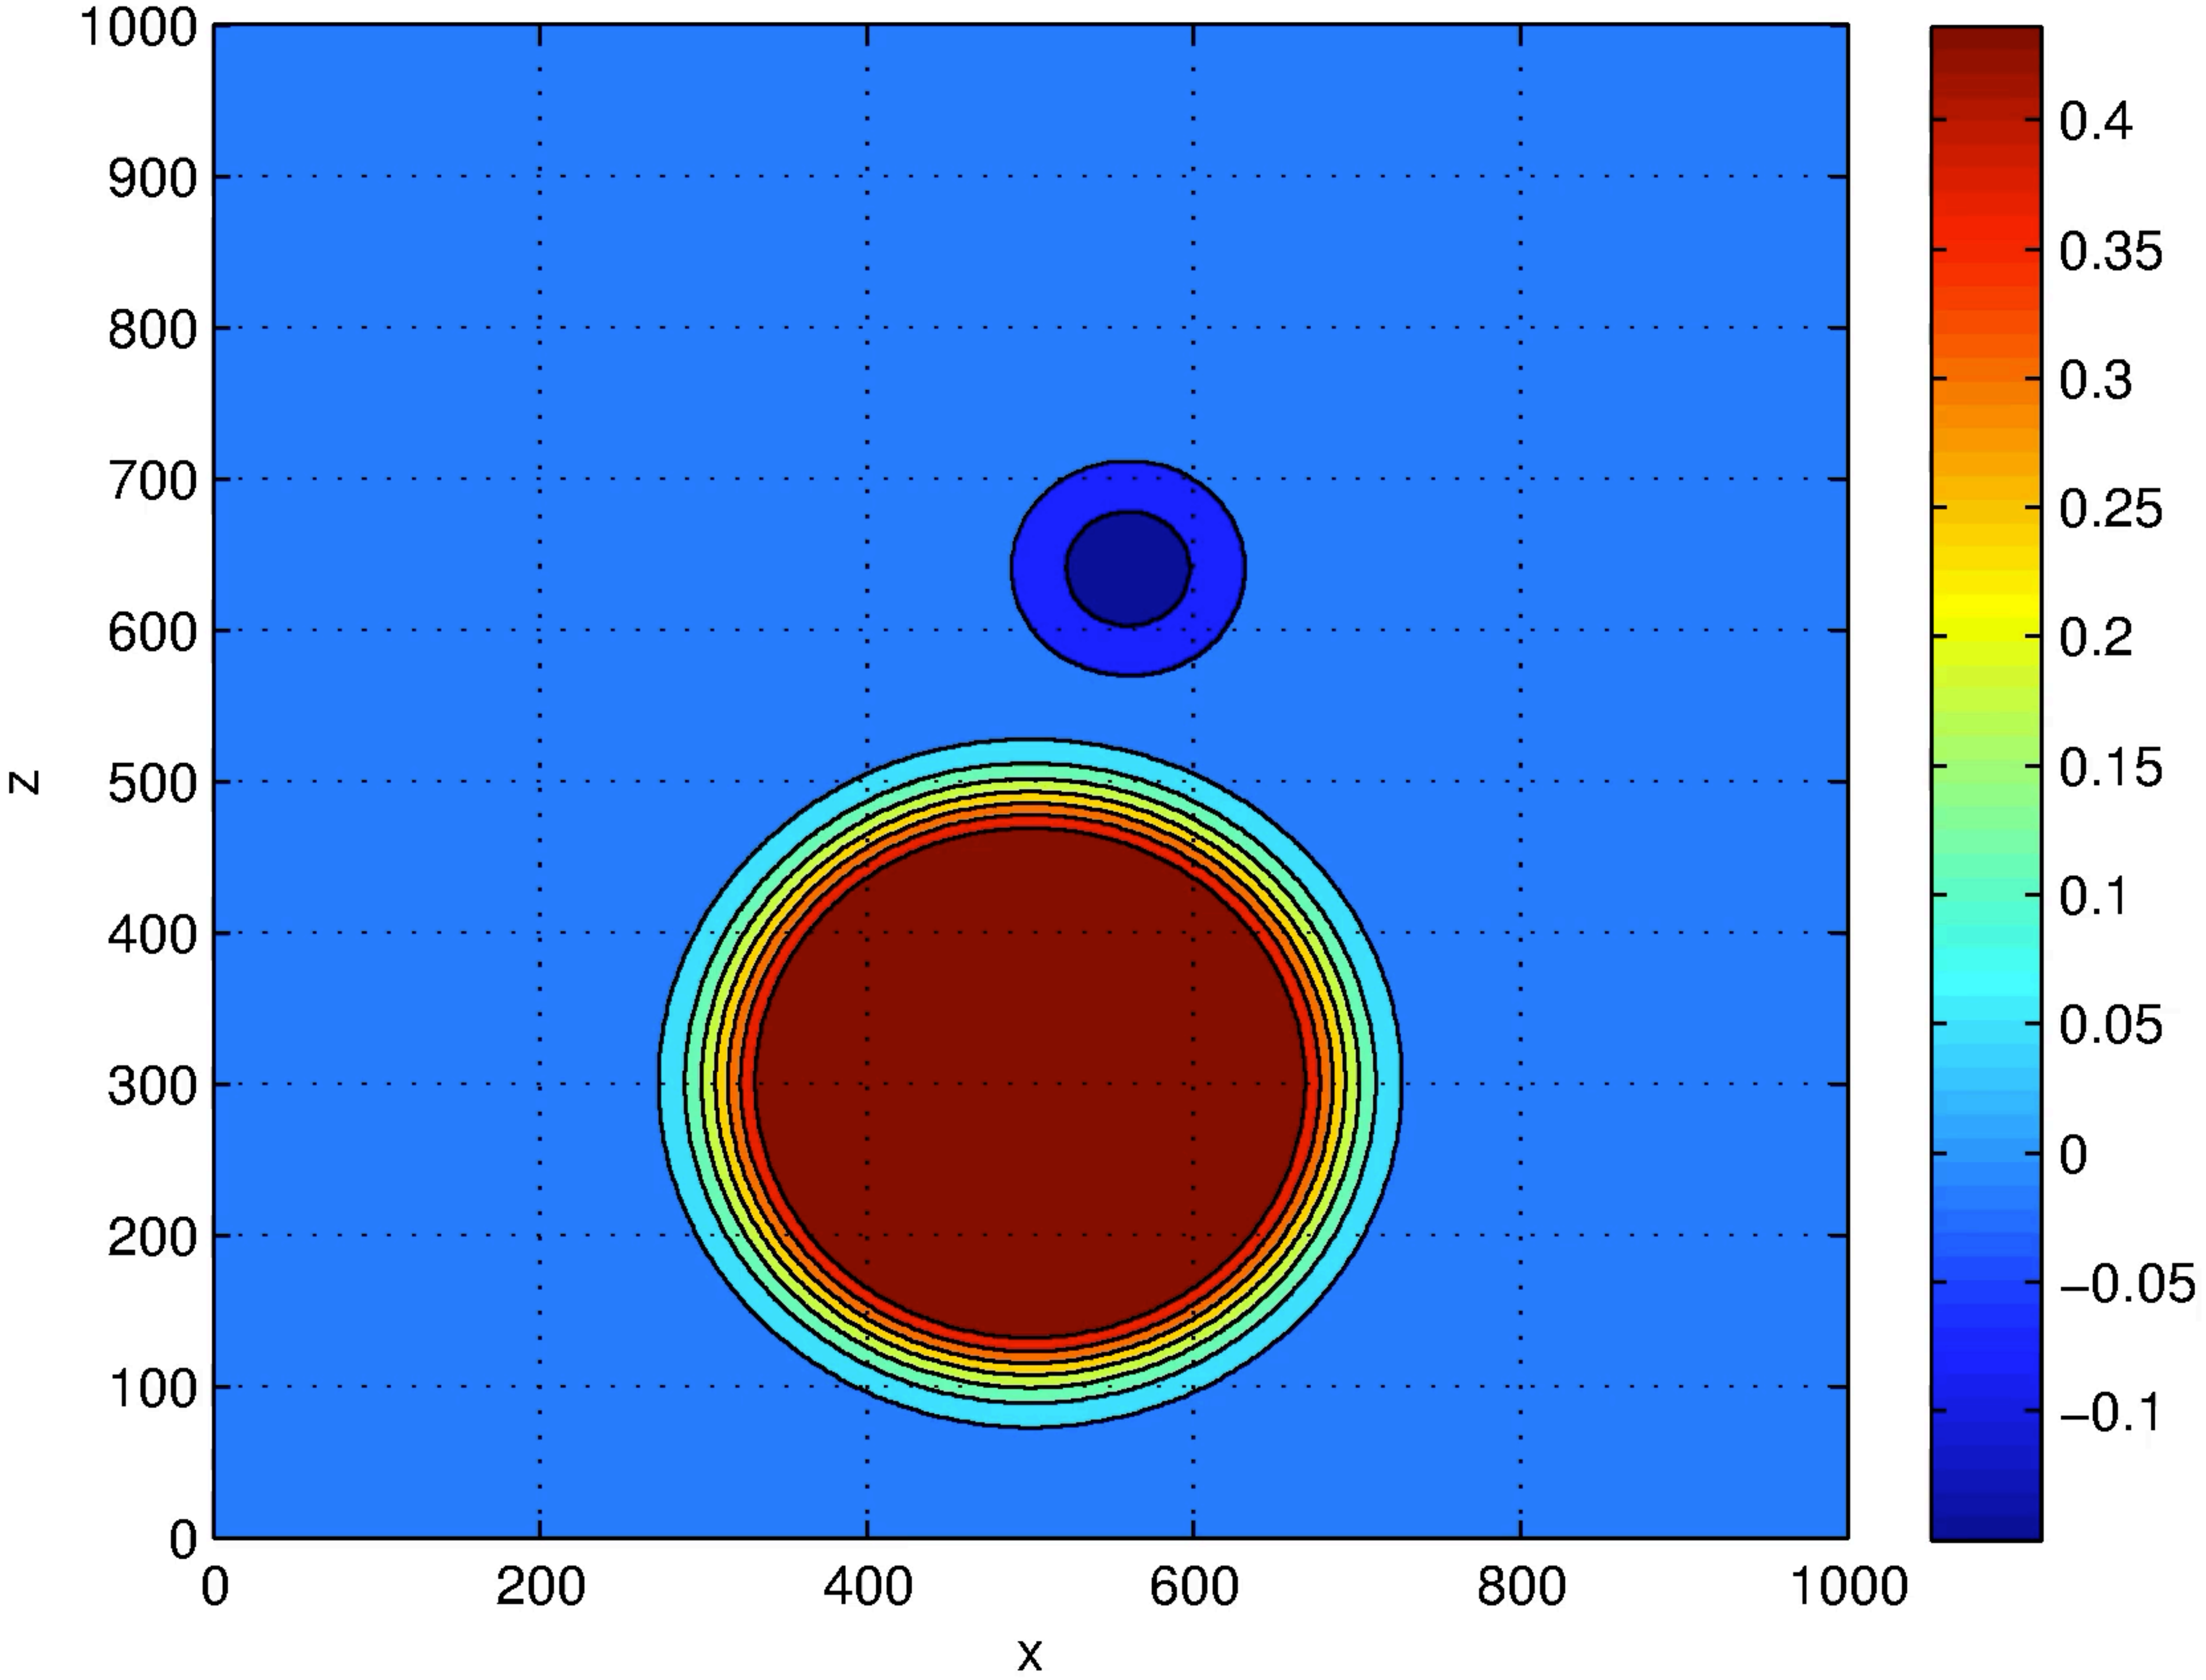

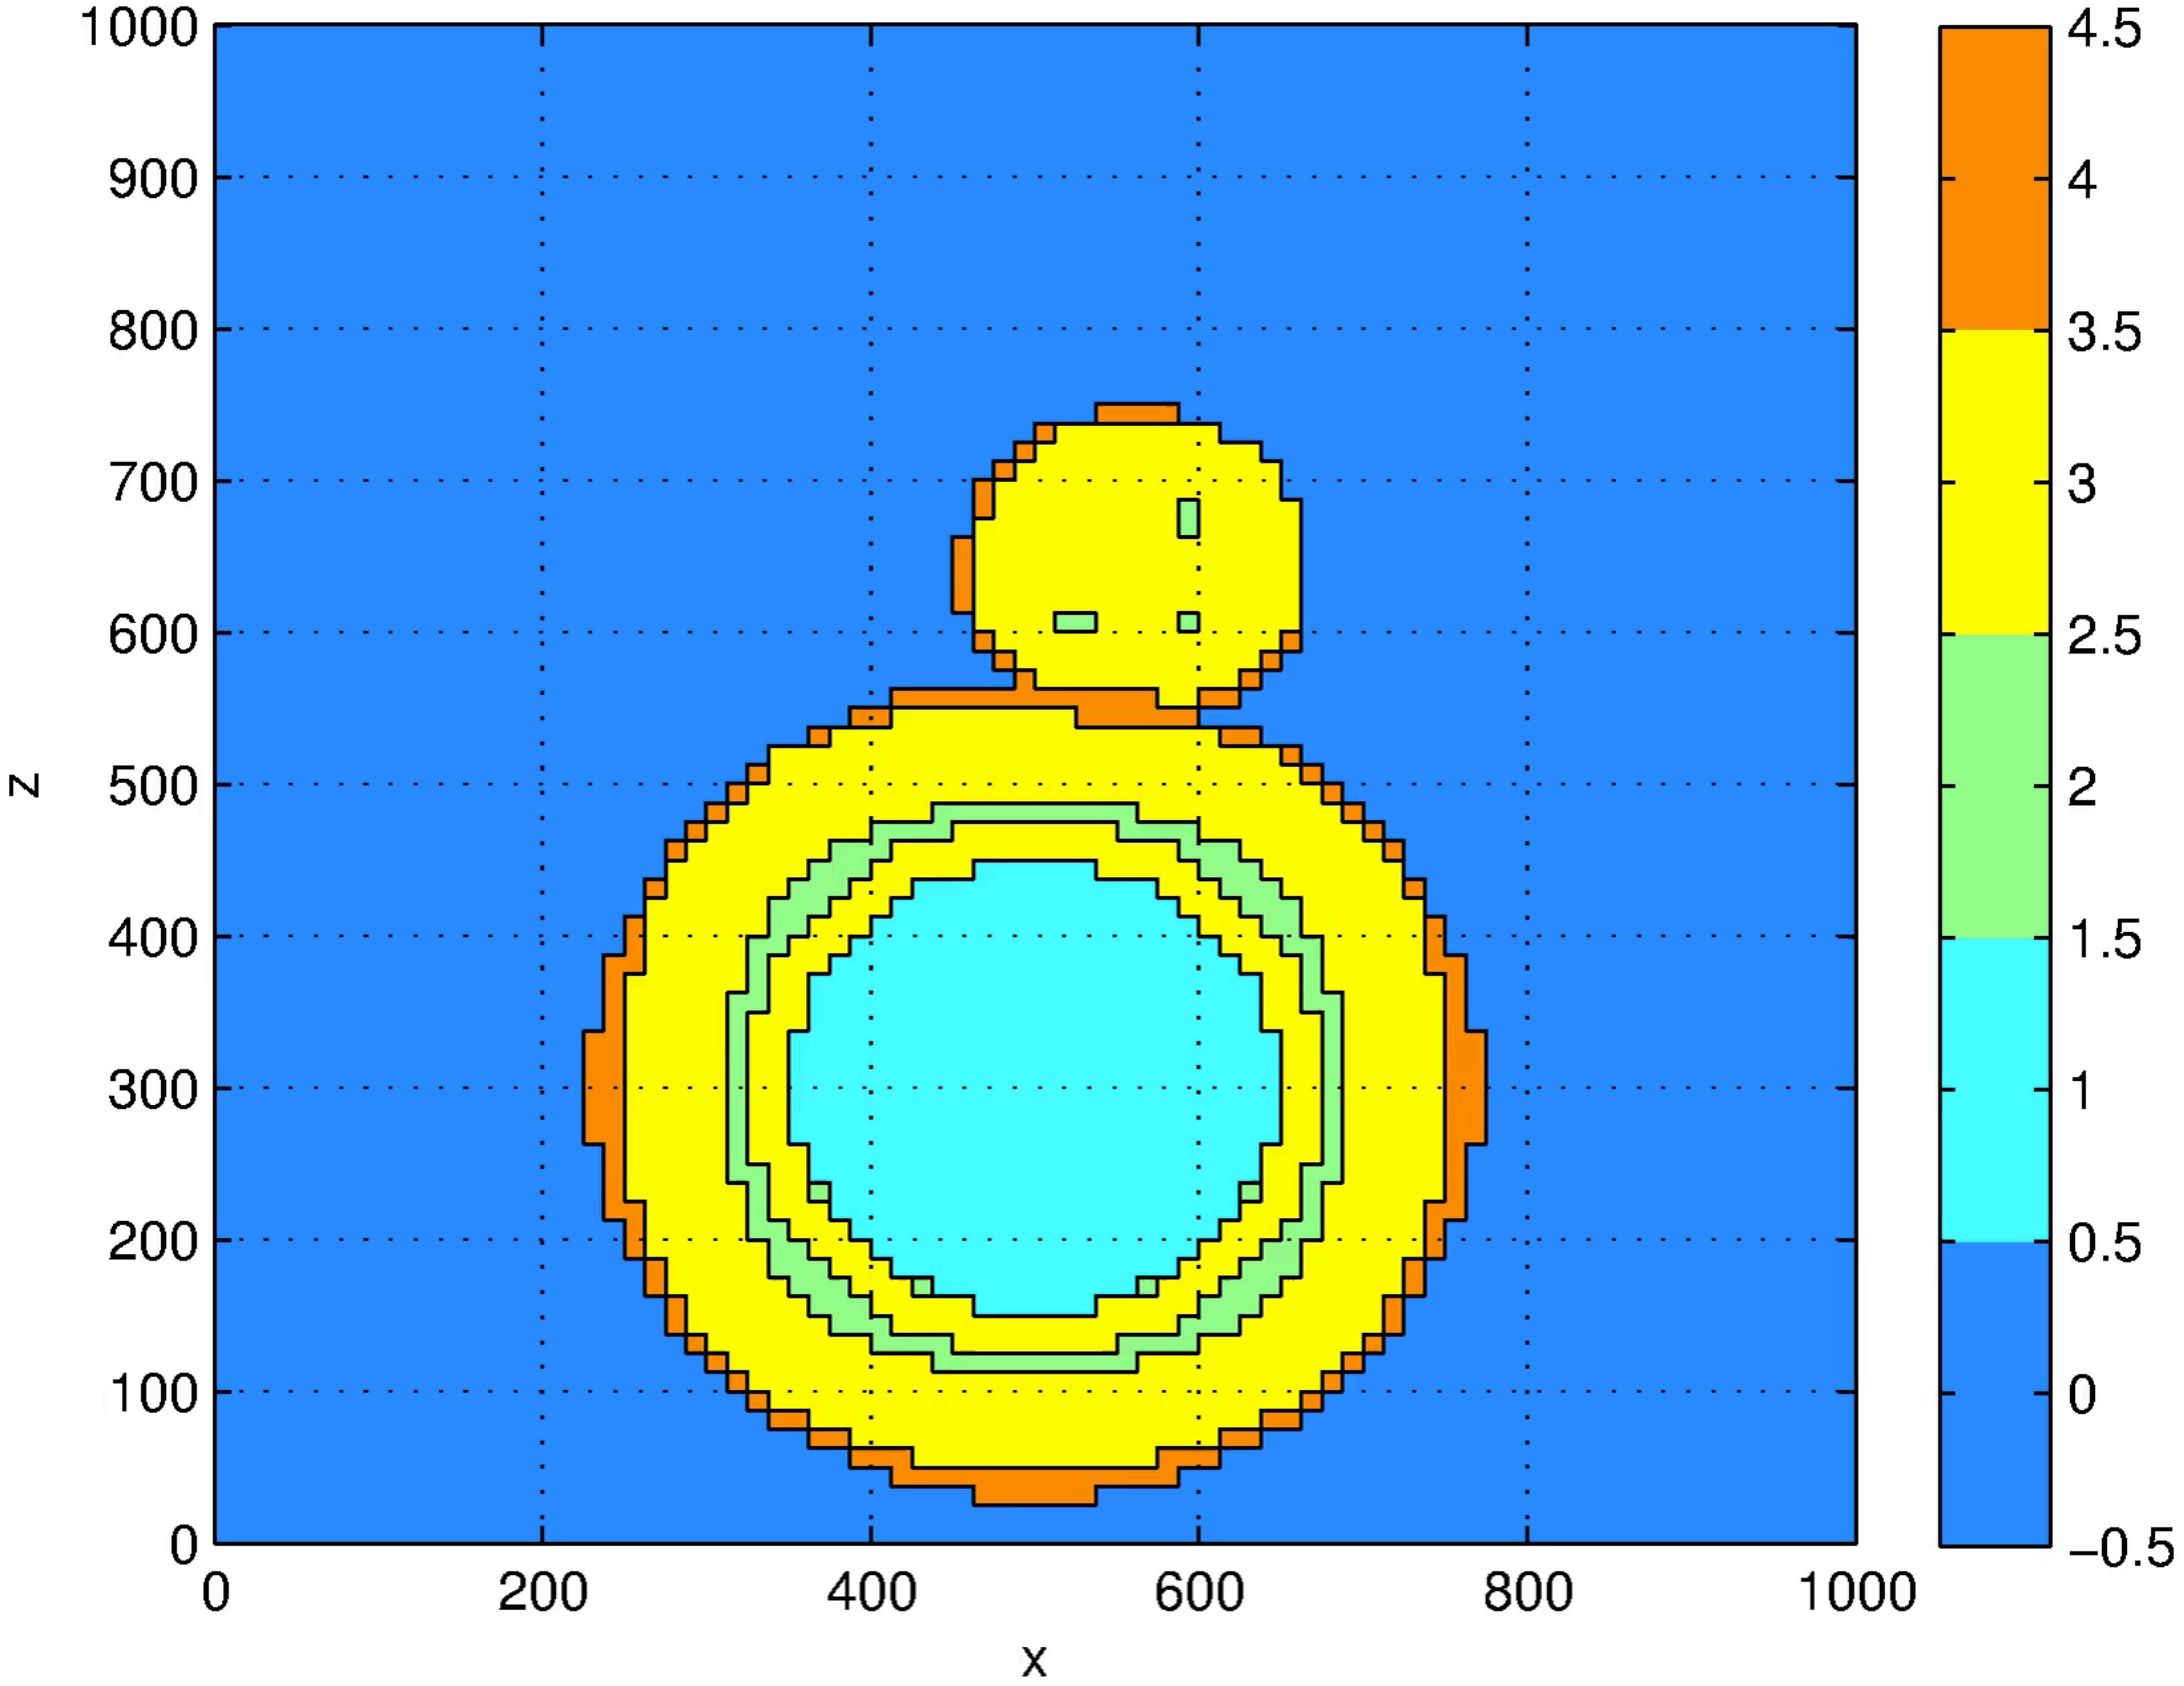

PANTHER: semi-Lagrangian DG with p-adaptivity colliding bubble test case from Robert (1993)

(PANTHER: P-Adaptive Numerical Tool for High order Efficient discRetizations, developed by Giovanni Tumolo)

PANTHER: semi-Lagrangian DG with p-adaptivity colliding bubble test case from Robert (1993)

(PANTHER: P-Adaptive Numerical Tool for High order Efficient discRetizations, developed by Giovanni Tumolo)

ESCAPE2: models and benchmarks

- IFS, IFS-FVM, SLDG, ICON-Ocean, NEMO
- Kronos: workload simulator developed at ECMWF
- goal: use dwarfs from many models of the atmospheric community to define weather and climate benchmarks
- start something like [top500.org](top500.org) that gives information on computer performance for weather and climate applications

Funded by the European Union

ECMWF Scalability Programme: External

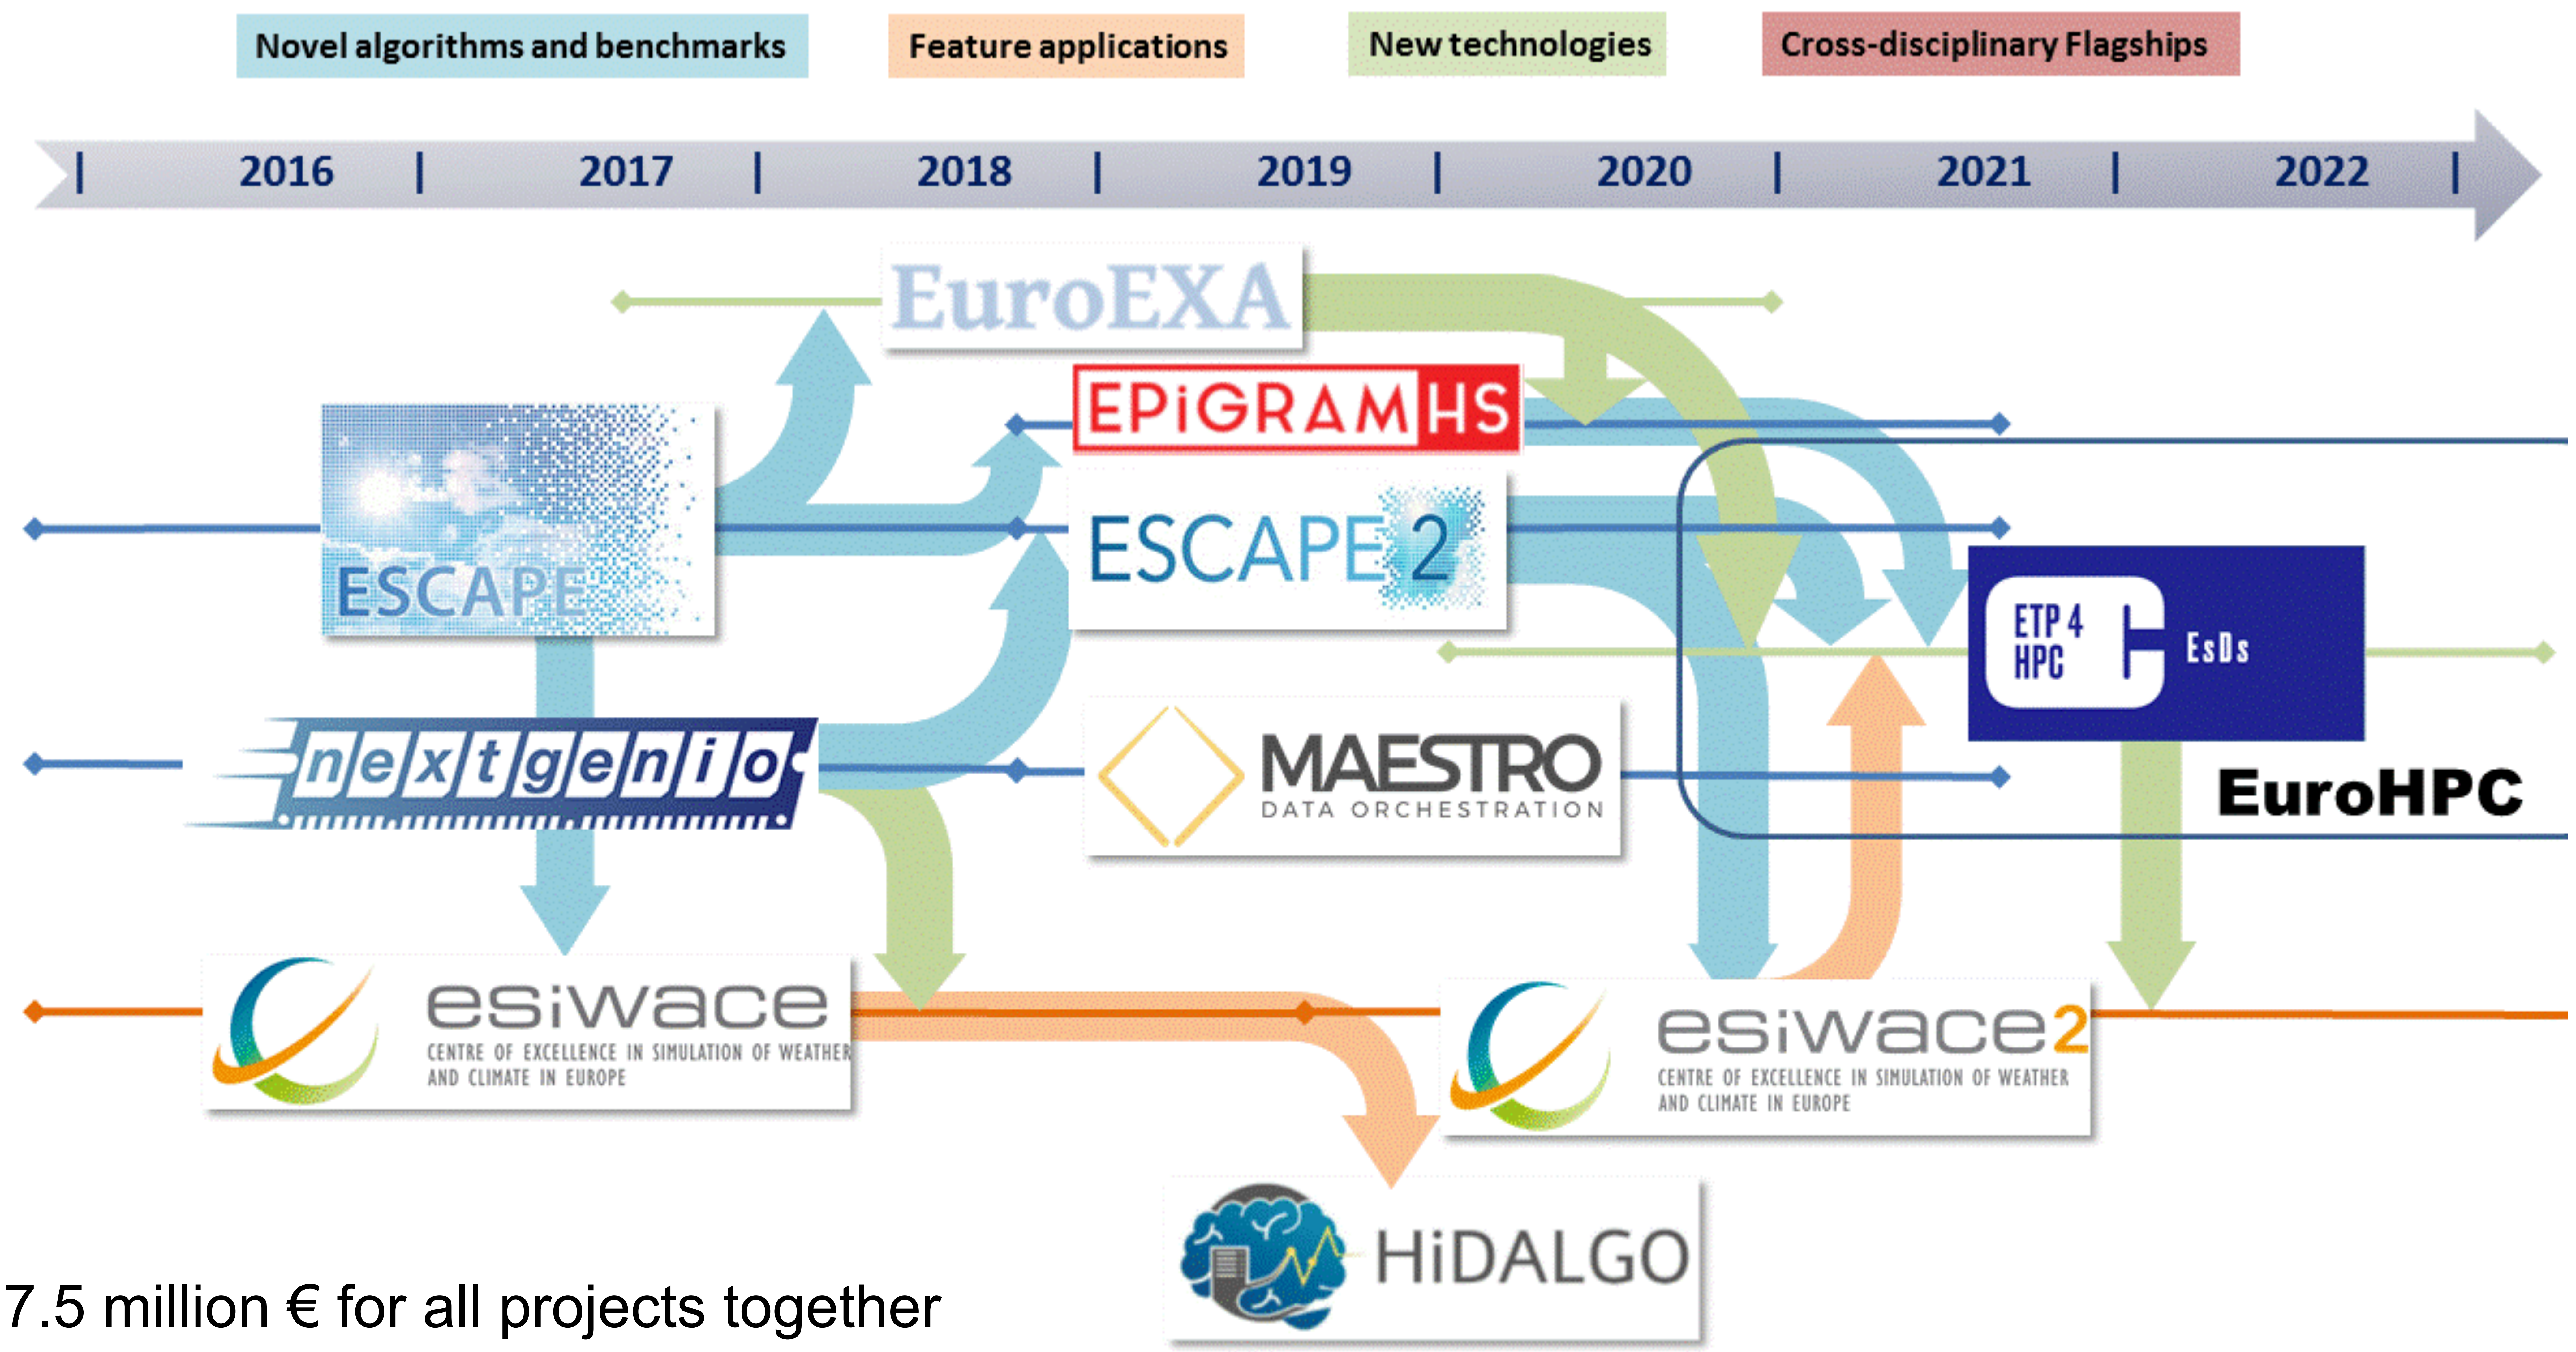

Roughly 7.5 million € for all projects together

One slide summary of ESCAPE

- dwarf concept: invite entire community to identify key patterns in terms of computation and communication which are crucial for the entire model performance
- demonstration of the domain specific language (DSL) concept as a route towards performance portability
- has been very useful for collaboration of vendors, operational weather prediction centres, HPC experts and academia

Image credits

- Images on slide 2 used under license from shutterstock.com