

# PySDM:

Bridging performance and pythonicity  
with Numba, Pythran and ThrustRTC

Piotr Bartman

# ABOUT THE PROJECT

Atmospheric Cloud Simulation Group  
Faculty of Mathematics and Computer Science,  
Jagiellonian University in Kraków, Poland

Super-droplet method (SDM) in Python: PySDM

Our projects:

<http://github.com/atmos-cloud-sim-uj/PySDM>



Author: Jan Mehlich

Jagiellonian University

# PySDM

CFD/Monte-Carlo simulations  
in aerosol-cloud-rain physics

Potential users:

- cloud physicists
- developers of parameterisations  
for weather & climate models

Motivation:

- novel modelling methods  
(probabilistic particle-based simulation)
- leveraging modern hardware and cloud computing
- maintainability and reproducibility requirements  
(of journals and scientific method)



# PySDM

## reproducibility goals

- “(…) code must be made accessible during the review process”
- „(…) no manual processing of the data: models are run by a script, and all pre- and post-processing is scripted”
- „(…) figures and tables must be scientifically reproducible from the scripts”
- „(…) if the code is not ready, then neither is the manuscript”

new 2019 GMD journal policy, doi:10.5194/gmd-12-2215-2019



## Geoscientific Model Development

An interactive open-access journal of the European Geosciences Union

# Pythonicity

Code is read much more often than it is written

Code readable - performance trade off

„‘two-language problem’ — researchers often prototype algorithms in a user-friendly language such as Python but then have to rewrite them in a faster language (...)”\*

```
def proj(vx, vy, vz, kx, ky, kz, inv_k_square_nozero):  
    tmp = (kx * vx + ky * vy + kz * vz) * inv_k_square_nozero  
    return vx - kx * tmp, vy - ky * tmp, vz - kz * tmp
```

==

```
subroutine proj(res, vx, vy, vz, kx, ky, kz, inv_k_square_nozero, N0, N1, N2)  
  
    implicit none  
  
    ! Input/Output  
    integer, intent(in) :: N0, N1, N2  
    double precision, intent(in) :: vx(2, N2, N1, N0), vy(2, N2, N1, N0), vz(2, N2, N1, N0)  
    double precision, intent(in) :: kx(N2, N1, N0), ky(N2, N1, N0), kz(N2, N1, N0)  
    double precision, intent(in) :: inv_k_square_nozero(N2, N1, N0)  
    double precision, intent(out) :: res(2, 3, N2, N1, N0)  
  
    ! Locals  
    double precision :: tmp(2)  
    integer :: i, j, k  
  
    do k = 1, N0  
        do j = 1, N1  
            do i = 1, N2  
                tmp(1:2) = (kx(i,j,k) * vx(1:2,i,j,k) &  
                    + ky(i,j,k) * vy(1:2,i,j,k) &  
                    + kz(i,j,k) * vz(1:2,i,j,k)) * inv_k_square_nozero(i,j,k)  
  
                res(1:2,1,i,j,k) = vx(1:2,i,j,k) - kx(i,j,k) * tmp(1:2)  
                res(1:2,2,i,j,k) = vy(1:2,i,j,k) - ky(i,j,k) * tmp(1:2)  
                res(1:2,3,i,j,k) = vz(1:2,i,j,k) - kz(i,j,k) * tmp(1:2)  
            enddo  
        enddo  
    enddo  
  
end subroutine proj
```


\*Nature 2019 “toolbox” column (on Julia), doi: 10.1038/d41586-019-02310-3

# PySDM

technological stack and workflows


General:

- open source
- automated tests
- code coverage



Python acceleration:

- LLVM
- multi-threading
- GPU computation

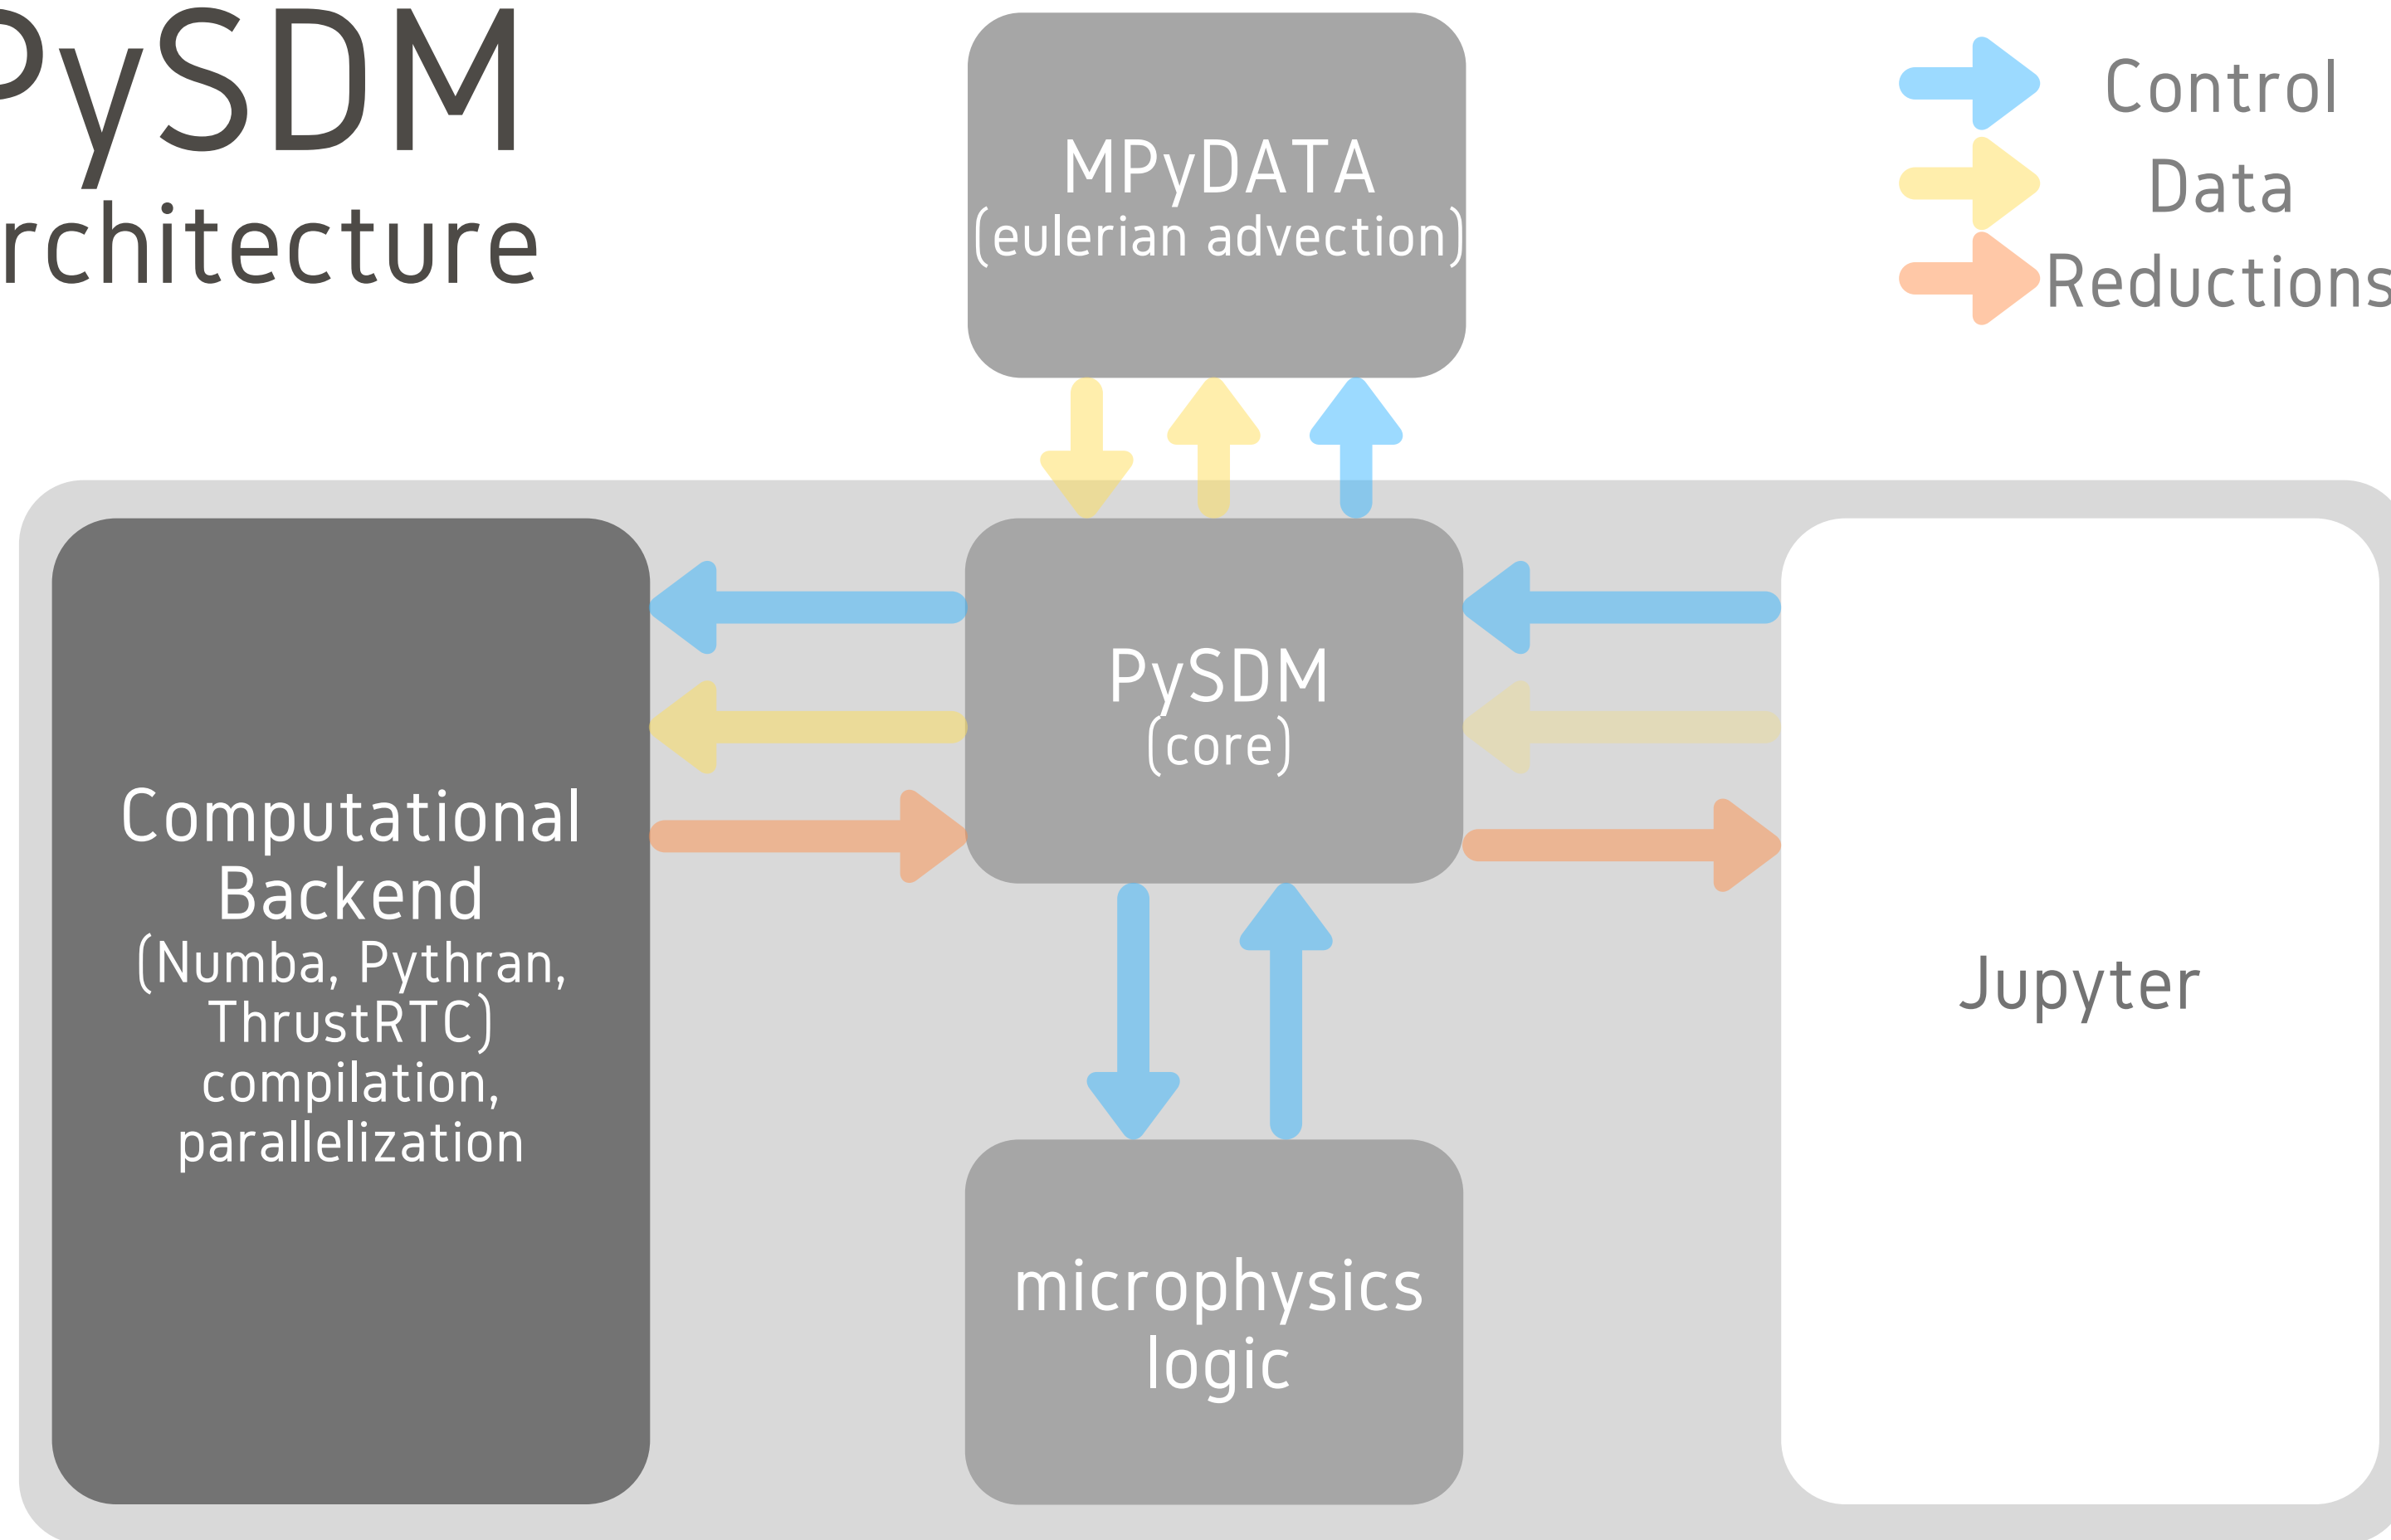


User interface:

- interactive examples & tutorials



# PySDM architecture




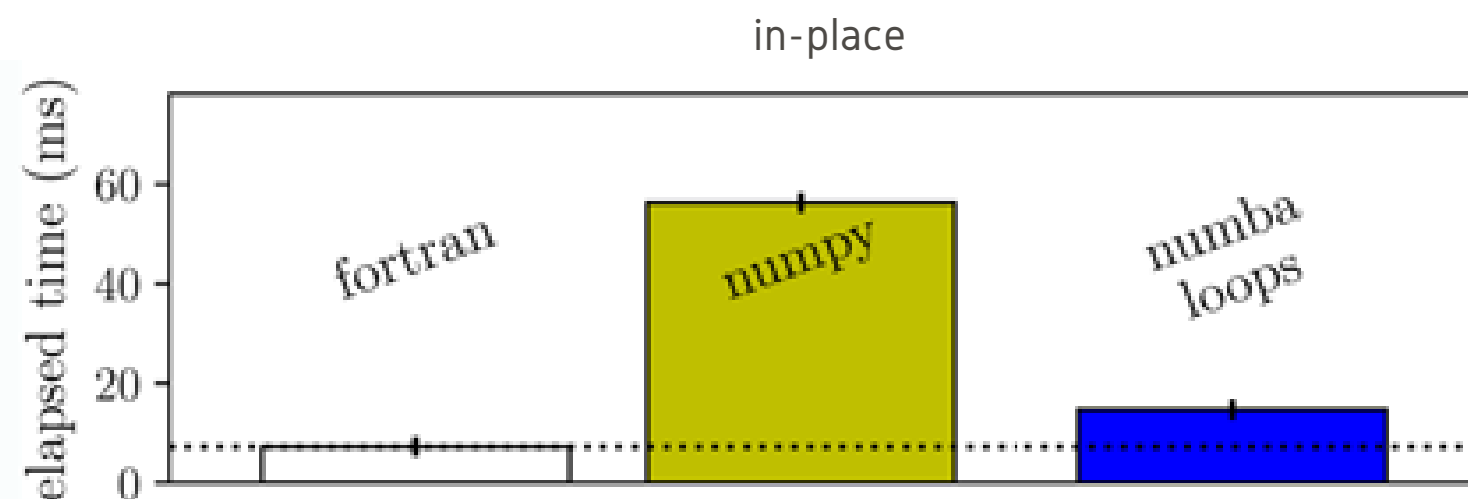
# Numba

[numba.pydata.org](http://numba.pydata.org) | [github.com/numba/numba/](https://github.com/numba/numba/)



- compiler for Python functions
- generates optimized machine code from pure Python code using the LLVM compiler infrastructure
- on-the-fly code generation
- native code generation for the CPU and GPU hardware
- integration with the Python scientific software stack (thanks to Numpy)

Mohanan et al. 2019, doi:10.5334/jors.238





# PySDM - Numba

maintainability & performance




```
166 @numba.njit(**conf.JIT_FLAGS)
167 def lv(T):
168     return const.l_tri + \
169         (const.c_pv - const.c_pw) * \
170         (T - const.T_tri)
```

```
37 def test_lv():
38     with DimensionalAnalysis():
39         # Arrange
40         si = constants.si
41         T = 300 * si.kelvins
42
43         # Act
44         latent_heat = formulae.lv(T)
45
46         # Assert
47         assert latent_heat.check('[energy] / [mass]')
```

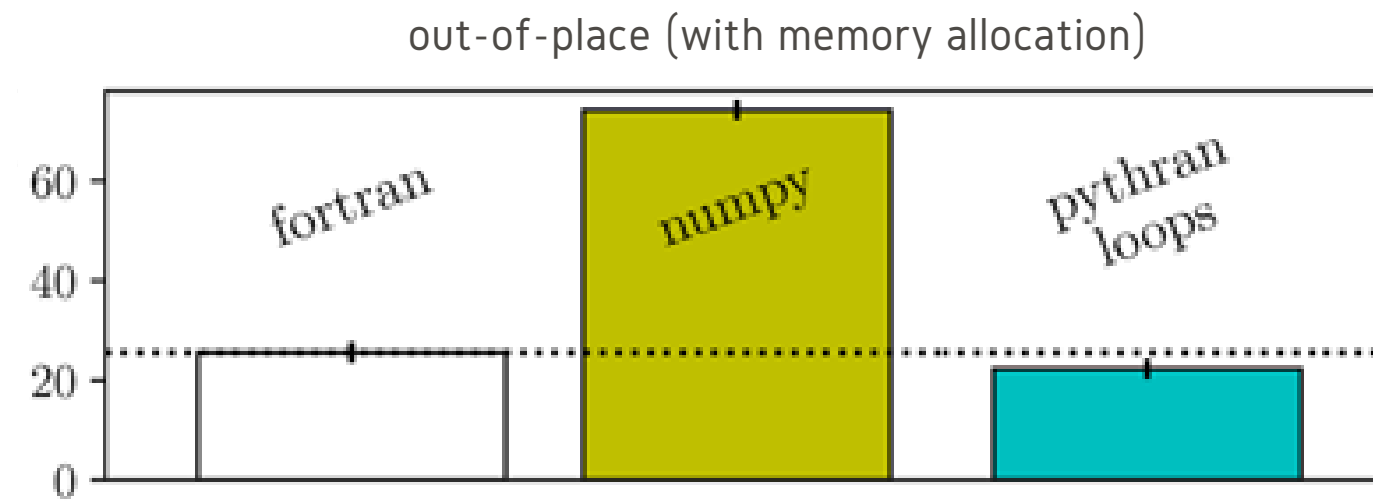
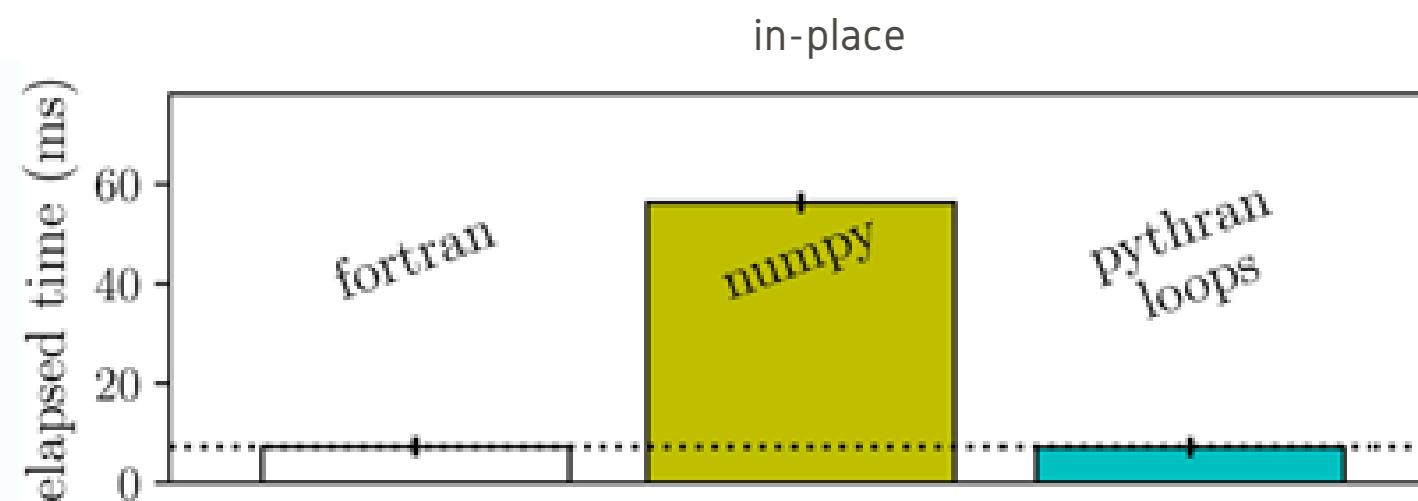
# Pythran

[pythran.readthedocs.io](http://pythran.readthedocs.io) | [github.com/serge-sans-paille/pythran](https://github.com/serge-sans-paille/pythran)



- ahead of time compiler for a subset of the Python language, with a focus on scientific computing
- takes a Python module annotated with a few interface description and turns it into a native Python module with the same interface
- meant to efficiently compile scientific programs
- focus on of multi-cores and SIMD instruction units

Mohanan et al. 2019, doi:10.5334/jors.238



# ThrustRTC + CURandRTC

[github.com/fynv/ThrustRTC](https://github.com/fynv/ThrustRTC) | [github.com/fynv/CurandRTC](https://github.com/fynv/CurandRTC)



- library of general GPU algorithms, functionally similar to Thrust, that can be used in non-C++ programming languages (Python)
- Thrust ([docs.nvidia.com/cuda/thrust](https://docs.nvidia.com/cuda/thrust)):
  - high-level interface enabling performance portability between GPUs and multicore CPUs
  - Interoperability with established technologies (such as CUDA, TBB, and OpenMP)
- CURandRTC: random number generator using the XORWOW algorithm

# PySDM - ThrustRTC

maintainability & performance



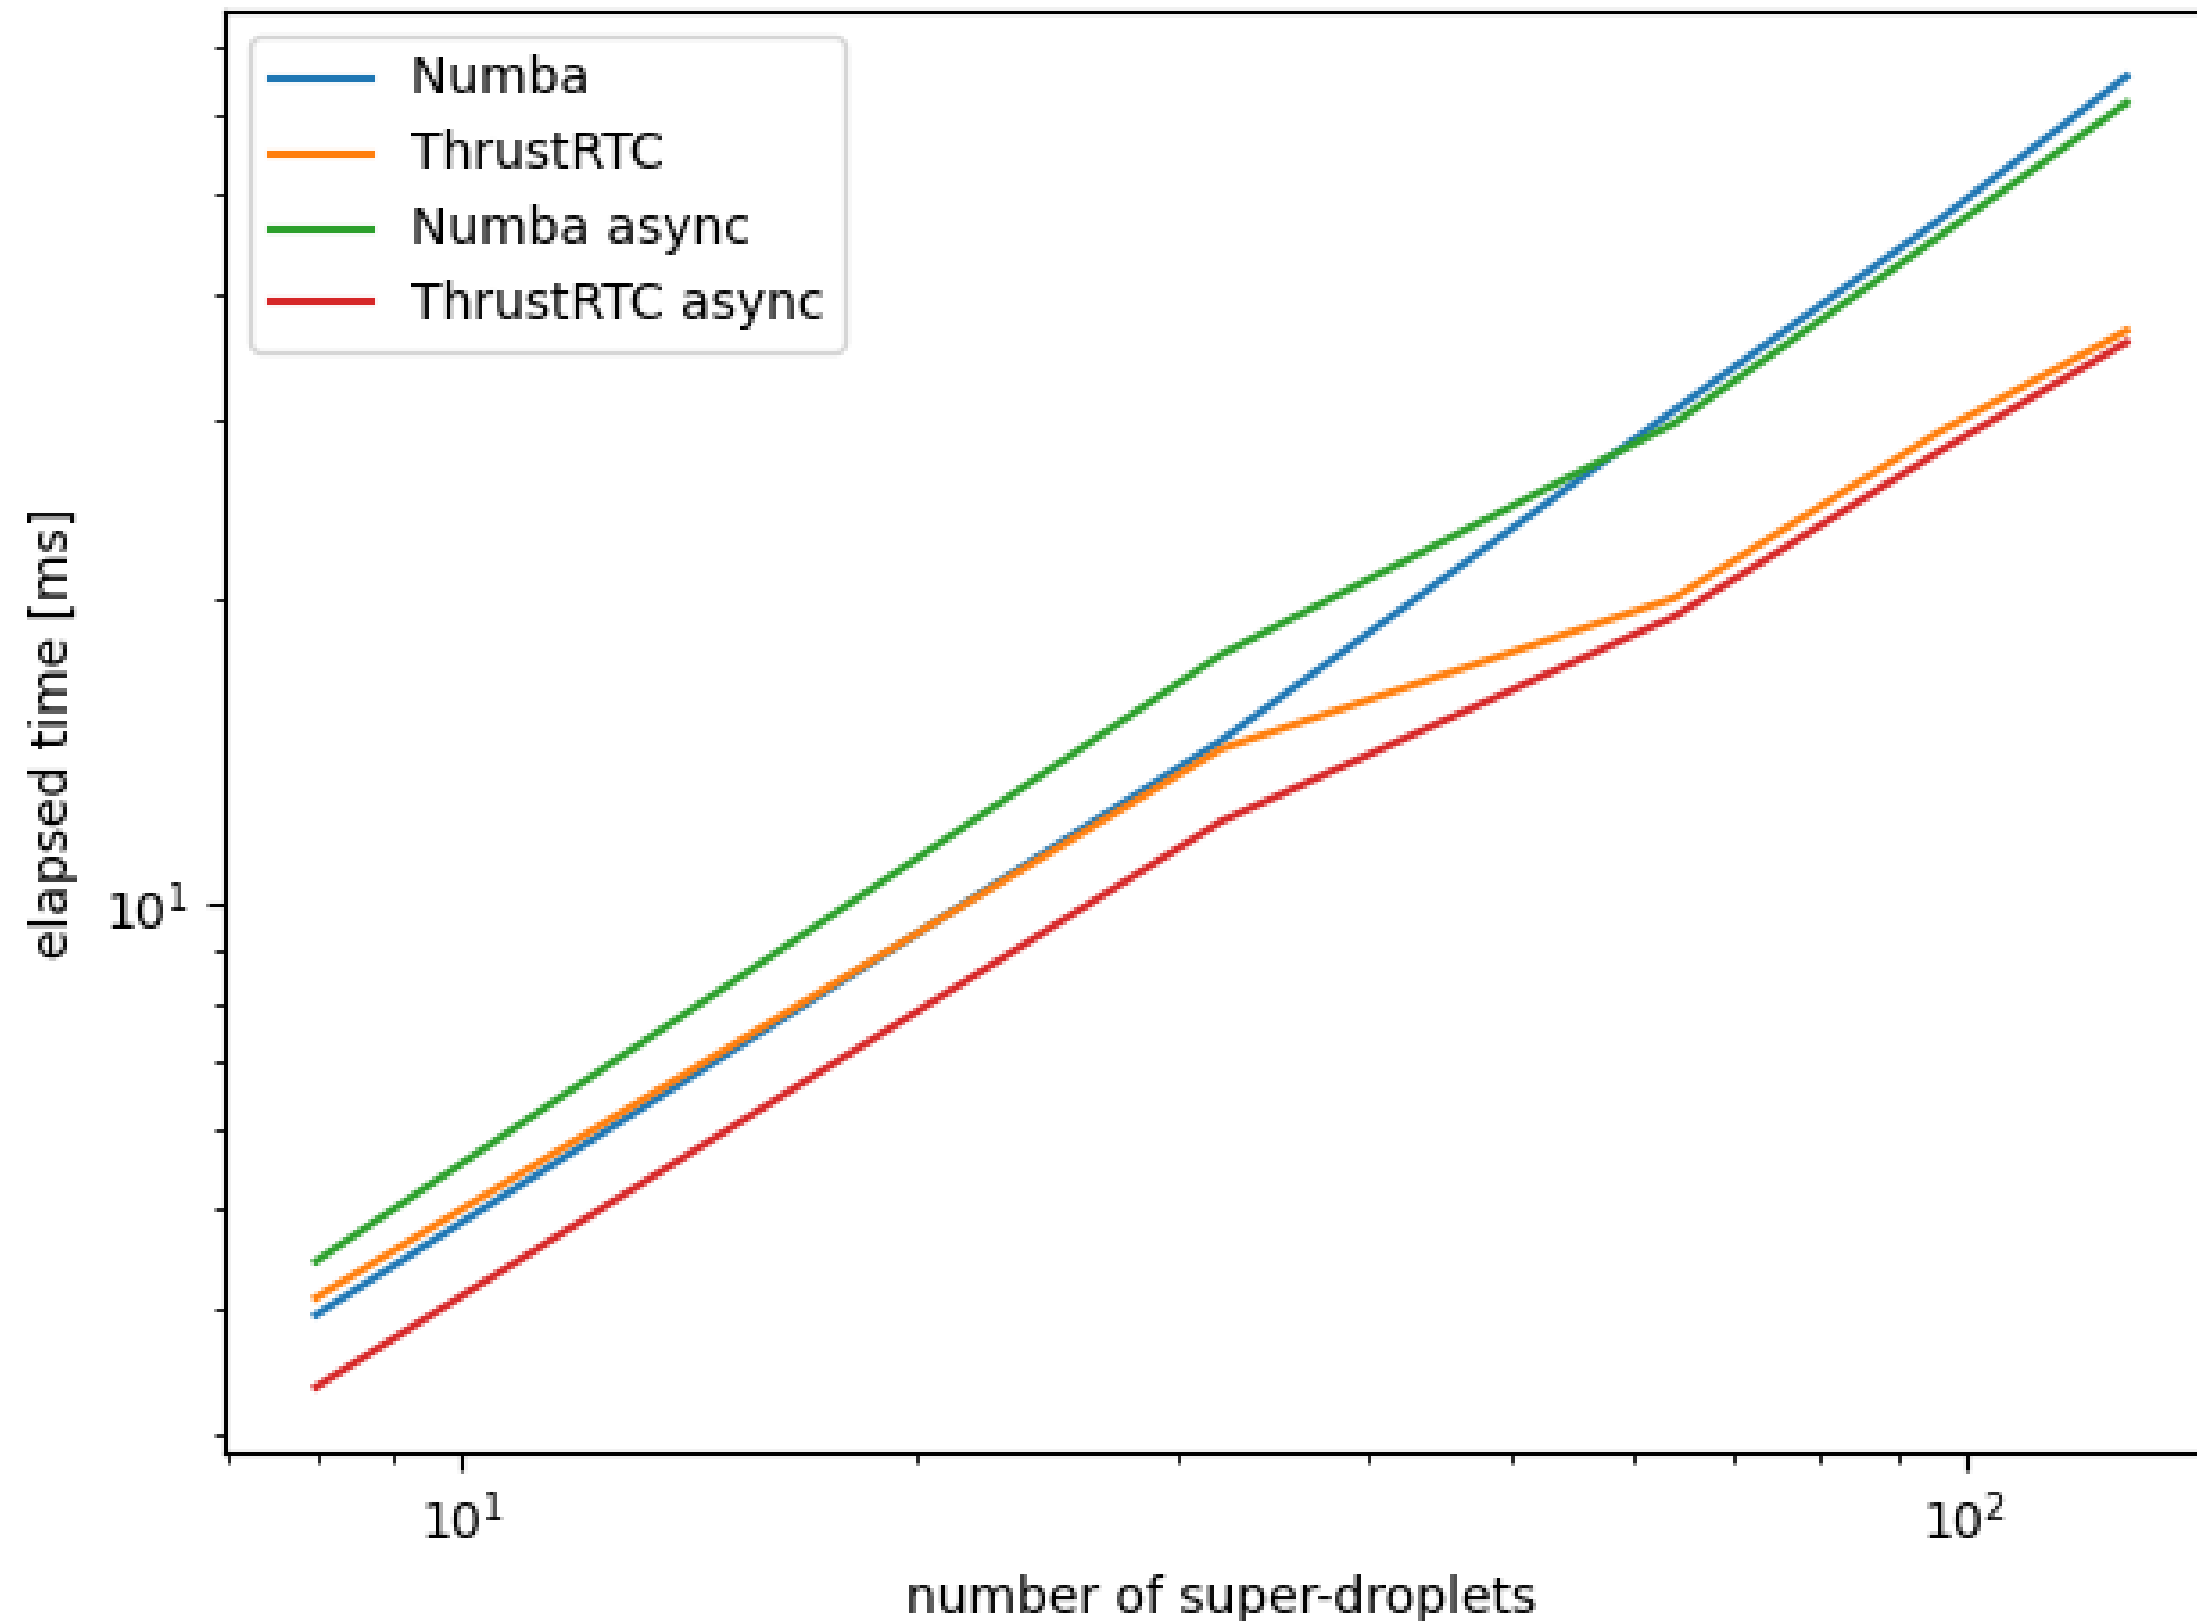
```
20 def add(output, addend):  
21     trtc.Transform_Binary(addend, output, output, trtc.Plus())
```

```
114 def flag_precipitated(cell, position, idx, length, healthy):  
115     idx_length = trtc.DVInt64(idx.size())  
116     n_dims = trtc.DVInt64(len(cell.shape))  
117     kernel = trtc.For(  
118         ['idx', 'idx_lth', 'n_dims', 'healthy', 'cell', 'position'], "i", ''  
119         id = idx_lth * (n_dims-1) + i;  
120         if (cell[id] == 0 && position[id] < 0) {  
121             idx[i] = idx_lth;  
122             healthy[0] = 0;  
123         }  
124         ''')  
125     kernel.launch_n(length, [idx, idx_length, n_dims, healthy, cell, position])
```



# ThrustRTC vs Numba

[github.com/atmos-cloud-sim-uj/PySDM/tree/master/PySDM\\_examples/ICMW\\_2012\\_case\\_1](https://github.com/atmos-cloud-sim-uj/PySDM/tree/master/PySDM_examples/ICMW_2012_case_1)



async - computations on the CPU  
are overlapped with  
computations on the accelerator






Numba - CPU multi-threading  
ThrustRTC - GPU resident

# PySDM

## portability

Build jobs

View config

✓ # 182.1	AMD64		Python 3.8 with newest packages on Linux	🕒 19 min 45 sec
✓ # 182.2	AMD64		Python 3.7 on Linux numba::parallel=False	🕒 21 min 10 sec
✓ # 182.3	AMD64		Python 3.7 on Linux numba::parallel=True	🕒 18 min 39 sec
✓ # 182.4	AMD64		Python 3.7 on OSX	🕒 30 min 11 sec
✓ # 182.5	AMD64		Python 3.7 on Windows	🕒 32 min 20 sec

build passing coverage 66%

# PySDM

PySDM simulates the dynamics of population of particles immersed in moist air using the particle-based (a.k.a. super-droplet) approach to represent aerosol/cloud/rain microphysics. The package features a Pythonic implementation of the Super-Droplet Method (SDM) Monte-Carlo algorithm for representing collisional growth (Shima et al. 2009), hence the name.

## Demos:

- Shima et al. 2009 Fig. 2 [launch](#) [binder](#)
- Arabas & Shima 2017 Fig. 5 [launch](#) [binder](#)
- Yang et al. 2018 Fig. 2: [launch](#) [binder](#)
- ICMW 2012 case 1 (work in progress) [launch](#) [binder](#)

## Tutorials:


- Introduction [launch](#) [binder](#)
- Coalescence [launch](#) [binder](#)

## Credits:

Development of PySDM is supported by the EU through a grant of the Foundation for Polish Science (POIR.04.04.00-00-5E1C/18).

# PySDM

## lessons learned



### Numba:

- portable and pure-python (yet not fully pythonic; OOP kills performance)
- easy debugging
- little-to-no control over (and huge performance sensitivity to) inlining/optimization

### Pythran:

- essentially equivalent performance and parallelization features as Numba (both based on LLVM and OpenMP)
- cross-platform support is low-priority



# PySDM

## lessons learned



### ThrustRTC (+CURandRTC):

- viable high-level Python abstractions for parallel GPU computations
- tricky debugging for custom kernels (CUDA-free C code in Python strings)
- no option to execute on CPU/threads (unlike original Thrust)

Special thanks for Numba, Pythran, ThrustRTC developers for quickly responses (2, 5, 1 fixed issues respectively)

**MORE:**

[github.com/atmos-cloud-sim-uj/PySDM](https://github.com/atmos-cloud-sim-uj/PySDM)

[github.com/piotr bartman](https://github.com/piotr bartman)

mail: [piotr.bartman@student.uj.edu.pl](mailto:piotr.bartman@student.uj.edu.pl)