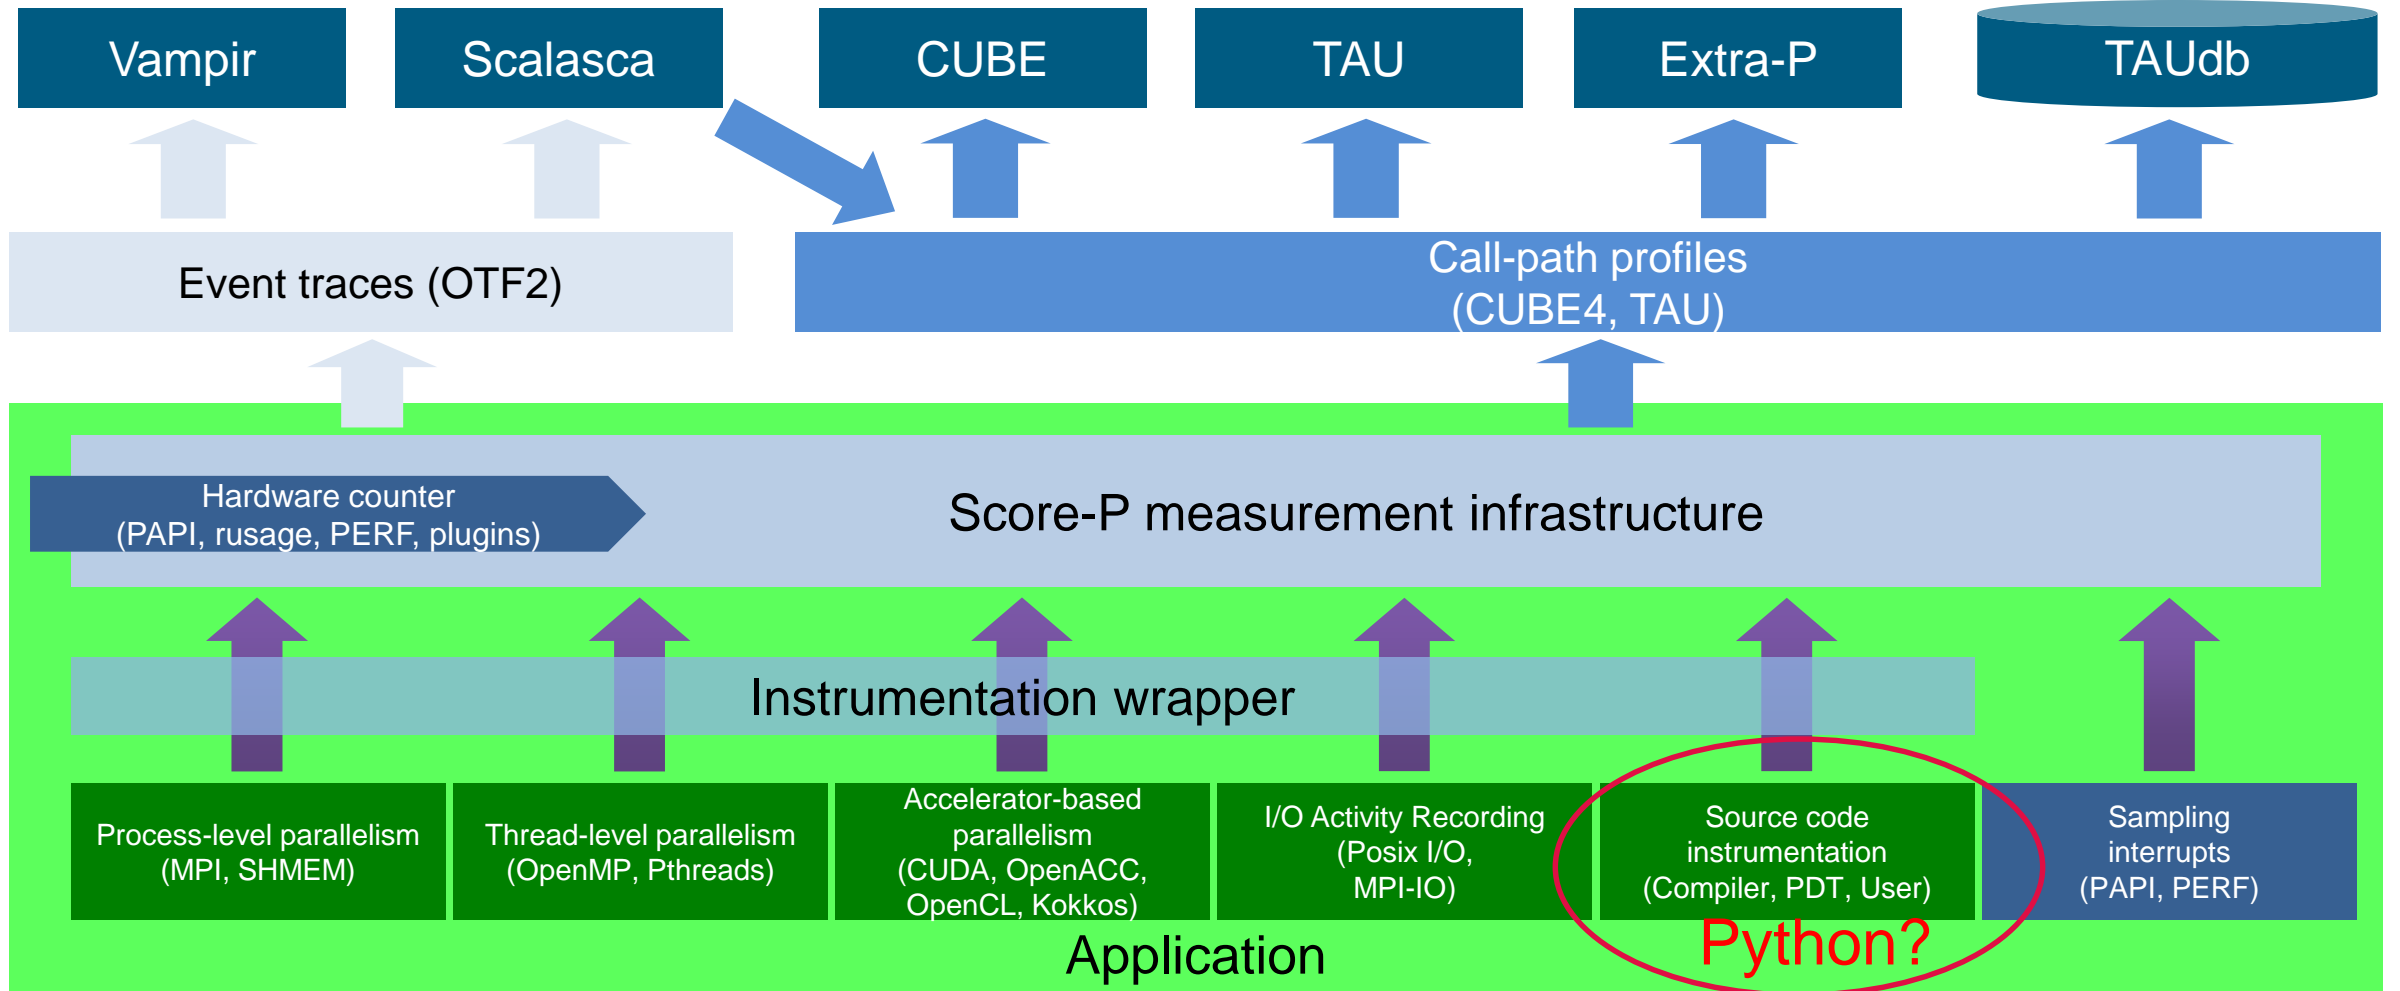




PERFORMANCE ANALYSIS OF PYTHON CODES USING THE SCORE-P ECOSYSTEM

APRIL 8, 2022 | MICHAEL KNOBLOCH

RECAP: Score-P ARCHITECTURE



PYTHON SUPPORT IN SCORE-P

- Follow the Python approach
 - No Score-P adapter for Python
 - Provide an external Python module for Score-P bindings
- Available at GitHub https://github.com/score-p/scorep_binding_python
- Requires Score-P build with `--enable-shared`
- Simple installation via pip (in virtual environment)

```
git clone https://github.com/score-p/scorep_binding_python
cd scorep_python_bindings/
pip3 install .
```

USAGE OF THE SCOREP MODULE

- Measure a Python script:

```
python -m scorep <script.py>
```

- Measure a Python script using MPI:

```
python -m scorep --mpp=mpi <script.py>
```

- Also works for multi-threaded code:

```
python -m scorep --mpp=mpi --thread=pthread <script.py>
```

- Measurement control via Score-P environment variables

```
SCOREP_ENABLE_TRACING=true
```

```
SCOREP_TOTAL_MEMORY=...
```

```
...
```

MEASUREMENT CONTROL

- By default, all Python functions are instrumented → huge overhead
- Can be disabled using the `--noinstrumenter` flag
- It is possible to enable/disable the instrumenter during program runtime

```
with scorep.instrumenter.disable():  
    do_something()
```

```
with scorep.instrumenter.enable():  
    do_something()
```

INSTRUMENTER EXAMPLES (1)

- Dot product:

```
import numpy as np

[...]  
c = np.dot(a,b)  
[...]
```

- Measures dot product and everything else
- Using intrumenter control (and run with `--noinstrumenter`):

```
import numpy as np  
import scorep

[...]  
with scorep.instrumenter.enable():  
    c = np.dot(a,b)  
[...]
```

- Only dot product (and everything below) is measured

INSTRUMENTER EXAMPLES (2)

- Score-P records every change of state of the instrumenter
 - Disabled regions can be named, that will create an event

```
[...]  
def fun_calls(n):  
    if (n>0):  
        fun_calls(n-1)  
  
with scorep.instrumenter.disable("my_fun_calls"):  
    fun_calls(1000000)  
[...]
```

- `my_fun_calls` will appear in measurement, `fun_calls` will not

```
[...]  
with scorep.instrumenter.disable():  
    with scorep.instrumenter.disable("my_fun_calls"):  
        fun_calls(1000000)  
[...]
```

- Neither call will appear in measurement as no state change happend

USER REGIONS

- Since v2.0, the Python bindings support context managers for user regions:

```
with scorep.user.region("region_name"):  
    do_something()
```

- Since v2.1, the Python bindings support also decorators for functions:

```
@scorep.user.region("region_name")  
def do_something():  
    #do some things
```

- If no region name is given, the function name will be used
- Traditional calls to define a region are also supported

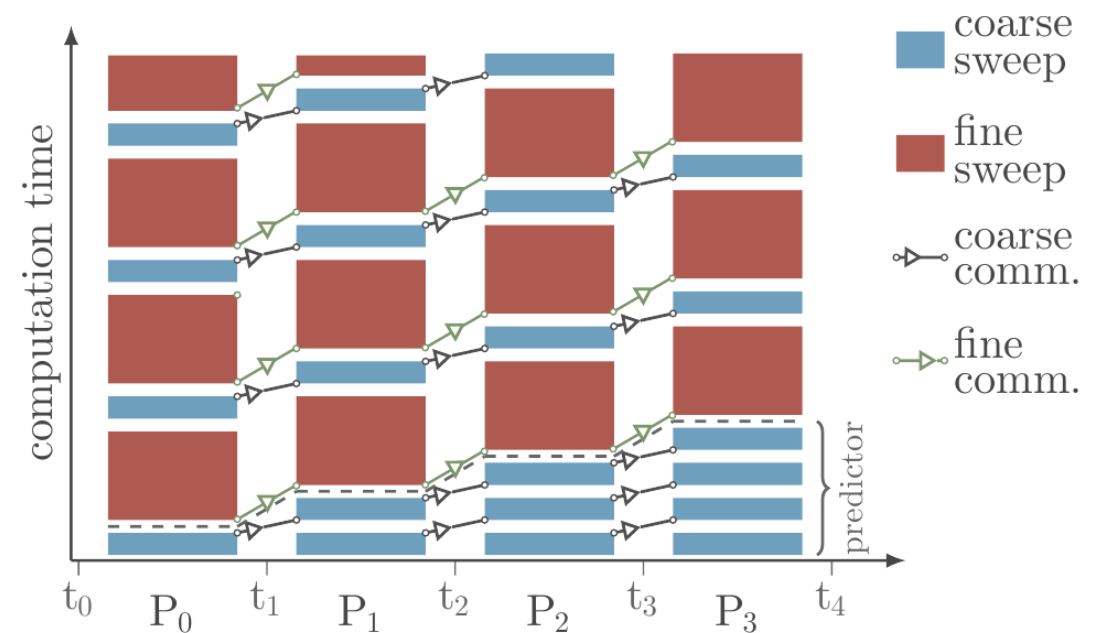
```
scorep.user.region_begin("region_name")  
scorep.user.region_end("region_name")
```


Case Study

PYSDC


PYSDC

- Python implementation of the spectral deferred correction (SDC) approach
- Intended for rapid prototyping and educational purposes
- Framework for testing, evaluating and applying different variants of SDC and PFASST:
 - available implementations of many variants of SDC, MLSDC and PFASST
 - many ordinary and partial differential equations already pre-implemented
 - coupling to FEniCS and PETSc, including spatial parallelism for the latter
 - tutorials to lower the bar for new users and developers
- MPI parallelization in time using **mpi4py**



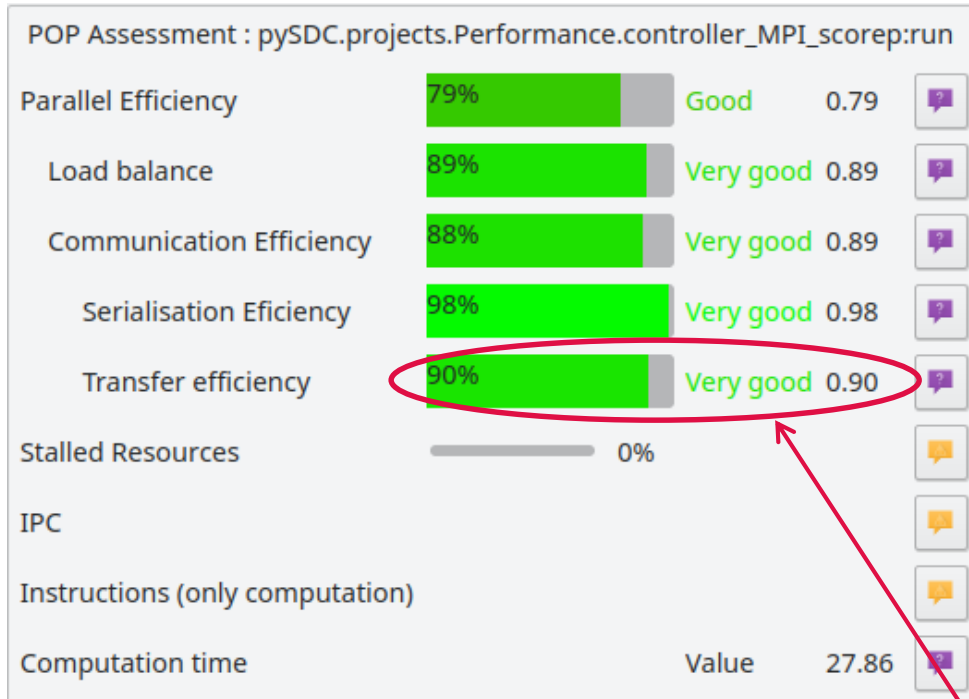
PYSDC FIRST MEASUREMENTS

- Plan: Use Intel compilers and Intel MPI
 - Failed due to an issue with the Intel compiler and the Score-P python bindings
 - Switched to GNU compilers and ParaStationMPI
- First trace was confusing – overwhelming with information
 - Switched to manual instrumentation

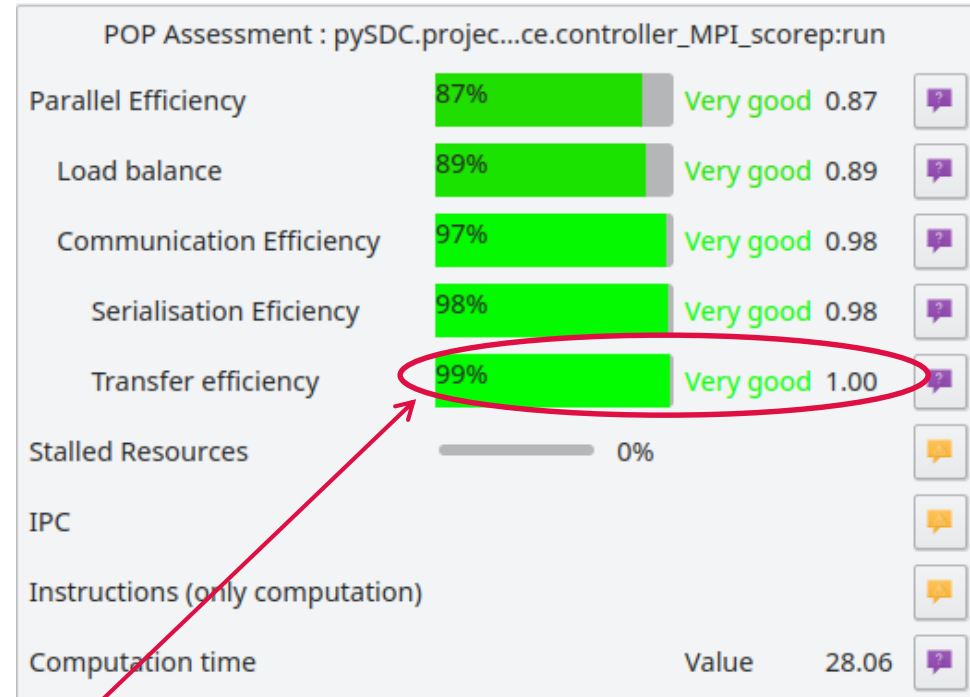


```
1  
2 from mpi4py import MPI  
3 from pySDC.core.Controller import controller  
4  
5 import scorep.user as spu  
6  
7 ...  
8  
9 def run_pfasset(*args, **kwargs):  
10     ...  
11  
12     while not done:  
13         ...  
14         name = f'REGION -- IT_FINE -- {my_rank}'  
15         spu.region_begin(name)  
16         controller.do_fine_sweep()  
17         spu.region_end(name)  
18         ...  
19         name = f'REGION -- IT_DOWN -- {my_rank}'  
20         spu.region_begin(name)  
21         controller.transfer_down()  
22         spu.region_end(name)  
23         ...  
24         name = f'REGION -- IT_COARSE -- {my_rank}'  
25         spu.region_begin(name)  
26         controller.do_coarse_sweep()  
27         spu.region_end(name)  
28         ...  
29         name = f'REGION -- IT_UP -- {my_rank}'  
30         spu.region_begin(name)  
31         controller.transfer_up()  
32         spu.region_end(name)  
33         ...  
34         name = f'REGION -- IT_CHECK -- {my_rank}'  
35         spu.region_begin(name)  
36         controller.check_convergence()  
37         spu.region_end(name)  
38         ...  
39     ...  
40     ...  
41     ...  
42     ...  
43     ...
```

ADVISOR COMPARISON – PS MPI VS. INTEL MPI




ParaStationMPI




Intel MPI

Why?

TO OVERLAP OR NOT TO OVERLAP?

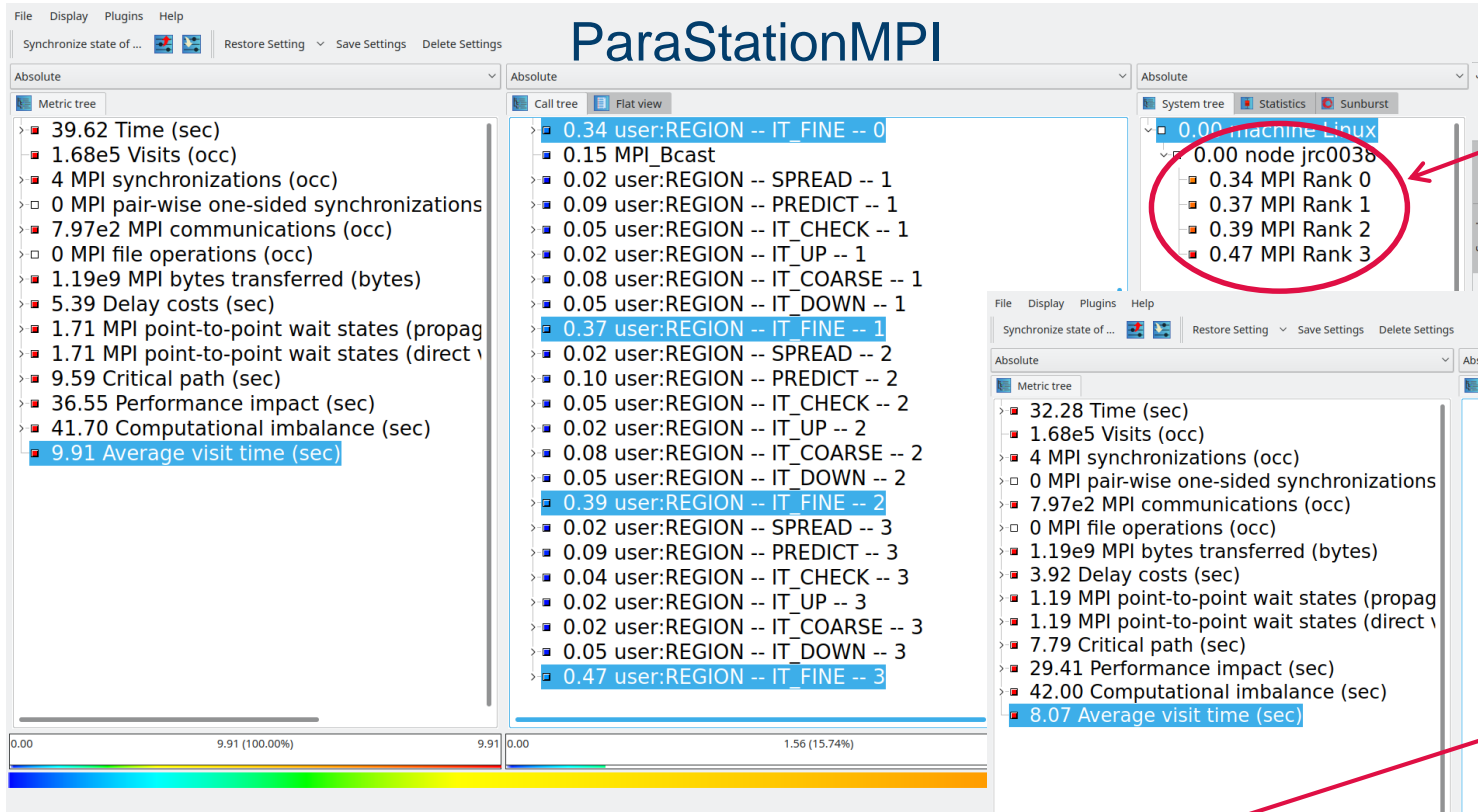


- Parastation MPI:
 - Significant time spend in MPI_Recv
 - No overlap of computation and communication

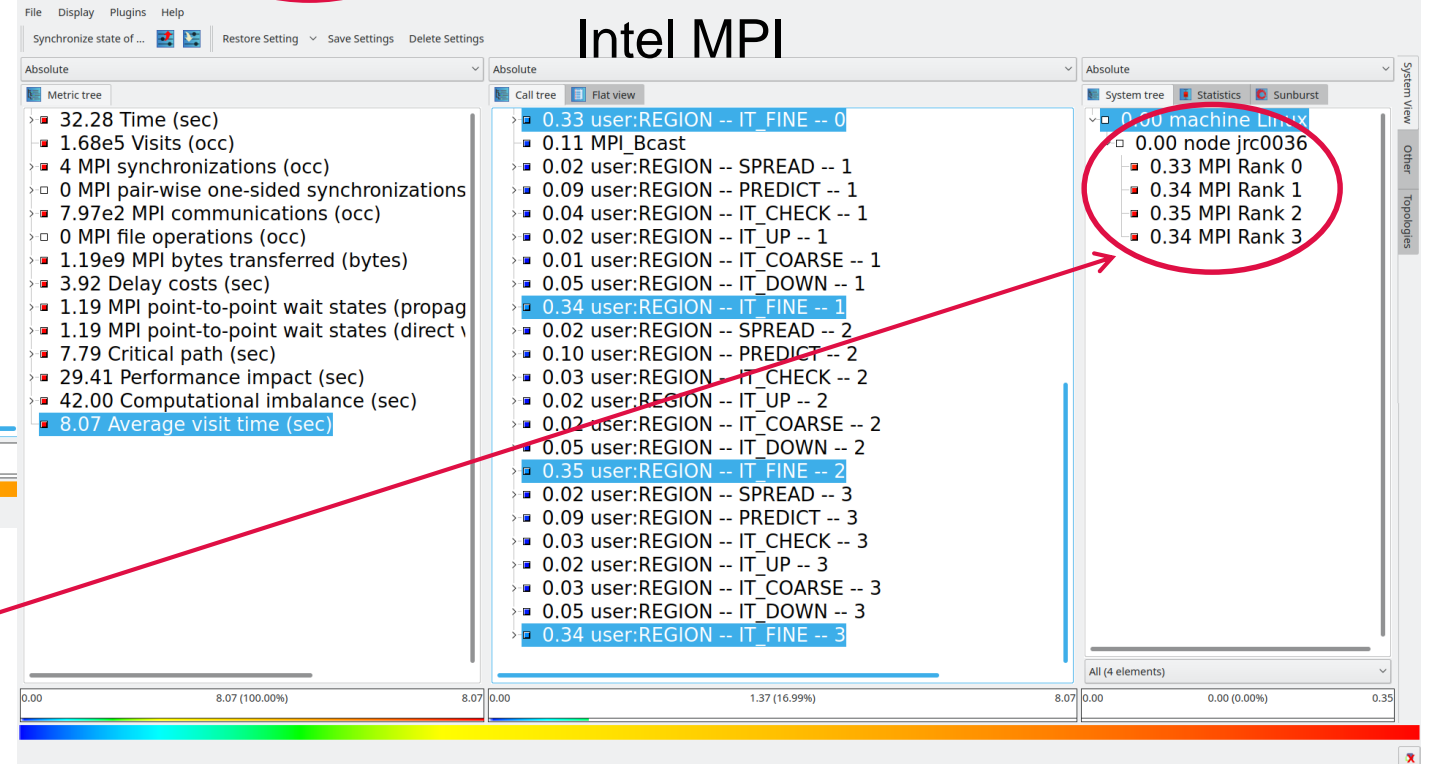


- Intel MPI:
 - Hardly any time spend in MPI_Recv
 - Good overlap of computation and communication

IS IT REALLY A PROBLEM?



- Runtime increases with MPI rank
- Likely a problem at scale



- Runtime well balanced
- Probably not a problem

QUESTIONS?

scalasca 

- <http://www.scalasca.org>
- scalasca@fz-juelich.de



 **Score-P**
Scalable performance measurement
infrastructure for parallel codes

- <http://www.score-p.org>
- support@score-p.org

