



arm

Improving Python- based workloads in HPC

Florent.Lebeau@arm.com

Outline

- Why using Python in HPC?
- Profiling Python
- Introduction to Arm MAP
- Optimizing Python-based workloads on single and many nodes

Python in HPC

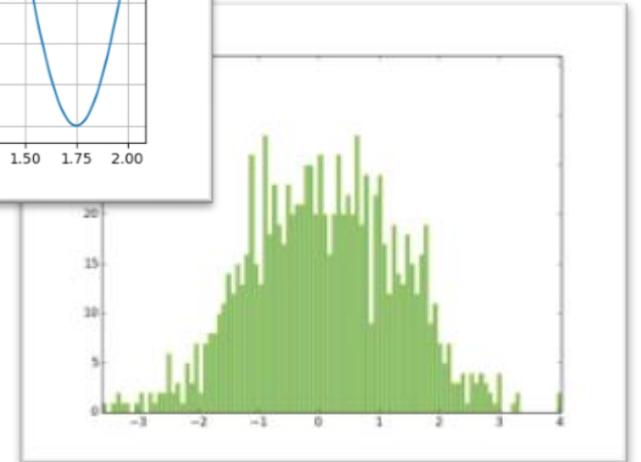
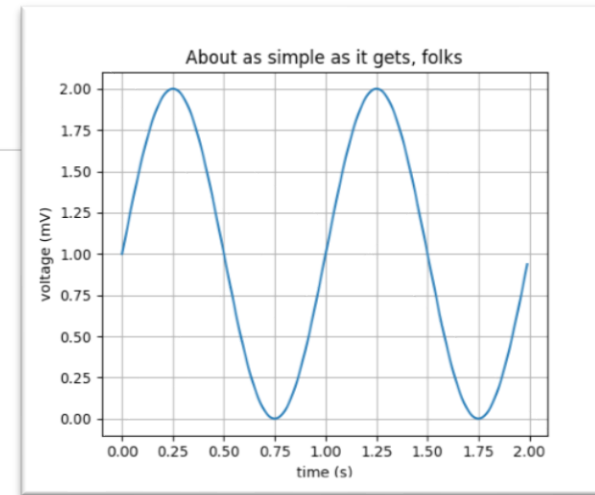
- High-level
- Readable
- Free
- Cross-platform
- Extensible
 - \$ pip install matplotlib
 - \$ conda install matplotlib



```
import matplotlib.pyplot as plt
import numpy as np

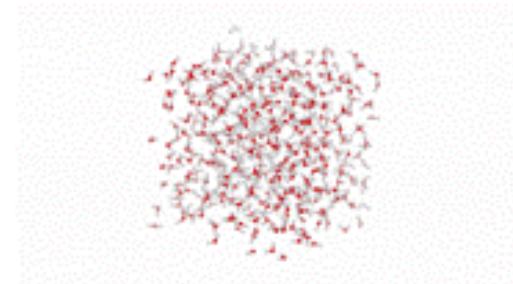
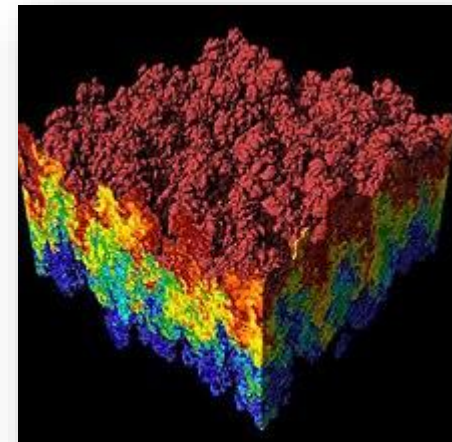
t = np.arange(0.0, 2.0, 0.01)
s = 1 + np.sin(2*np.pi*t)
plt.plot(t, s)

plt.xlabel('time (s)')
plt.ylabel('voltage (mV)')
plt.title('About as simple as it gets, folks')
plt.grid(True)
plt.savefig("test.png")
plt.show()
```



Python in HPC

- Essential modules:
 - **NumPy**: support of large multi-dimensional arrays and matrices
 - **SciPy**: support for linear algebra, integration, interpolation, FFT, ...
 - **MPI4Py**: provides bindings of the MPI standard
- Rely on highly-optimized libraries
 - Written in lower-level languages: C, FORTRAN, ...
 - BLAS, LAPACK, FFTW, ...
- Can easily be interfaced with other languages



Profiling Python

- time

```
import time

[...]

t = time.clock()

[...]

print "time spent:", time.clock() - t, "seconds"
```

- cProfile

- `$ python -m cProfile -s tottime ./laplace1.py slow 100 500`

```
127857 function calls (124922 primitive calls) in 25.961 seconds

Ordered by: internal time
ncalls  tottime  percall  cumtime  percall filename:lineno(function)
    500   25.568    0.051   25.631    0.051 laplace1.py:101(slowTimeStep)
  49706    0.064    0.000    0.064    0.000 {range}
      1    0.034    0.034   25.961   25.961 laplace1.py:23(<module>)

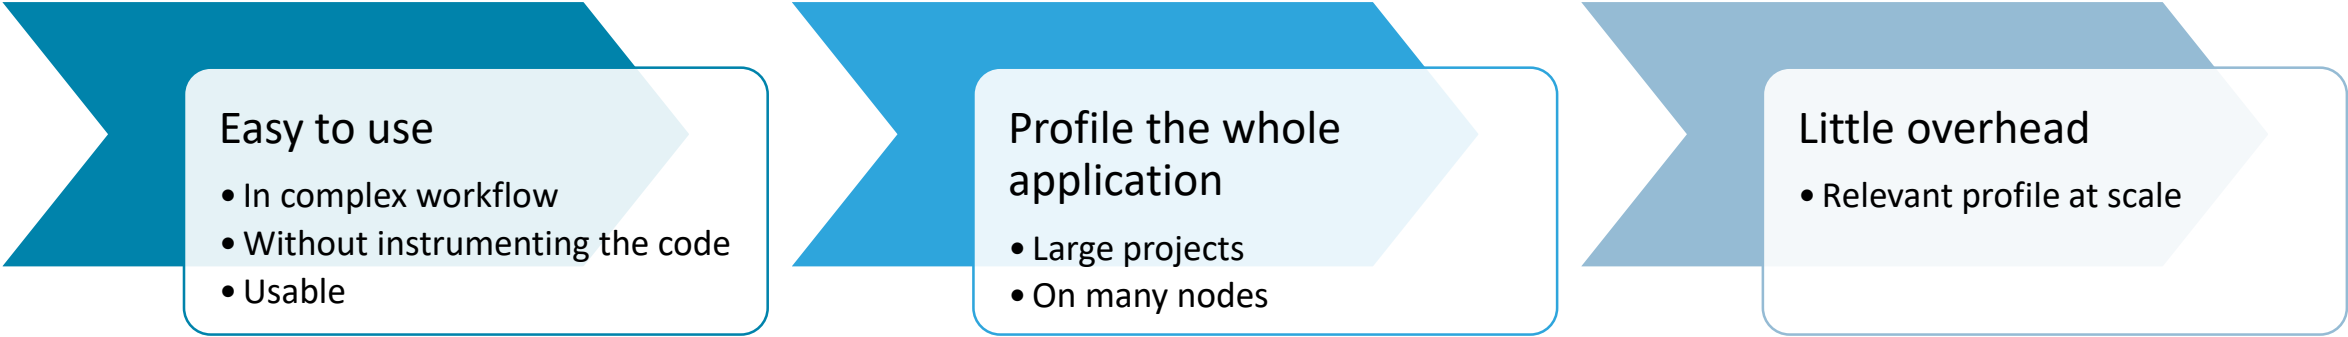
[...]
```

- Line_profiler

- `$ kernprof -l -v ./laplace1.py slow 100 500`

Line #	Hits	Time	Per Hit	% Time	Line Contents
99					@profile
100					def slowTimeStep(self, dt=0.0):
[...]					
109	49500	35571.0	0.7	0.1	for i in range(1, nx-1):
110	4851000	3698877.0	0.8	9.6	for j in range(1, ny-1):
111	4802000	4984464.0	1.0	12.9	tmp = u[i,j]
112	4802000	8034699.0	1.7	20.8	u[i,j] = ((u[i-1, j] + u[i+1, j])*dy2 +
113	4802000	10799762.0	2.2	28.0	(u[i, j-1] + u[i, j+1])*dx2)*dnr_inv
[...]					

Profiling Python: challenges



Easy to use

- In complex workflow
- Without instrumenting the code
- Usable

Profile the whole application

- Large projects
- On many nodes

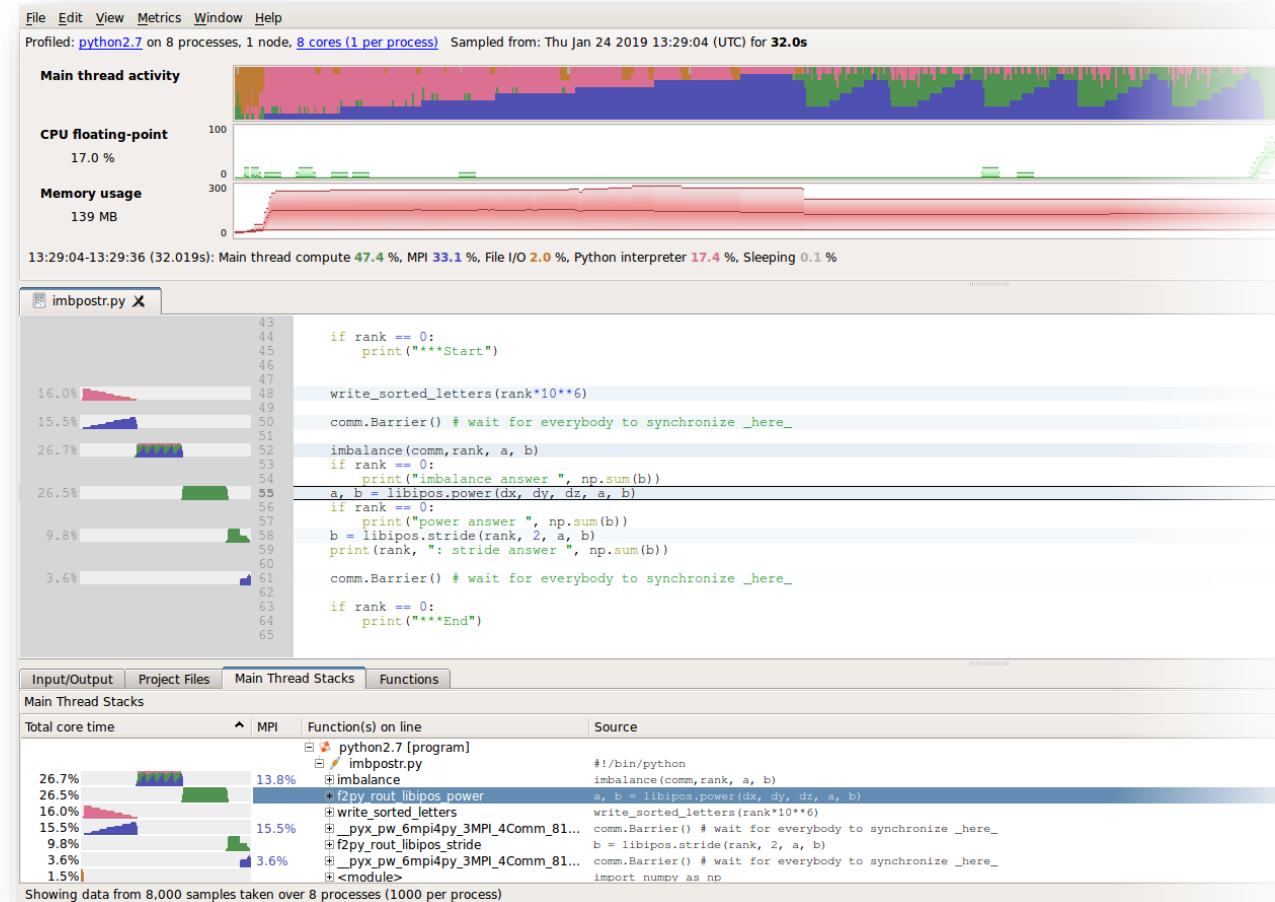
Little overhead

- Relevant profile at scale

Arm MAP profiler

arm MAP

- Arm MAP is a lightweight multi-node profiling tool
 - Shows processes and threads activity over time
 - Source code is annotated
 - Information aggregated by stacks and function
- **Compute**, **IO**, **MPI**, **Python interpreter**
- Part of Arm Forge



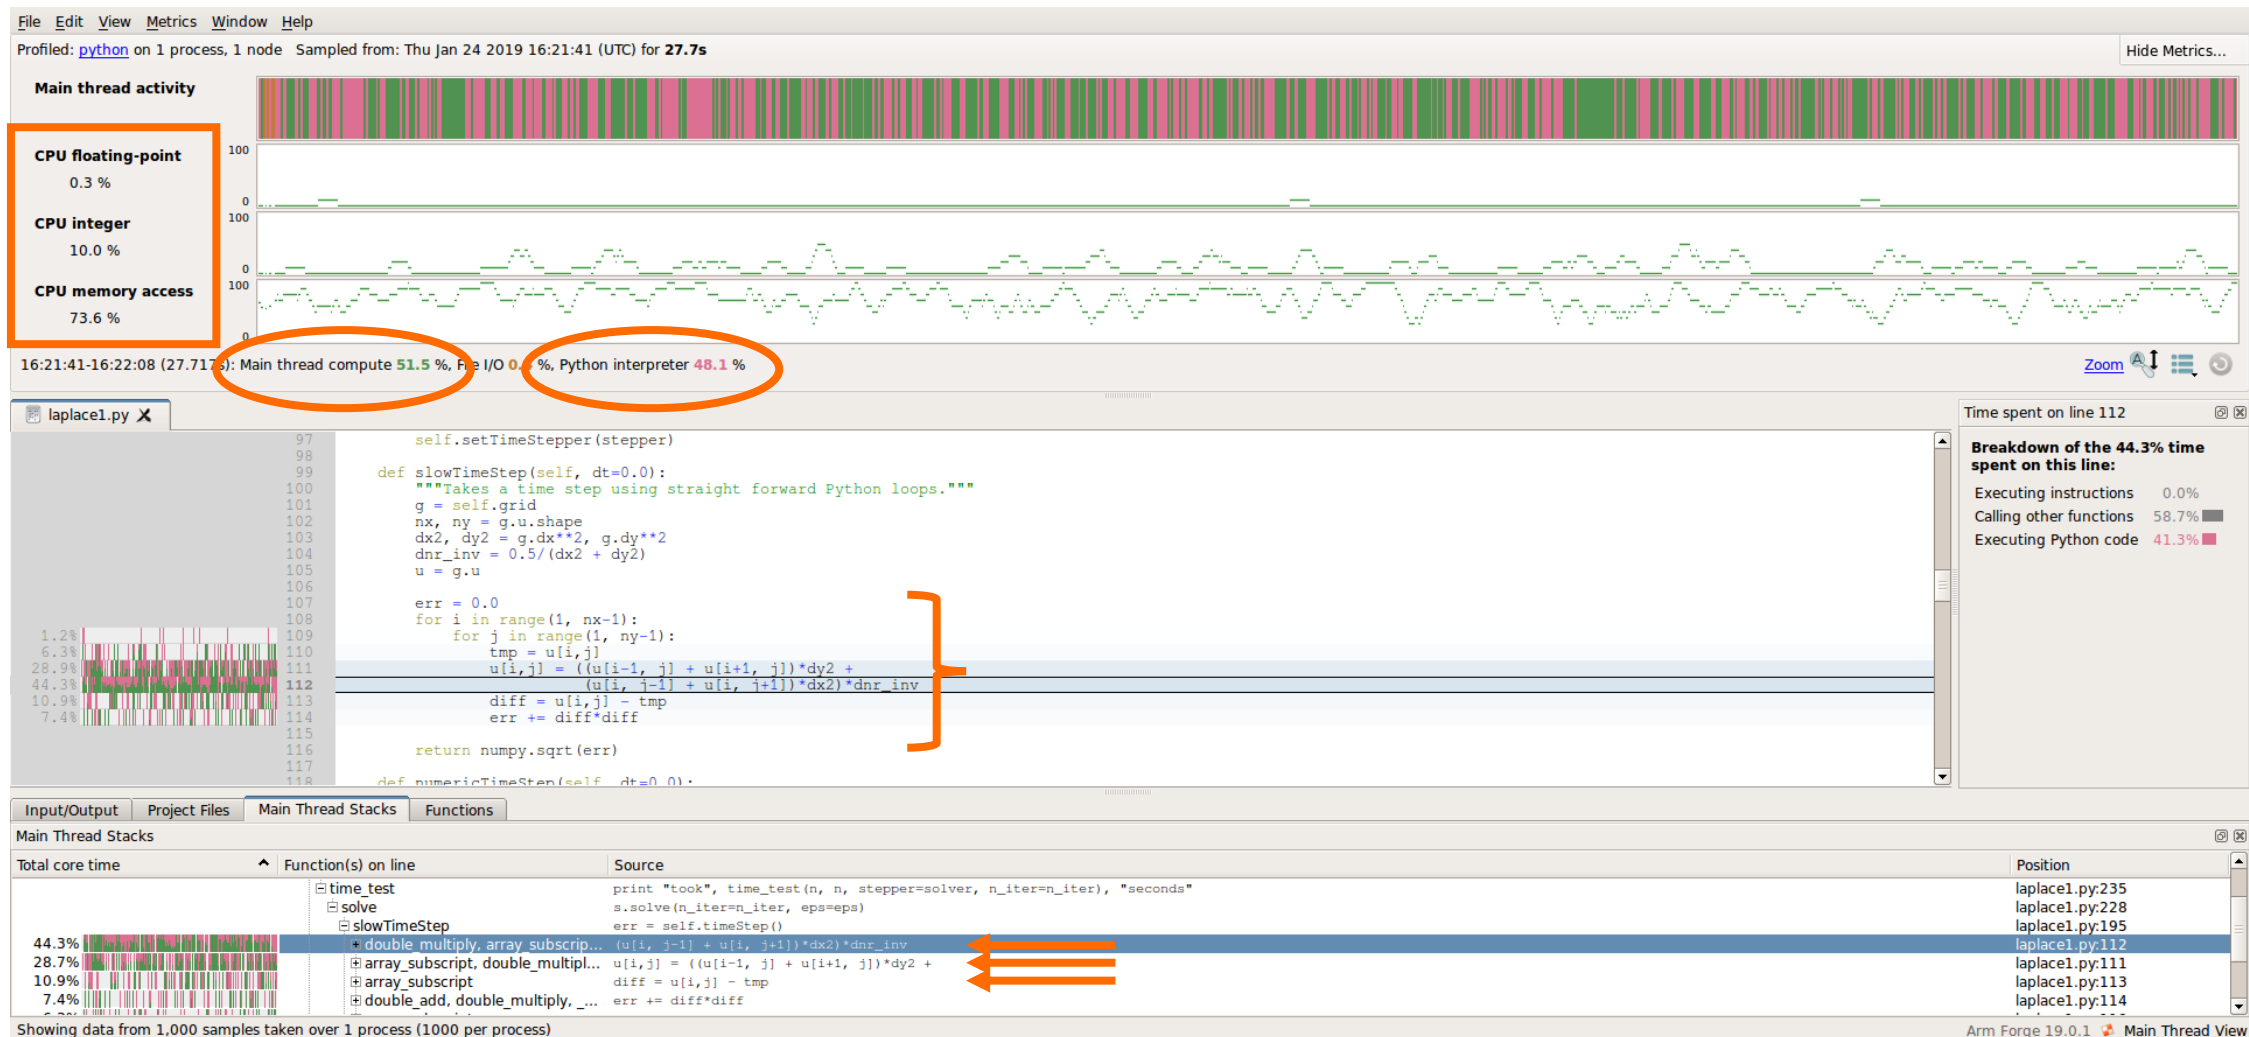
Arm MAP: usage

- Launch command
 - \$ **python** ./laplace1.py slow 100 100
- Profiling command
 - \$ **map --profile python** ./laplace1.py slow 100 100
 - --profile: non-interactive mode
 - --output: name of output file
- Display profiling results
 - \$ **map** laplace1.map

Laplace1.py

```
[...]
err = 0.0
for i in range(1, nx-1):
    for j in range(1, ny-1):
        tmp = u[i,j]
        u[i,j] = ((u[i-1, j] + u[i+1, j])*dy2 +
                  (u[i, j-1] + u[i, j+1])*dx2)*dnr_inv
        diff = u[i,j] - tmp
        err += diff*diff
return numpy.sqrt(err)
[...]
```


Naïve Python loop




Optimizing computation on NumPy arrays

Naïve Python loop

```
err = 0.0
for i in range(1, nx-1):
    for j in range(1, ny-1):
        tmp = u[i,j]
        u[i,j] = ((u[i-1, j] + u[i+1, j])*dy2 +
                  (u[i, j-1] + u[i, j+1])*dx2)*dnr_inv
        diff = u[i,j] - tmp
        err += diff*diff
return numpy.sqrt(err)
```

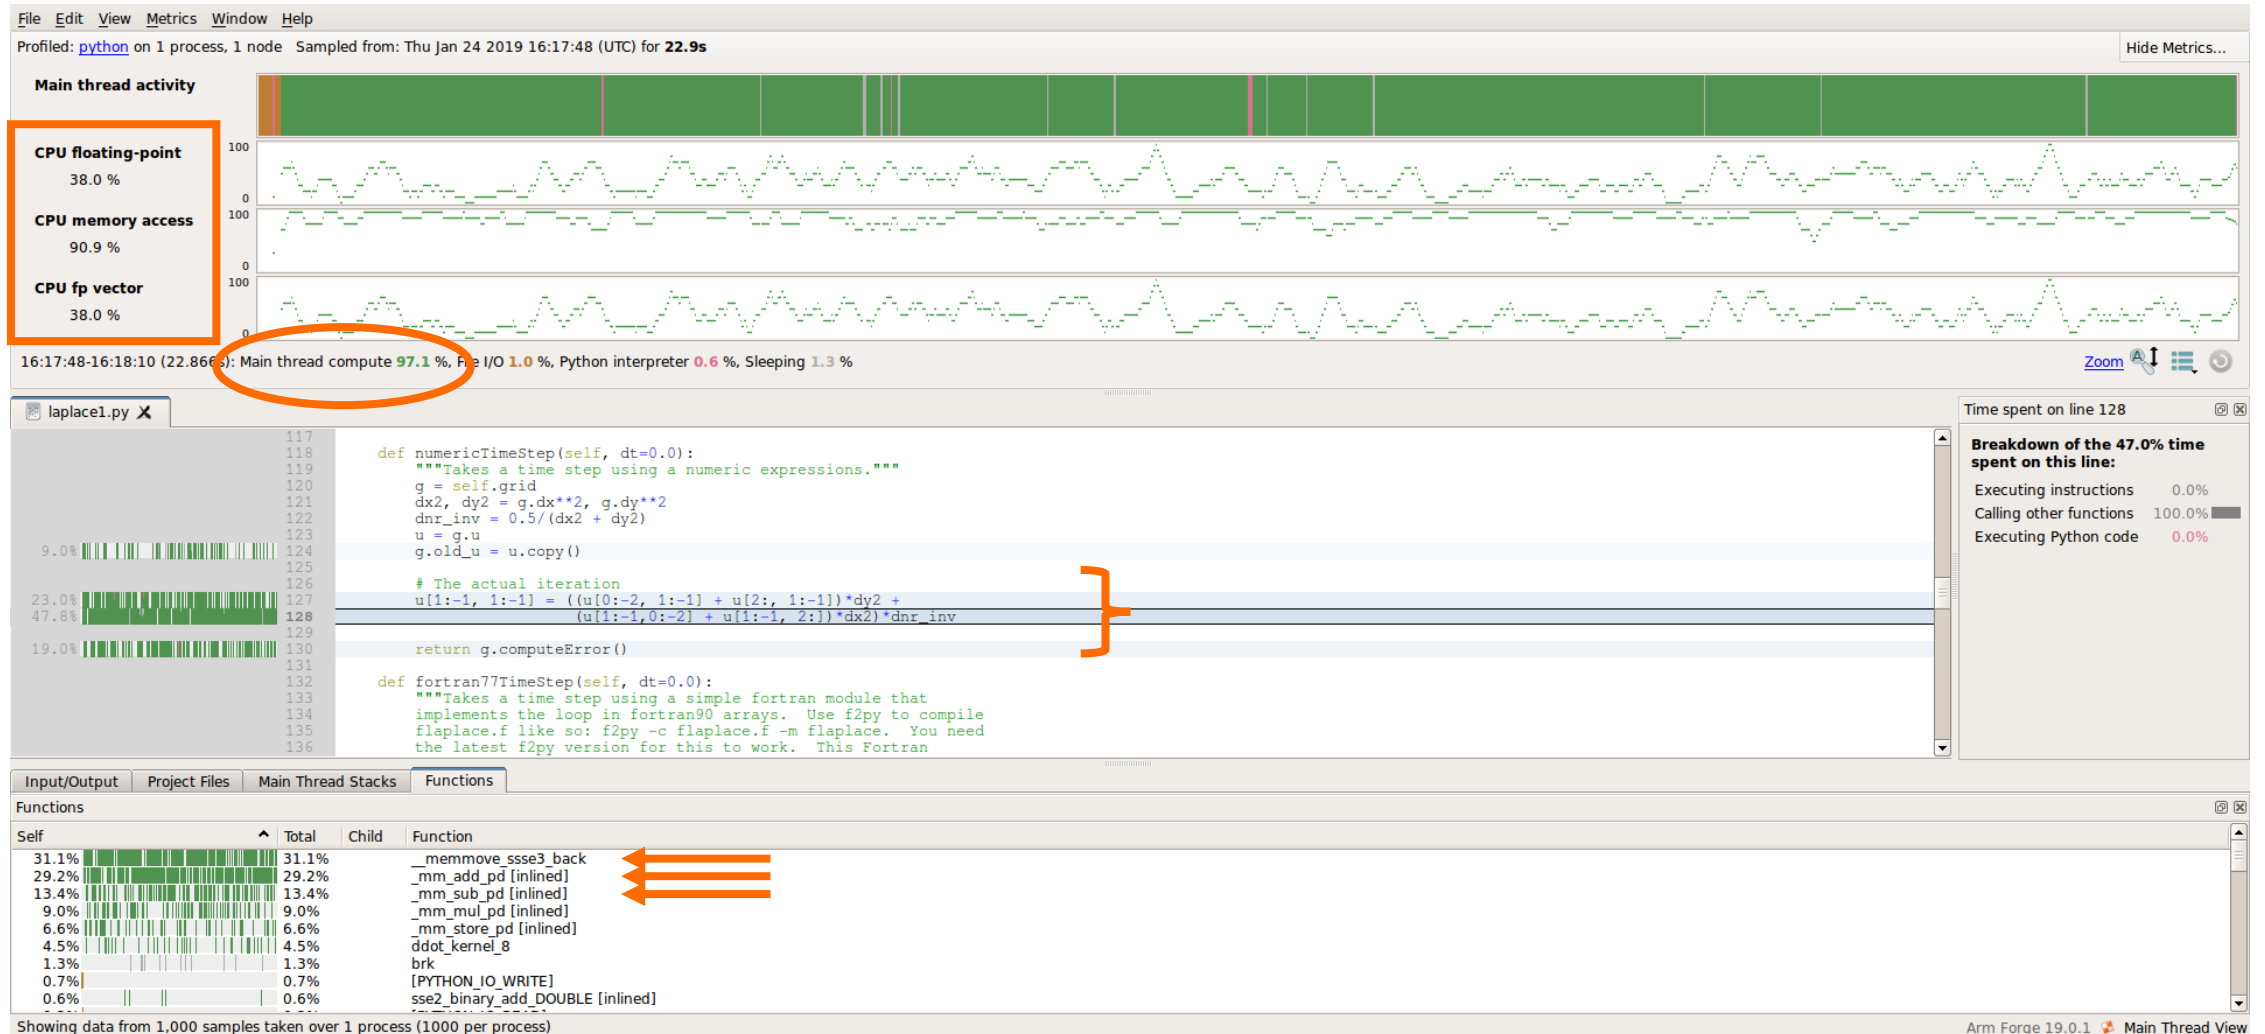
NumPy loop



```
u[1:-1, 1:-1] =
    ((u[0:-2, 1:-1] + u[2:, 1:-1])*dy2 +
     (u[1:-1, 0:-2] + u[1:-1, 2:])*dx2)*dnr_inv

return g.computeError()
```

NumPy array notation



Use FORTRAN in Python applications

- F2PY: FORTRAN to Python interface generator
- Part of NumPy
- Compile with debugging flag for profiling
 - `$ f2py --debug -c flaplace90_arrays.f90 -m flaplace90_arrays`
 - Relies on underlying compiler: GCC, IFORT, PGI
 - Generates a *.so library imported in the Python script:
`import flaplace90_arrays`
 - `--debug`: enables debug information

Use FORTRAN in Python applications

FORTRAN loop

```
subroutine timestep(u,n,m,dx,dy,error)
[...]
!f2py intent(in) :: dx,dy
!f2py intent(in,out) :: u
!f2py intent(out) :: error
!f2py intent(hide) :: n,m
[...]

u(1:n-2, 1:m-2)=((u(0:n-3, 1:m-2) + u(2:n-1, 1:m-2))*dy2 + &
    (u(1:n-2,0:m-3) + u(1:n-2, 2:m-1))*dx2)*dnr_inv

error=sqrt(sum((u-diff)**2))
end subroutine
```

Python script

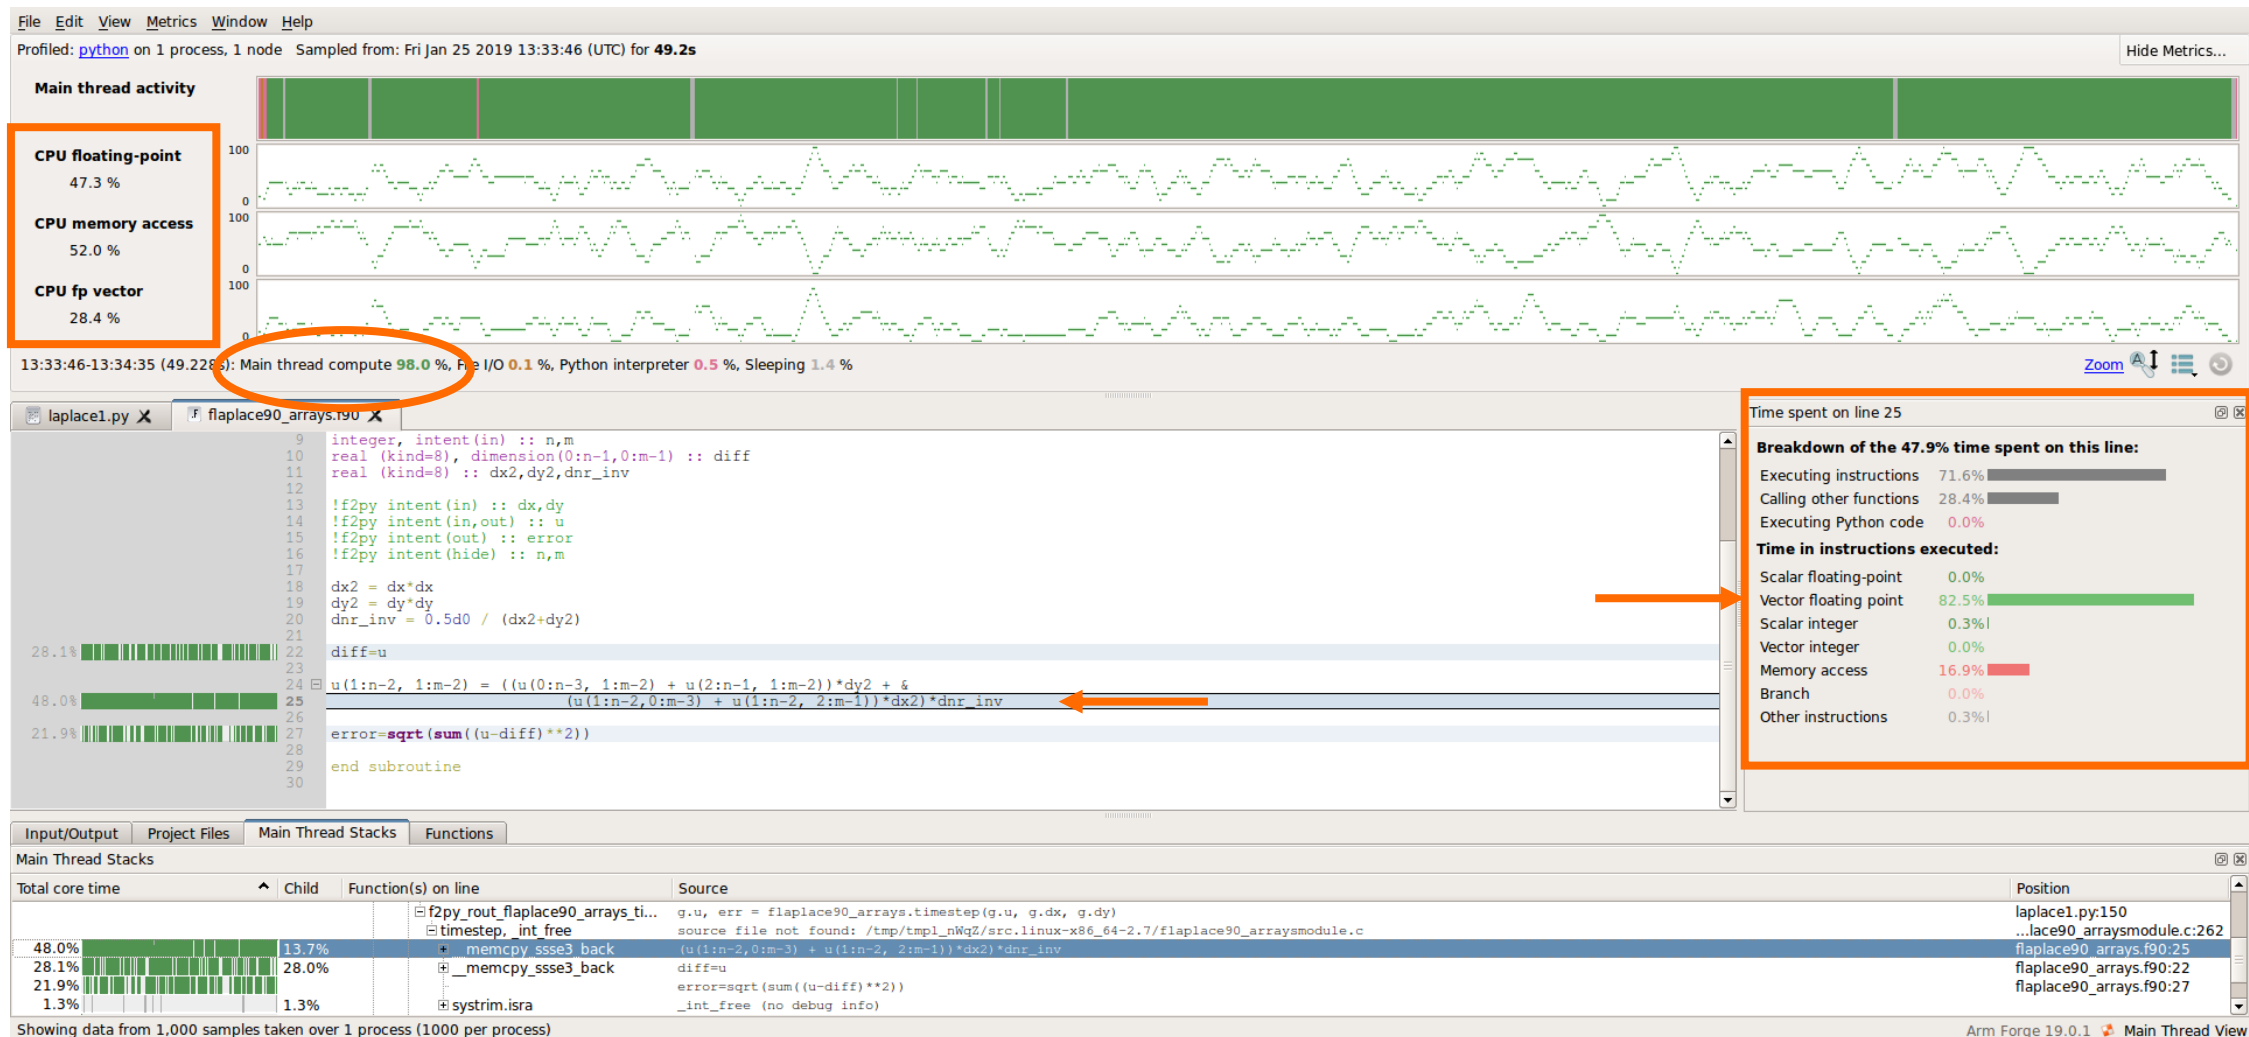
```
import flaplace90_arrays

[...]

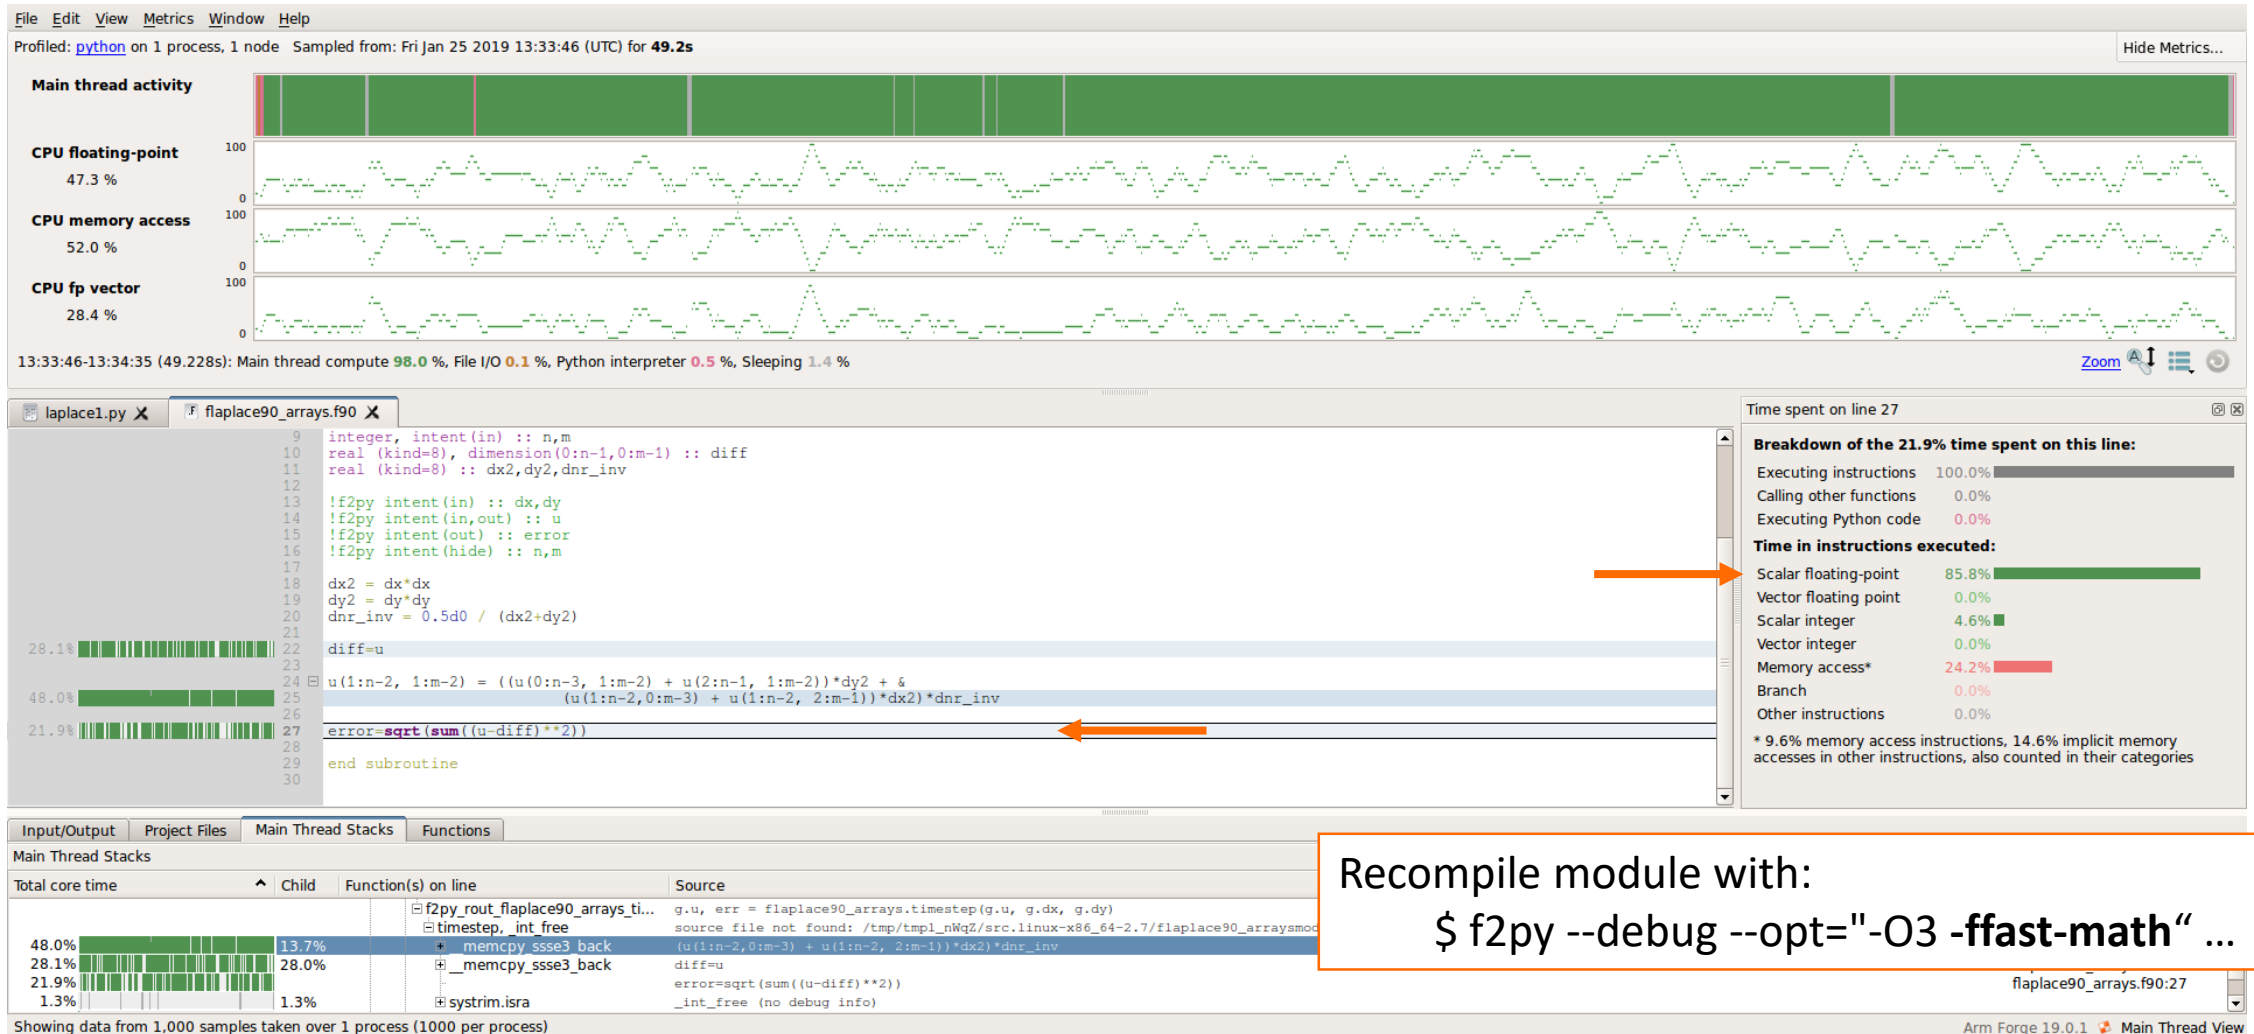
def fortran90TimeStep(self, dt=0.0):
    g = self.grid
    g.u,err =
        flaplace90_arrays.timestep(g.u, g.dx, g.dy)
    return err

[...]
```

FORTRAN code

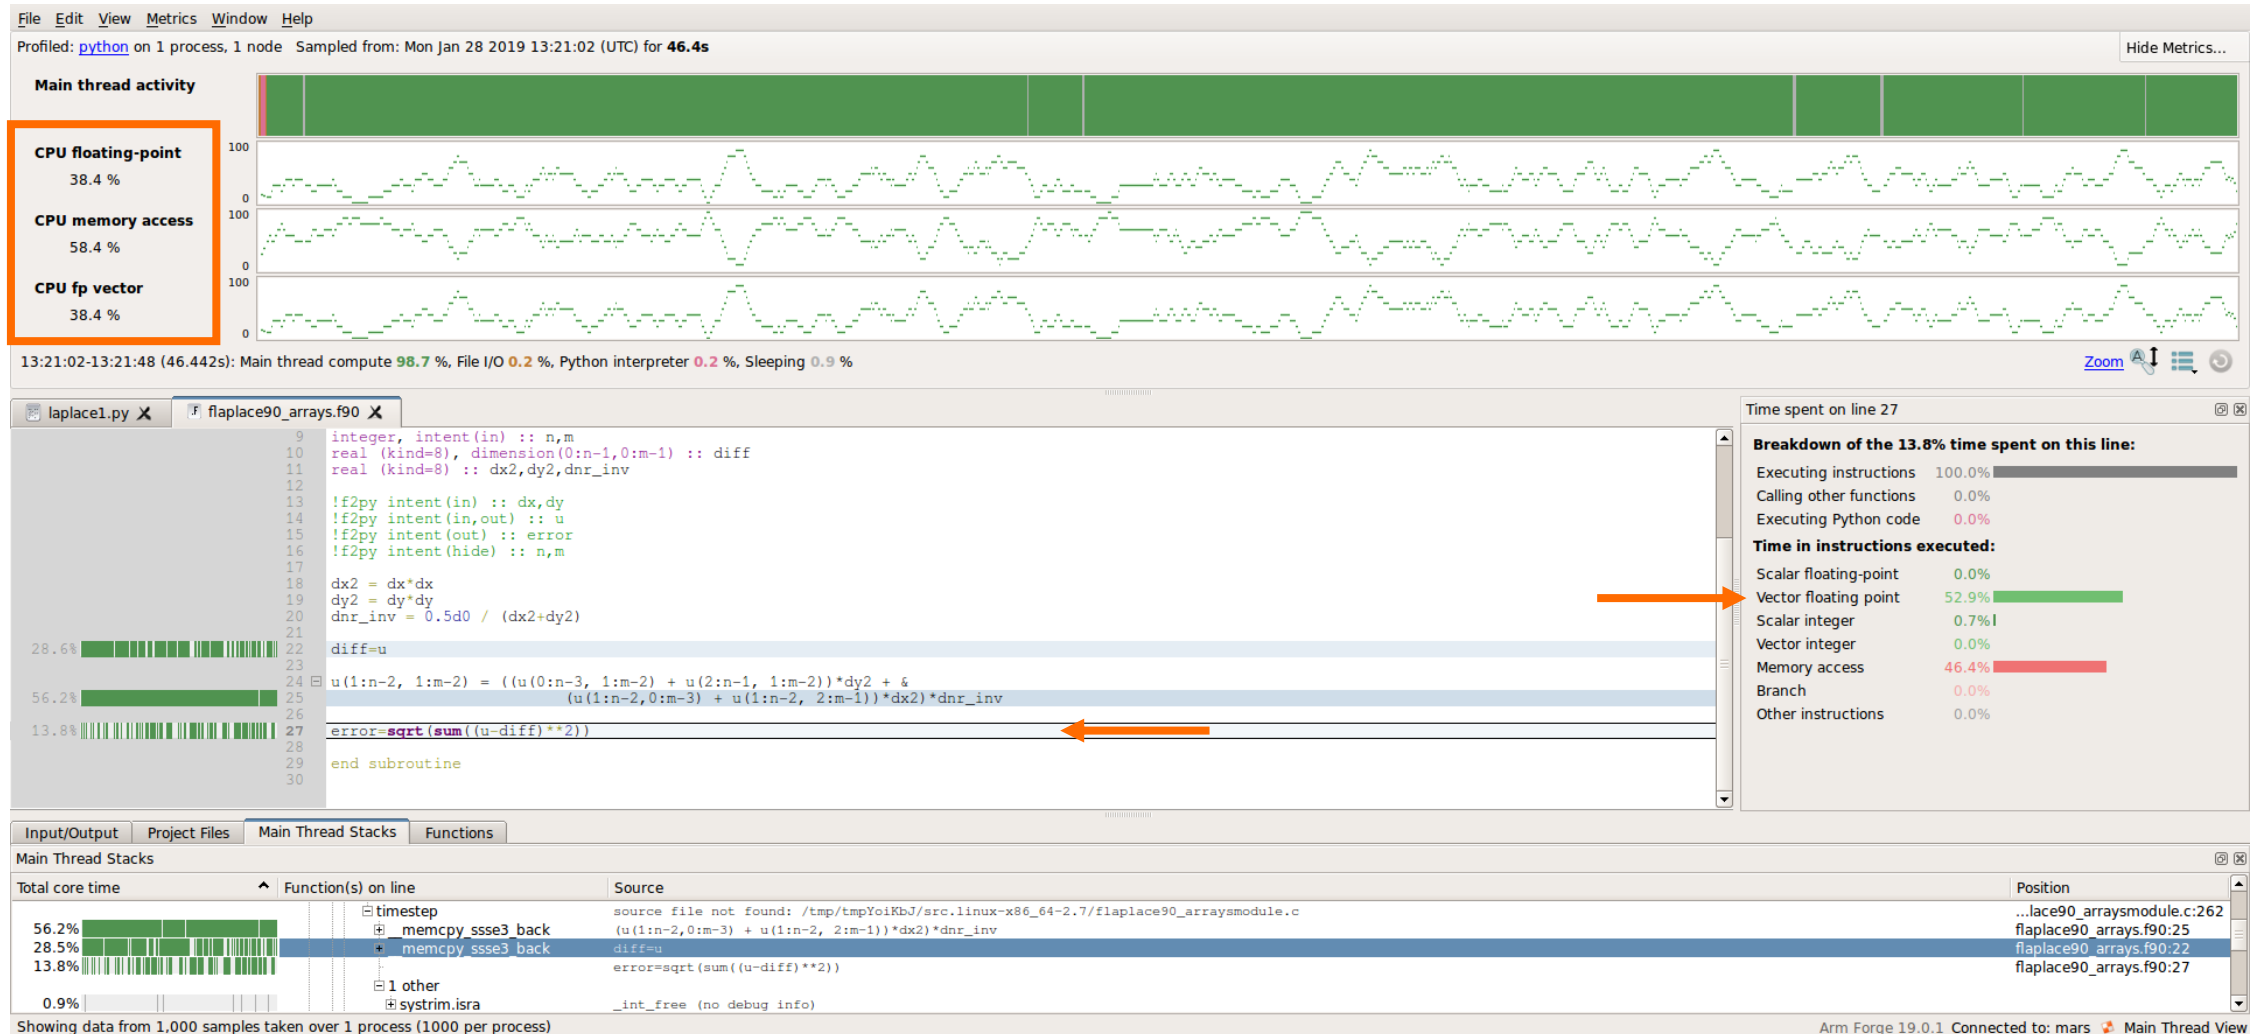


FORTRAN code



Recompile module with:
\$ f2py --debug --opt="-O3 -ffast-math" ...

FORTRAN code fast-math



Multi-processing in Python

- MPI4Py
 - Provides Python bindings of the MPI standard
 - MPI: Message Passing Interface
- Rely on existing MPI infrastructure
 - MPICC must be in path when installing module
 - MPIRUN enables to launch the application
- Profiling command
 - \$ **map --profile** mpirun -n 8 python ./mmprod.py

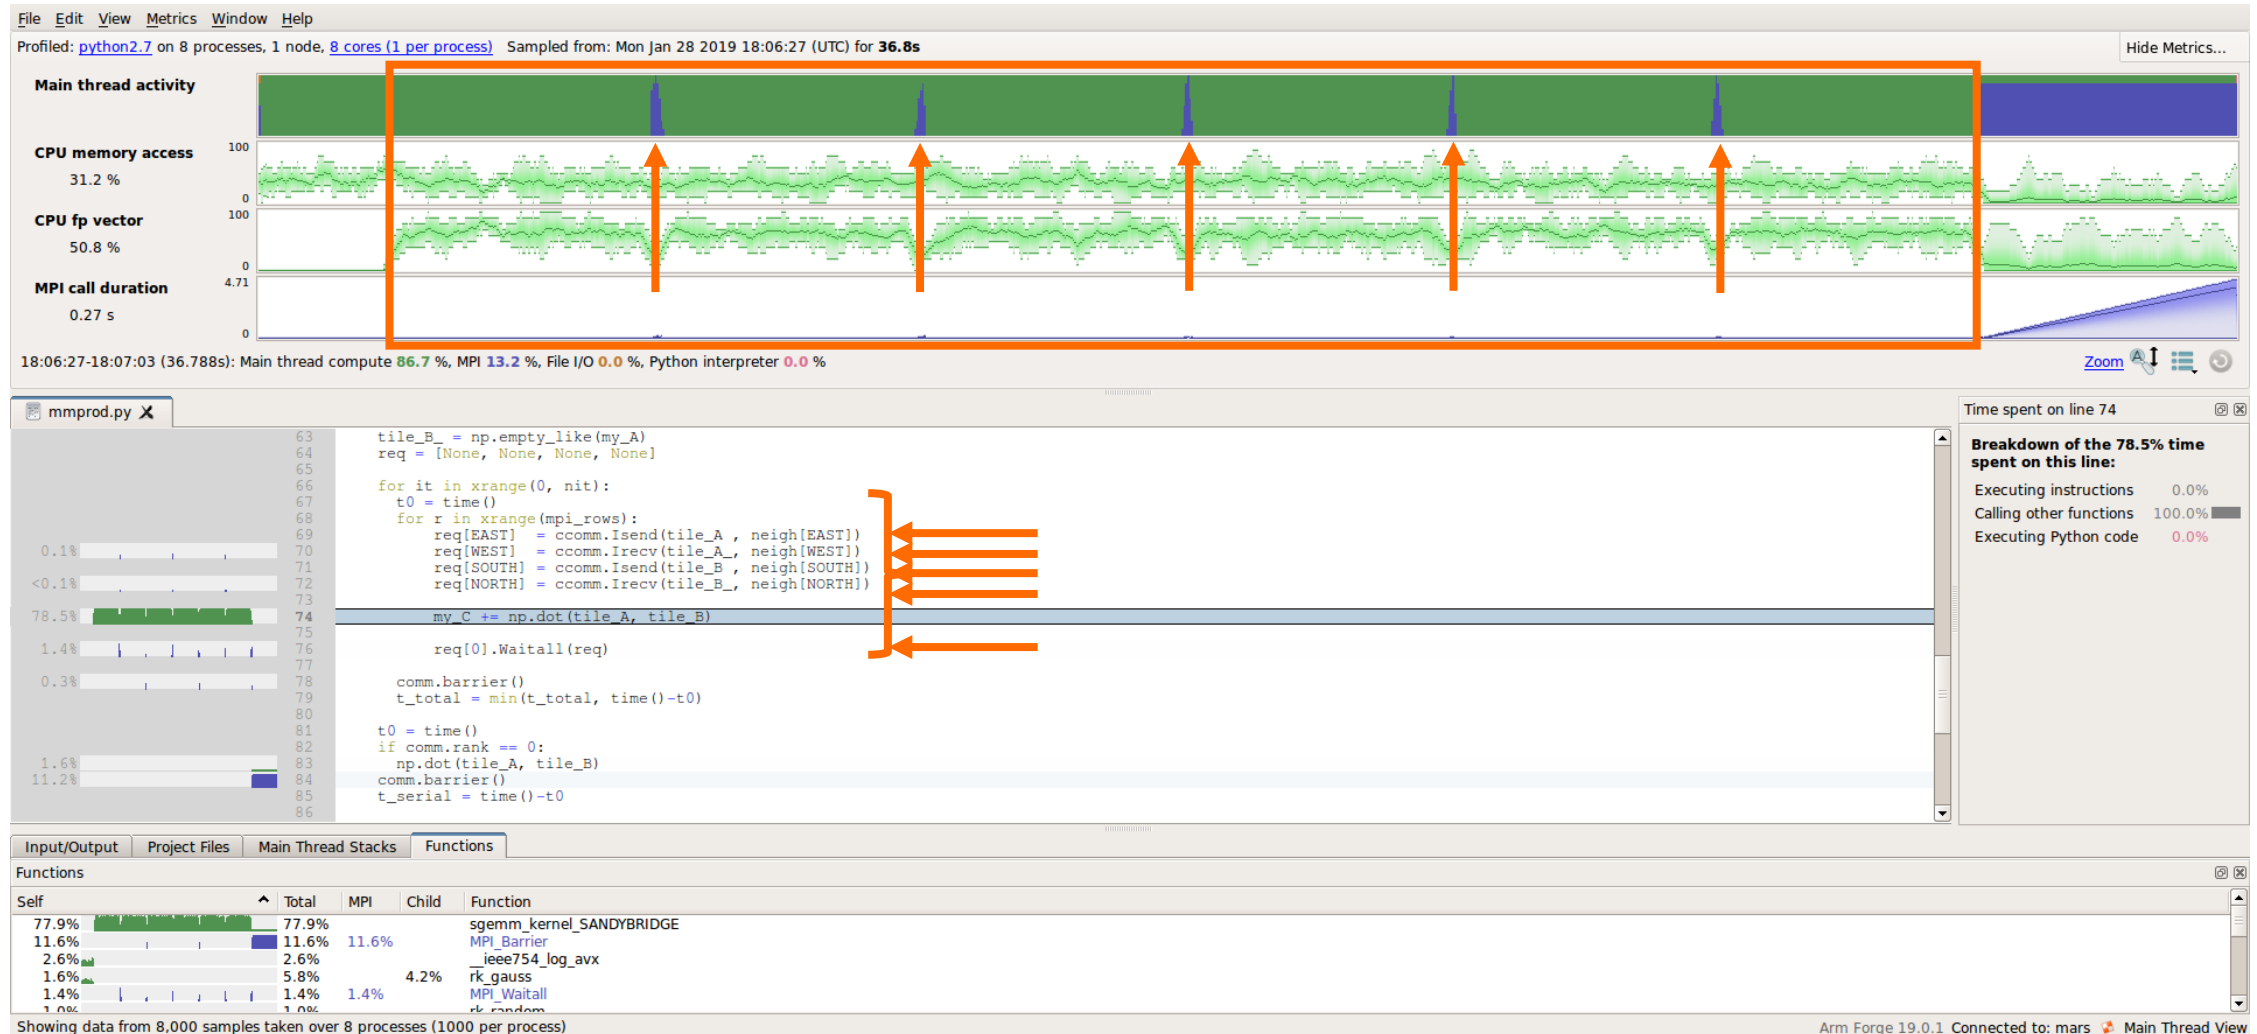
MPI4Py example

```
from mpi4py import MPI
import numpy

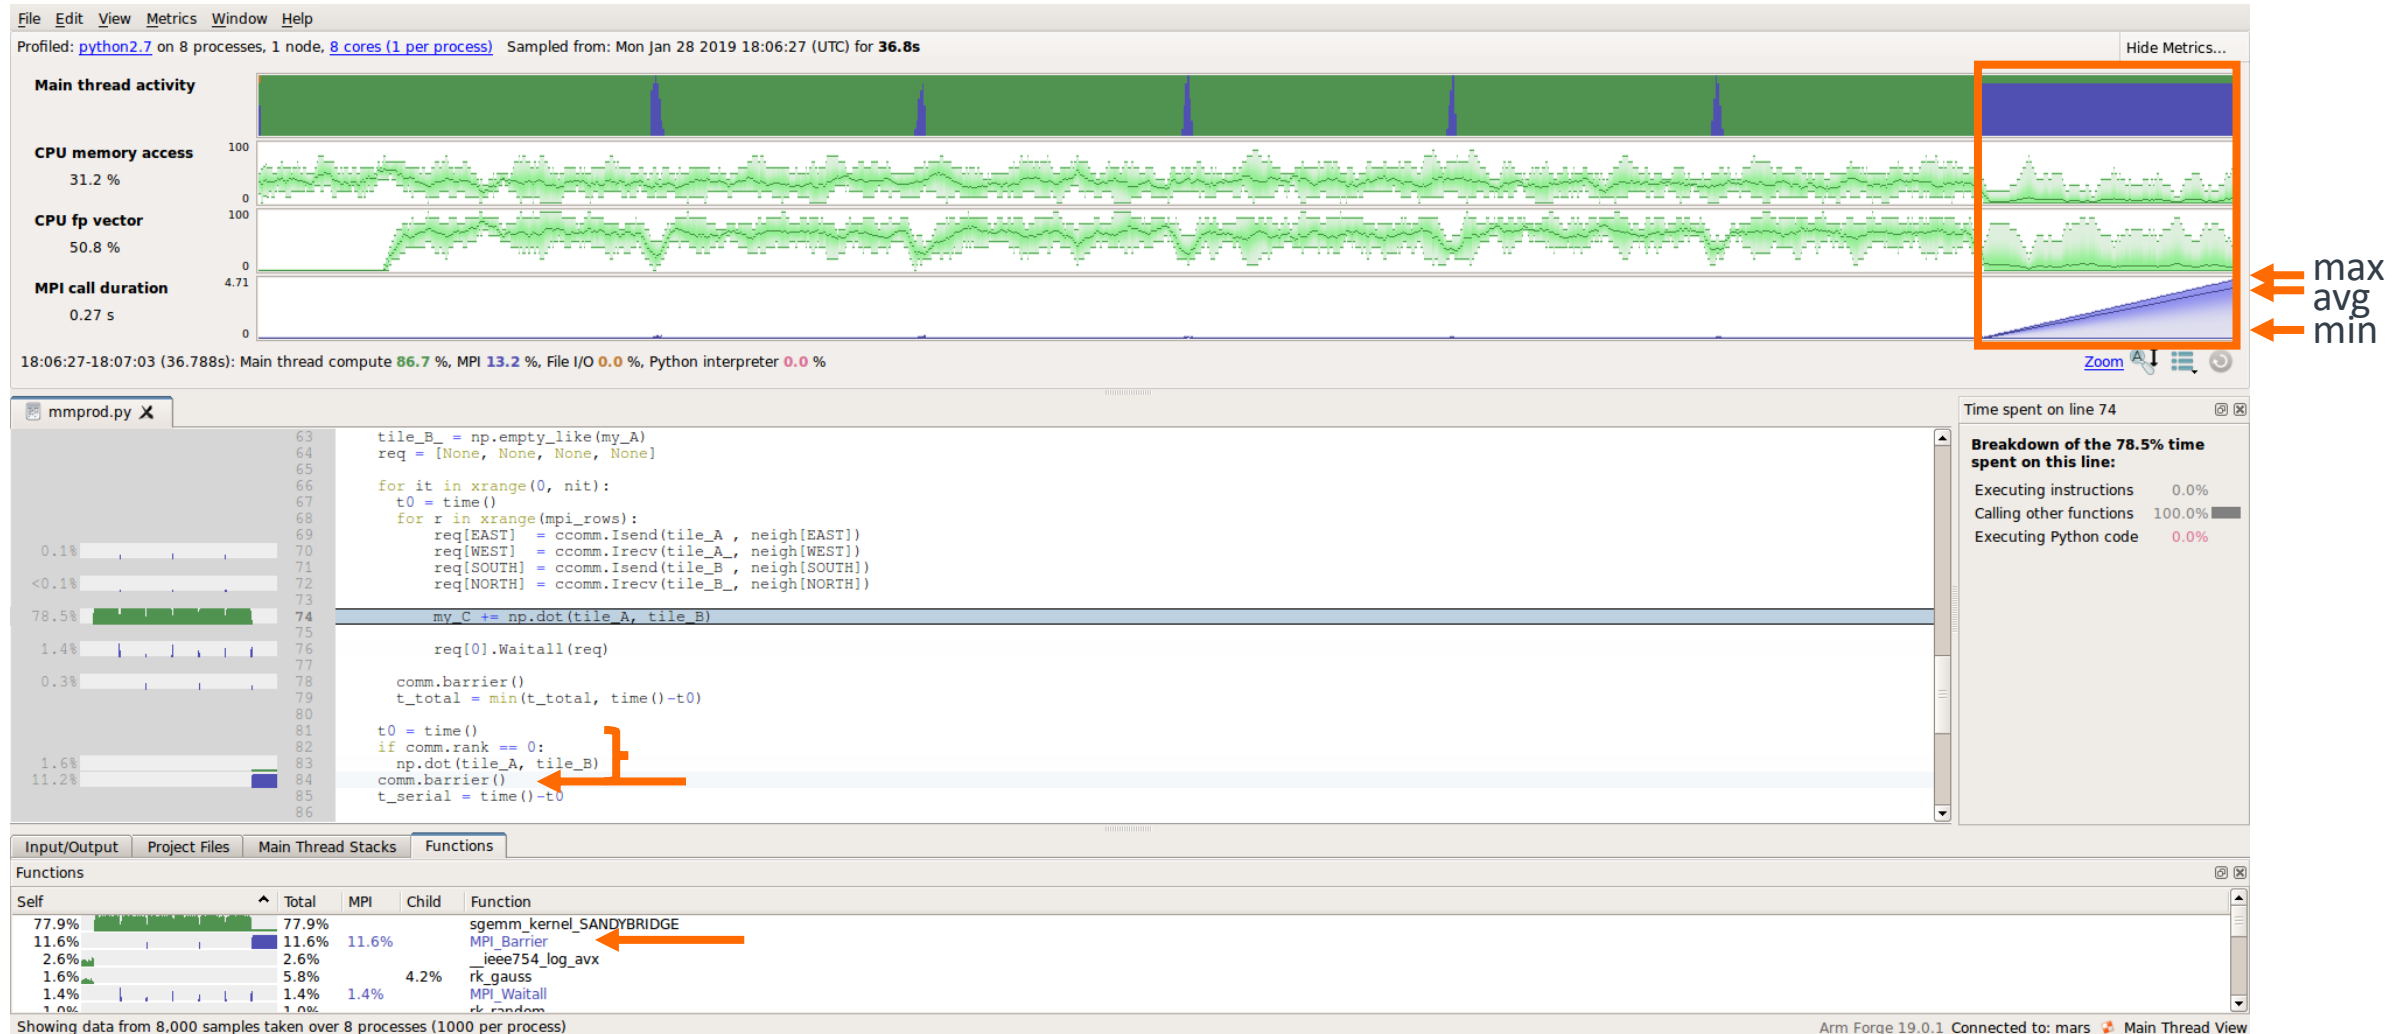
comm = MPI.COMM_WORLD
rank = comm.Get_rank()

if rank == 0:
    data = {'a': 7, 'b': 3.14}
    comm.send(data, dest=1)
elif rank == 1:
    data = comm.recv(source=0)
    print('On process 1, data is ', data)
```

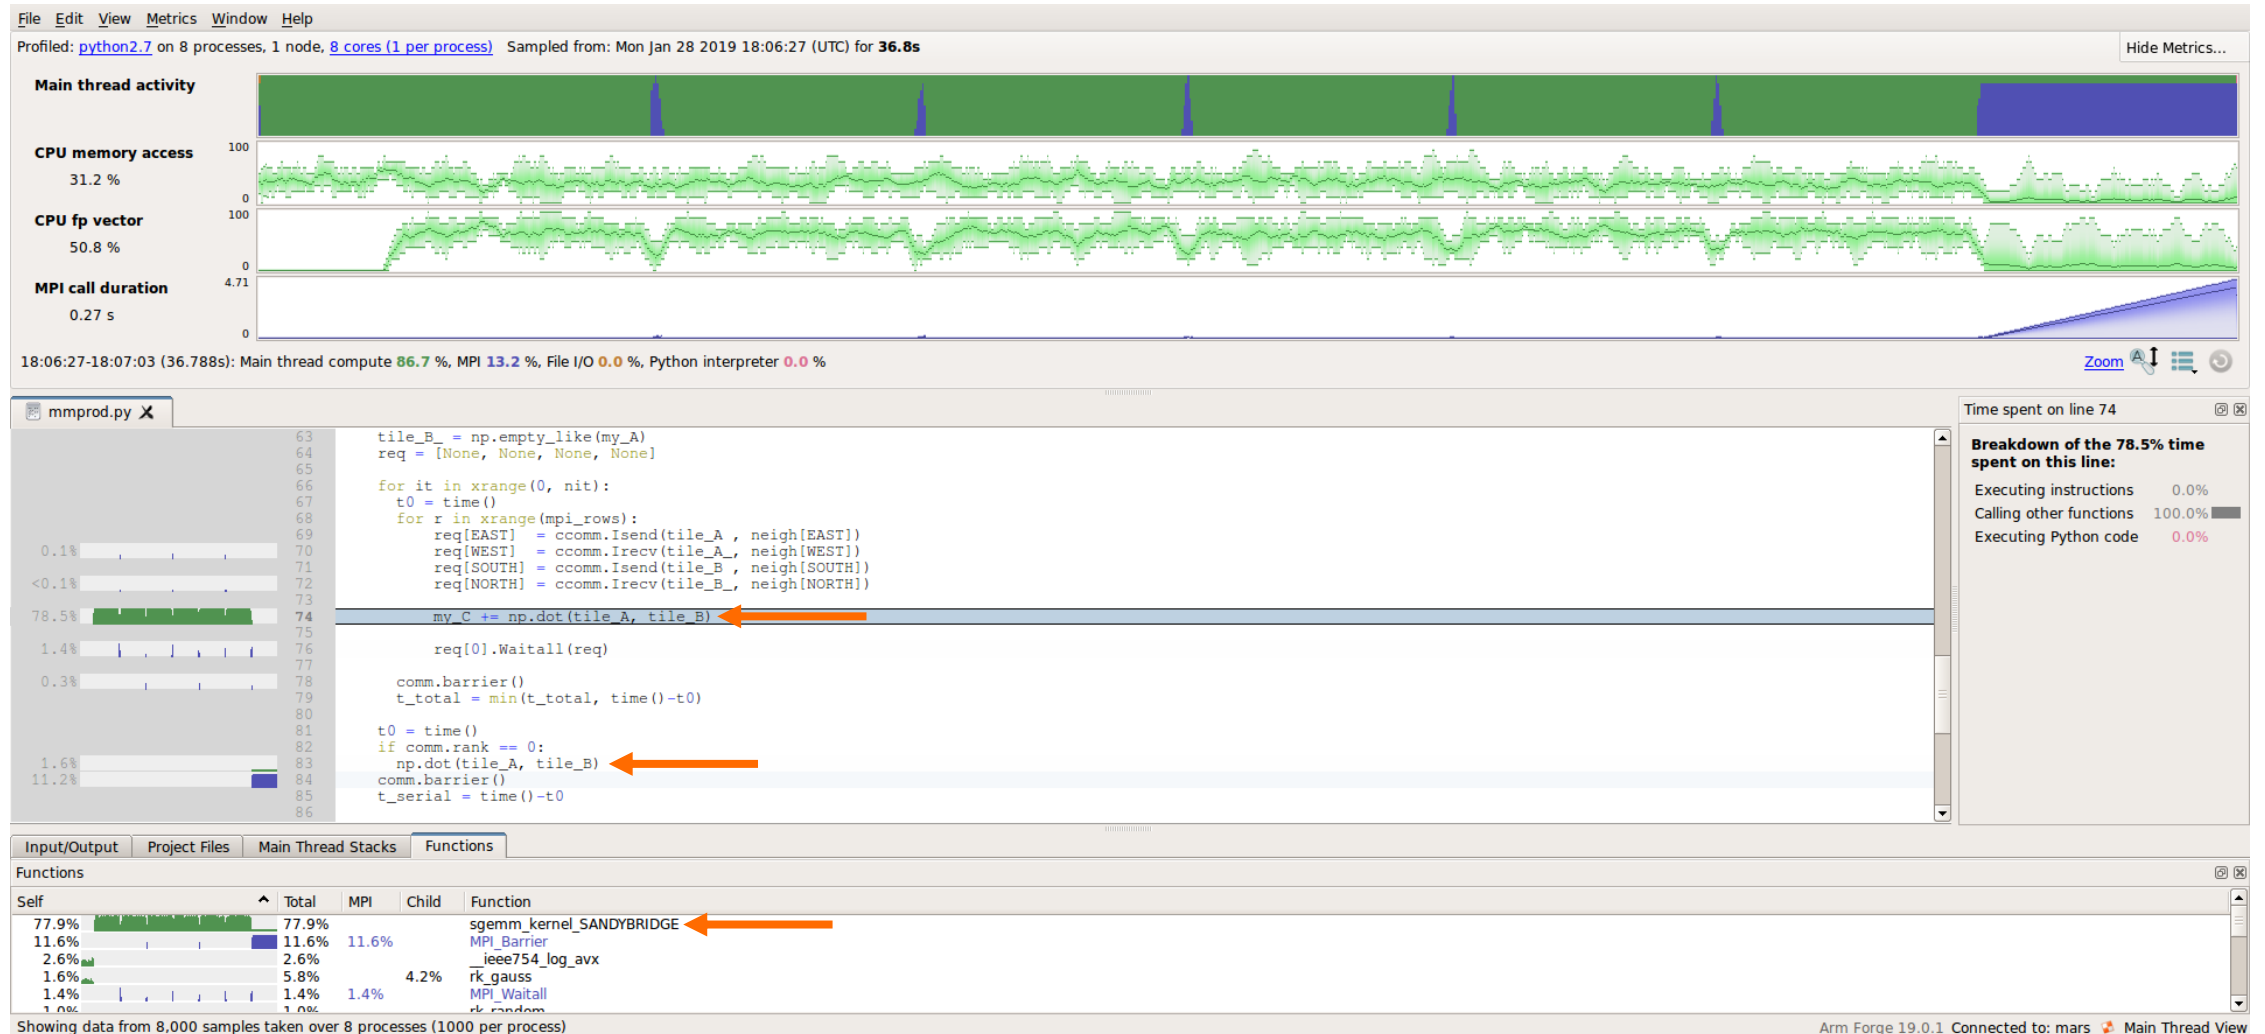
MPI-parallel matrix multiplication



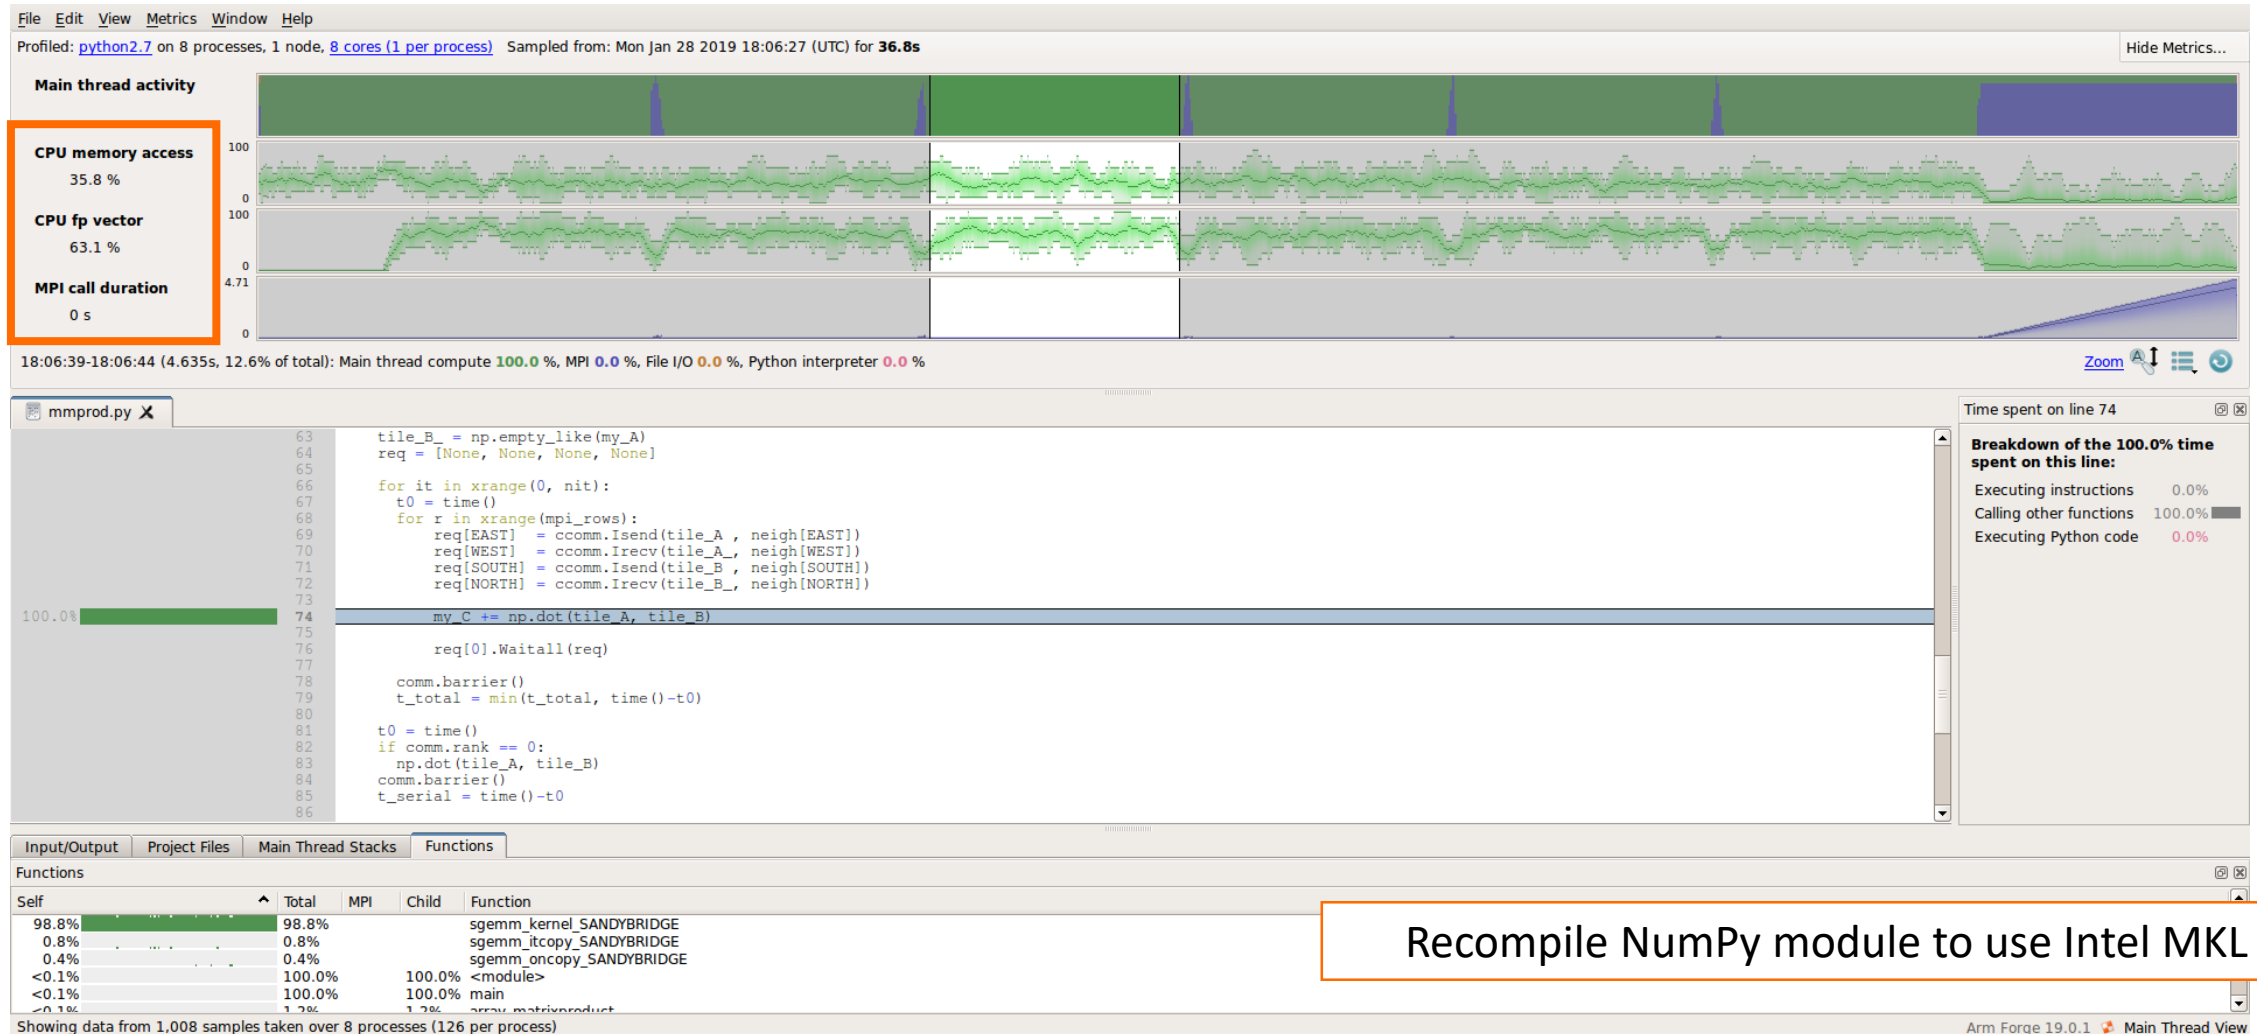
MPI-parallel matrix multiplication



MPI-parallel matrix multiplication

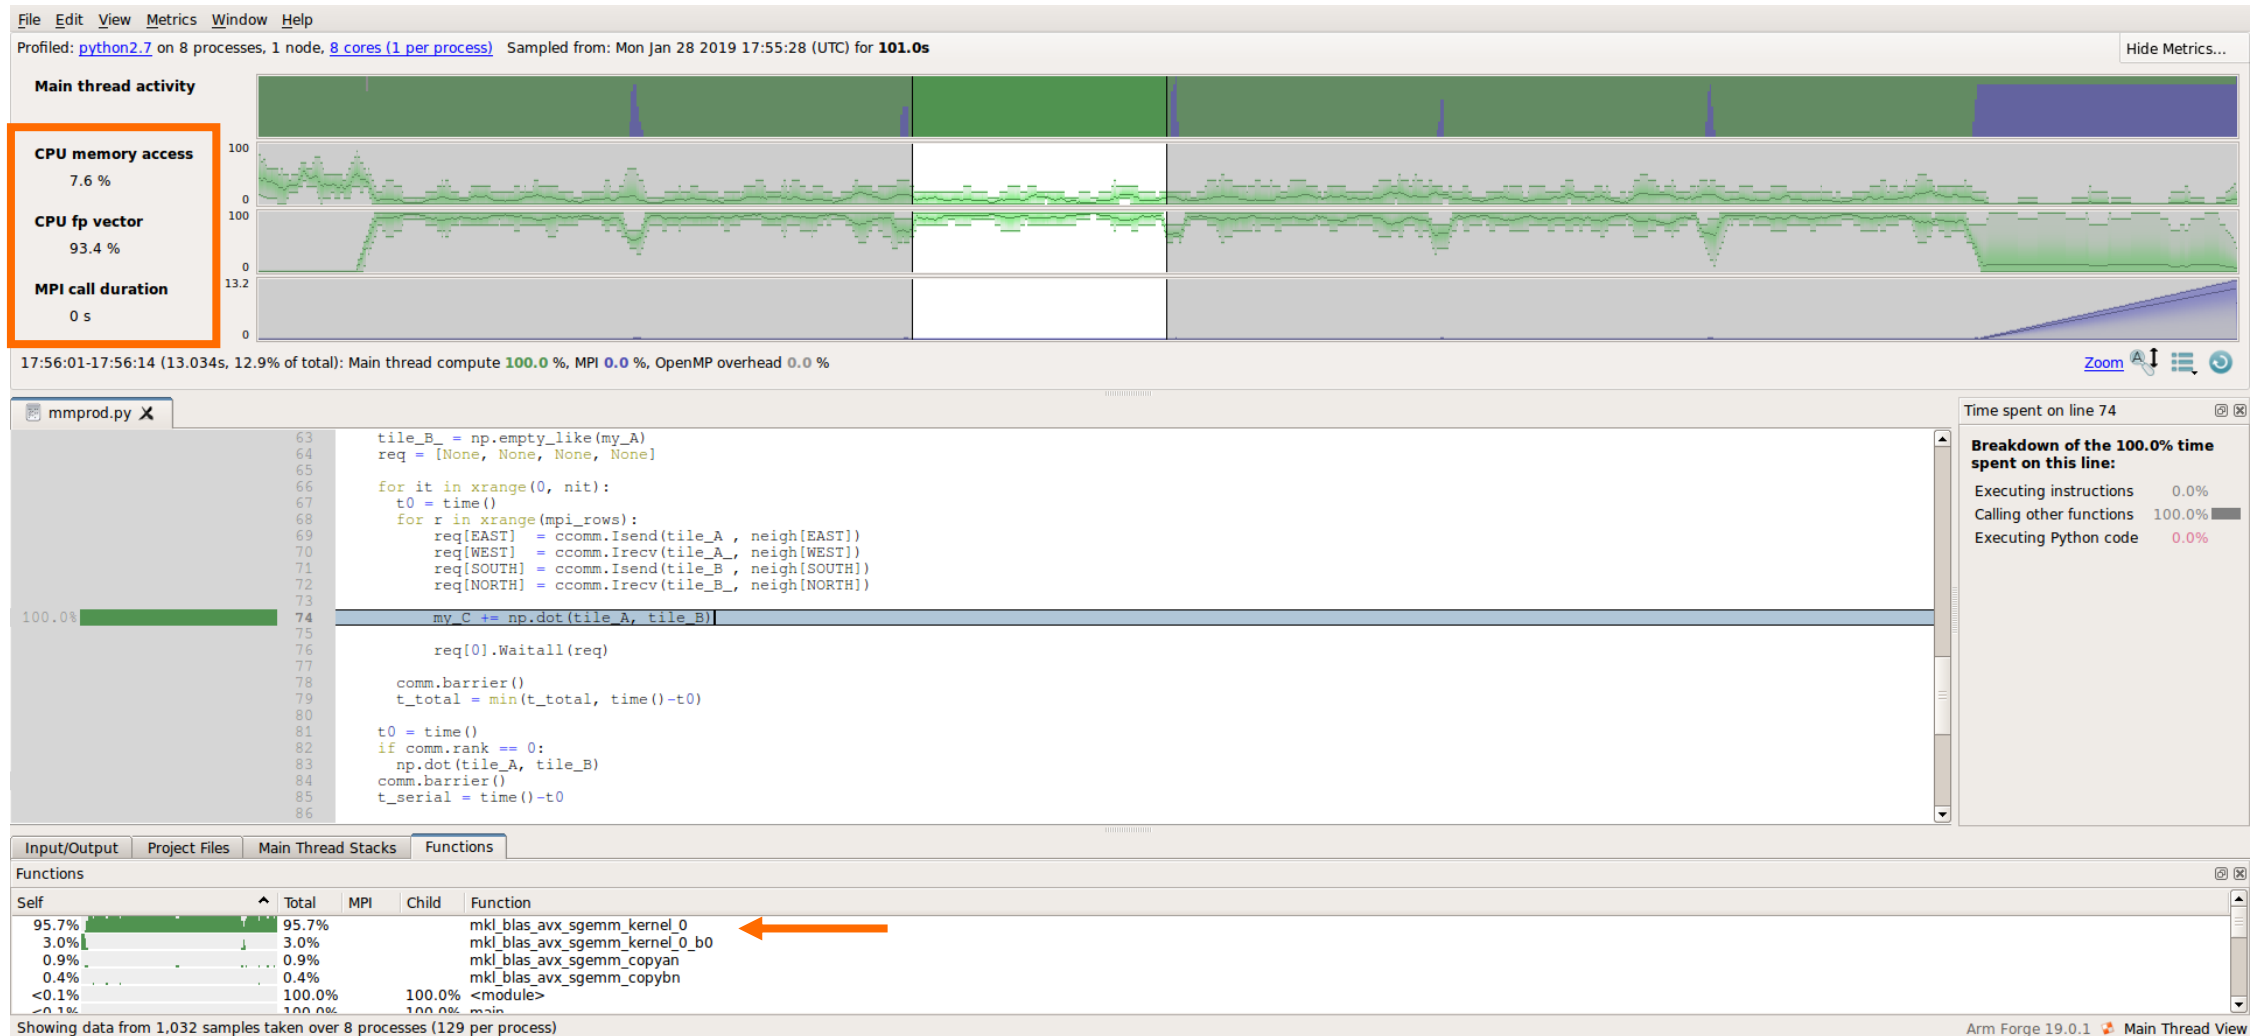


MPI-parallel matrix multiplication: OpenBLAS

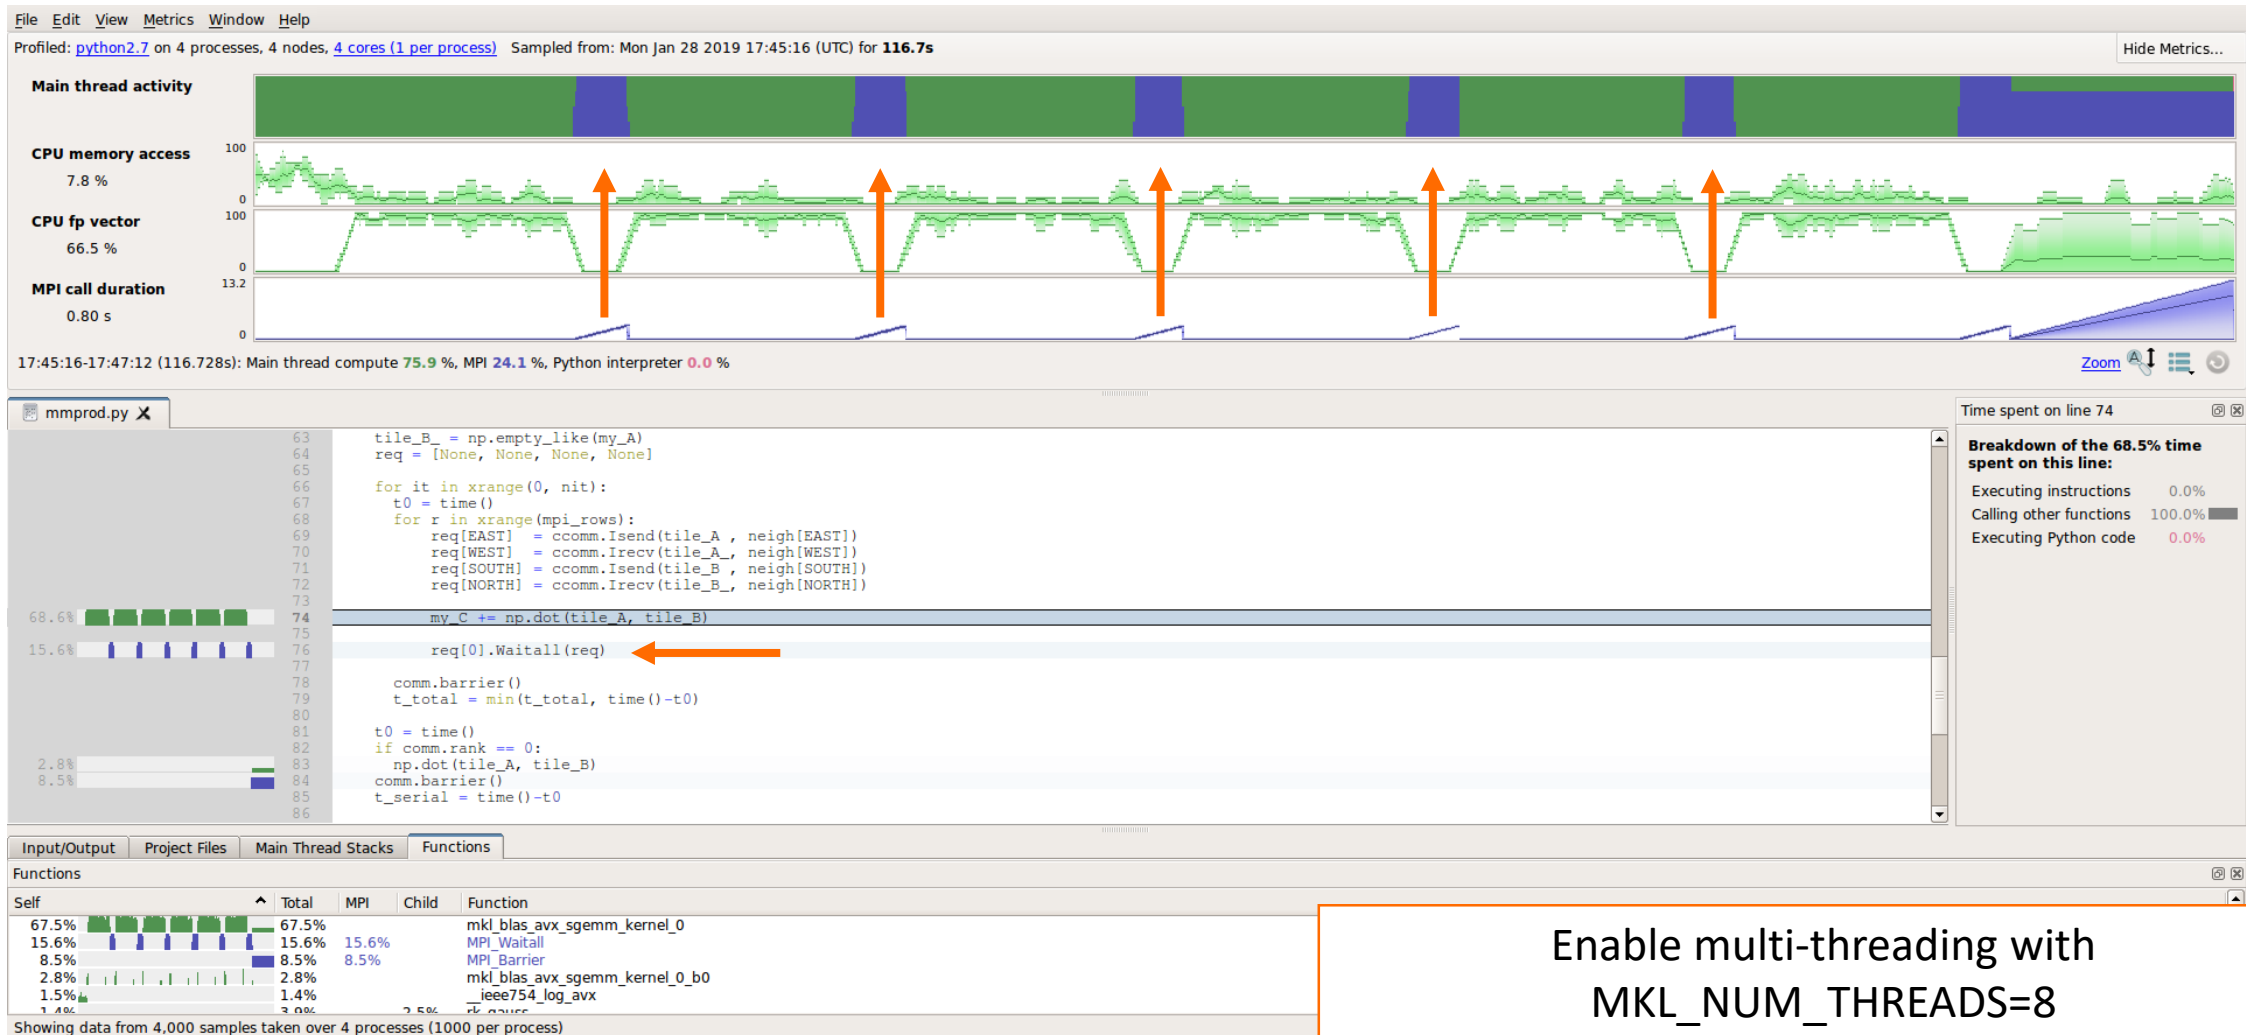


Recompile NumPy module to use Intel MKL

MPI-parallel matrix multiplication: MKL



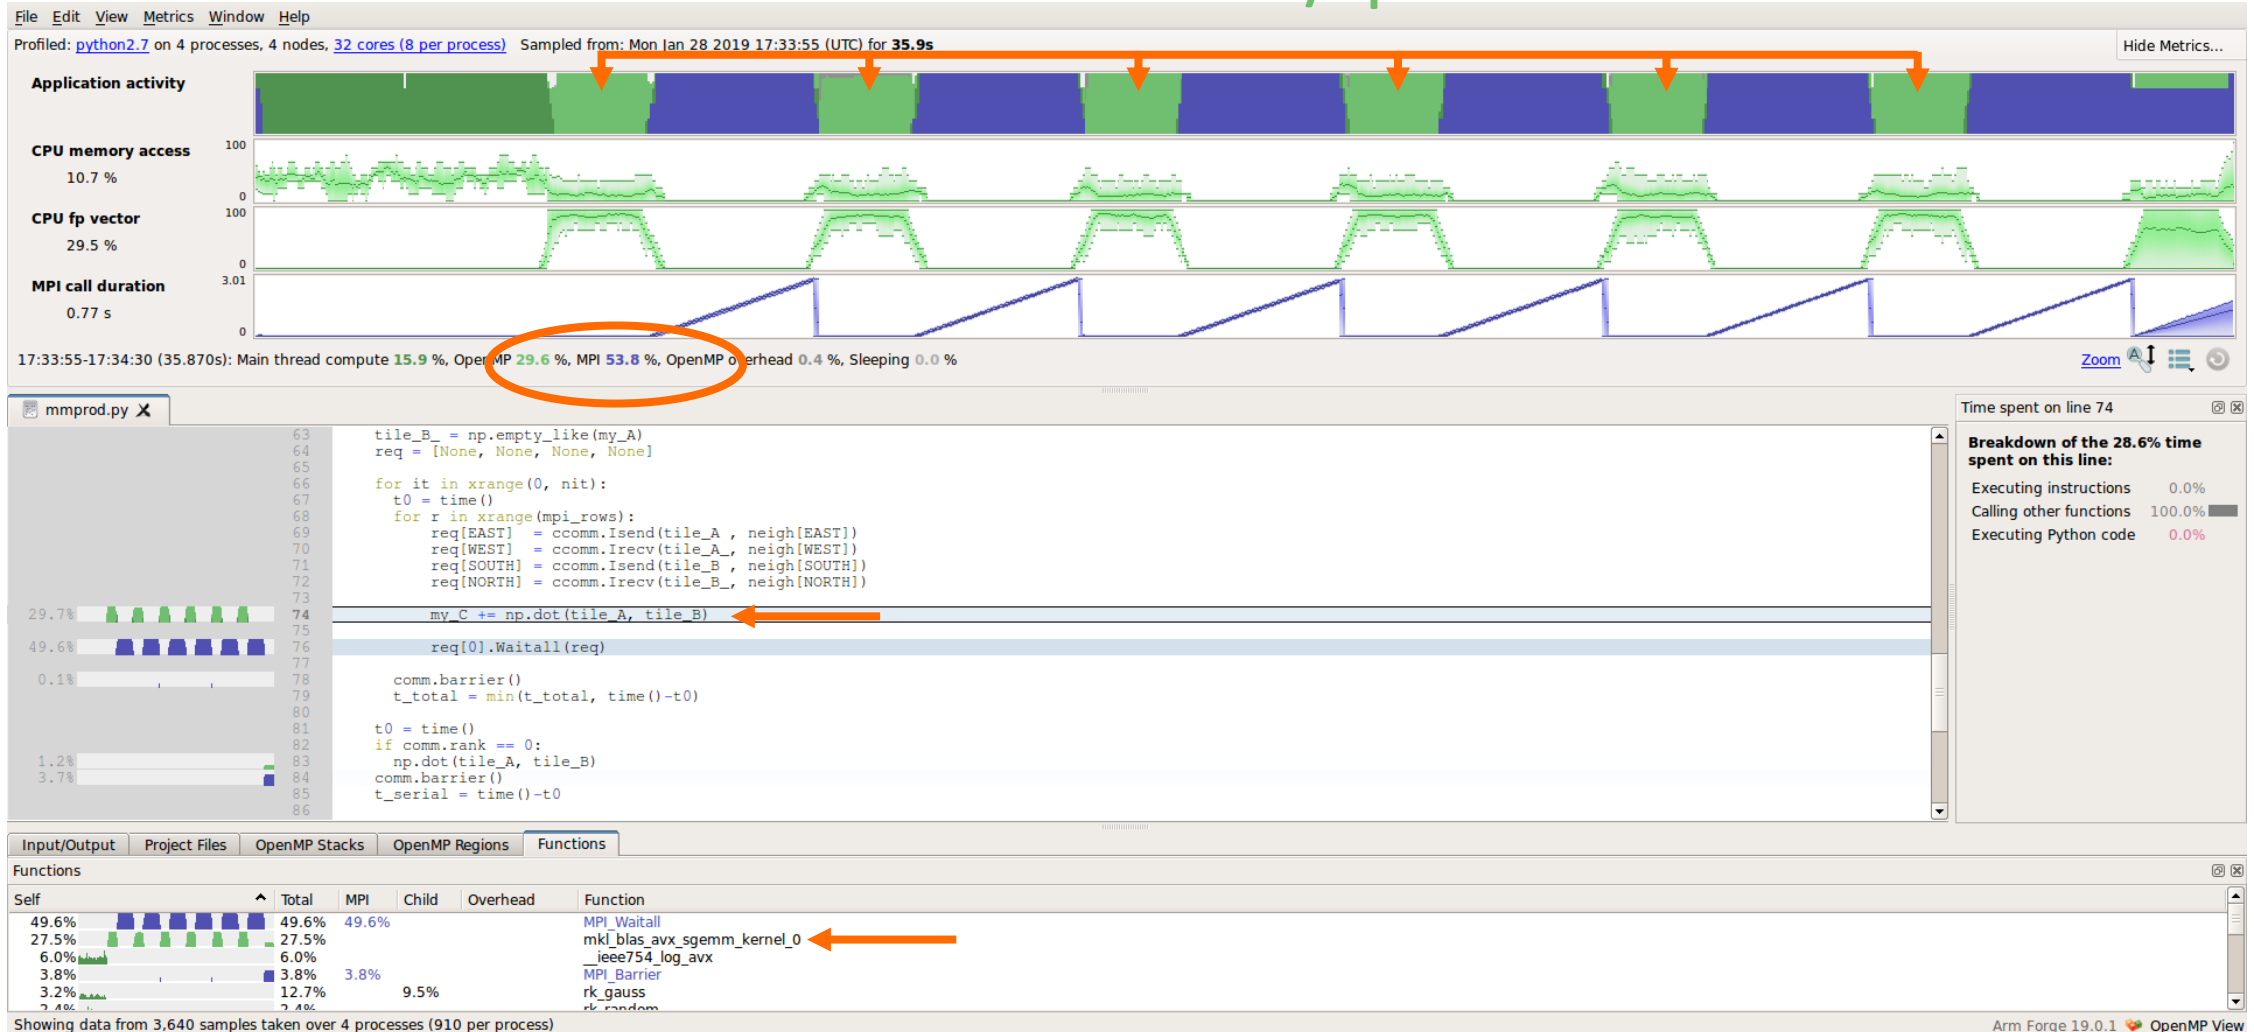
Multi-node matrix multiplication: MKL



Enable multi-threading with
MKL_NUM_THREADS=8

Hybrid OMP/MPI matrix multiplication

Multithreaded/OpenMP MKL

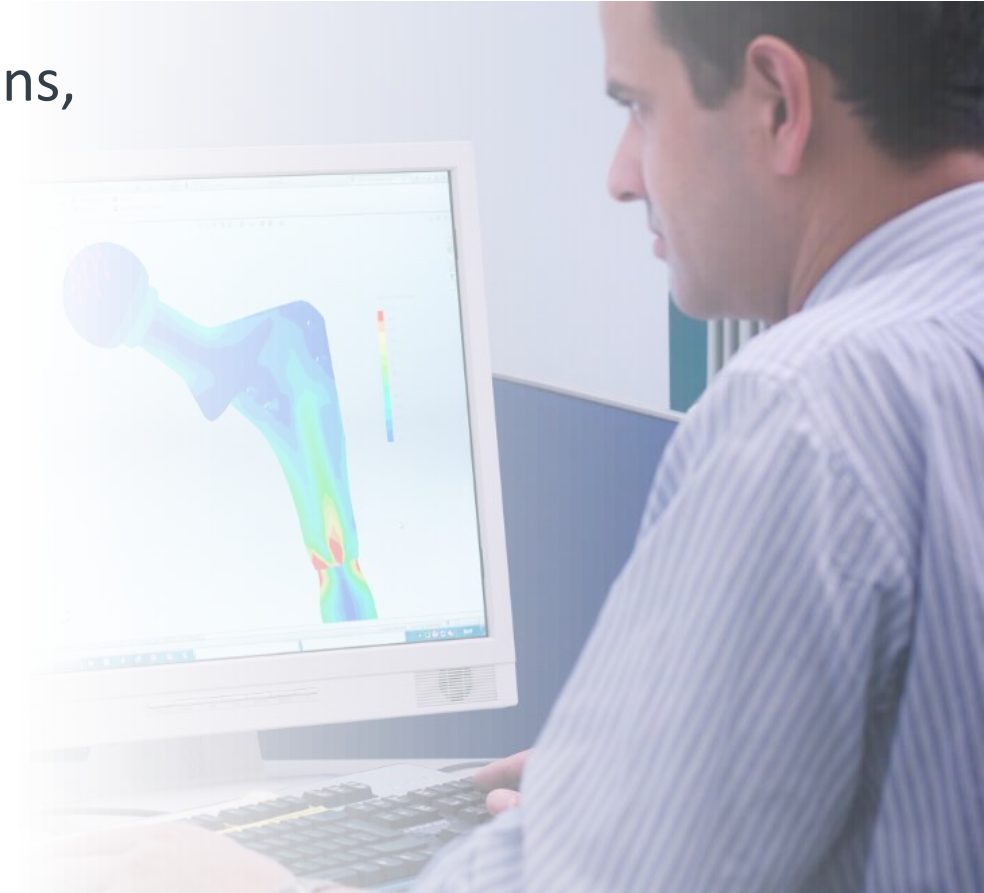


Conclusion: topics covered

- Profiling Python applications on HPC systems easily with Arm MAP
- How to analyze MAP profiling results
- How to optimize Python
 - Using NumPy array notation and minimizing time spent in interpreter
 - Integrating FORTRAN code
 - Maximizing HPC library efficiency
 - Using multi-threading

Request a trial license

- If you would like to profile your own Python applications, visit our [website](#)
- A trial licence of Arm Forge enables:
 - Multi-node and multi-threaded profiling with Arm MAP
 - On all HPC architectures
 - On C/C++, FORTRAN and Python codes
 - Debugging with Arm DDT
 - Access to our support
- Available resources online for your trial



Thank You

Danke

Merci

谢谢

ありがとう

Gracias

Kiitos

감사합니다

धन्यवाद

תודה

arm

Questions