



Cineca
TRAINING
High Performance
Computing 2017

IntraNode Optimization

eric.pascolo@cineca.it


mirko.cestari@cineca.it

SuperComputing Applications and Innovation Department





Intranode Optimization






Are you ready?



Today you become
Software Mechanic!!



Explore the engine!



Intel Vtune Profiler

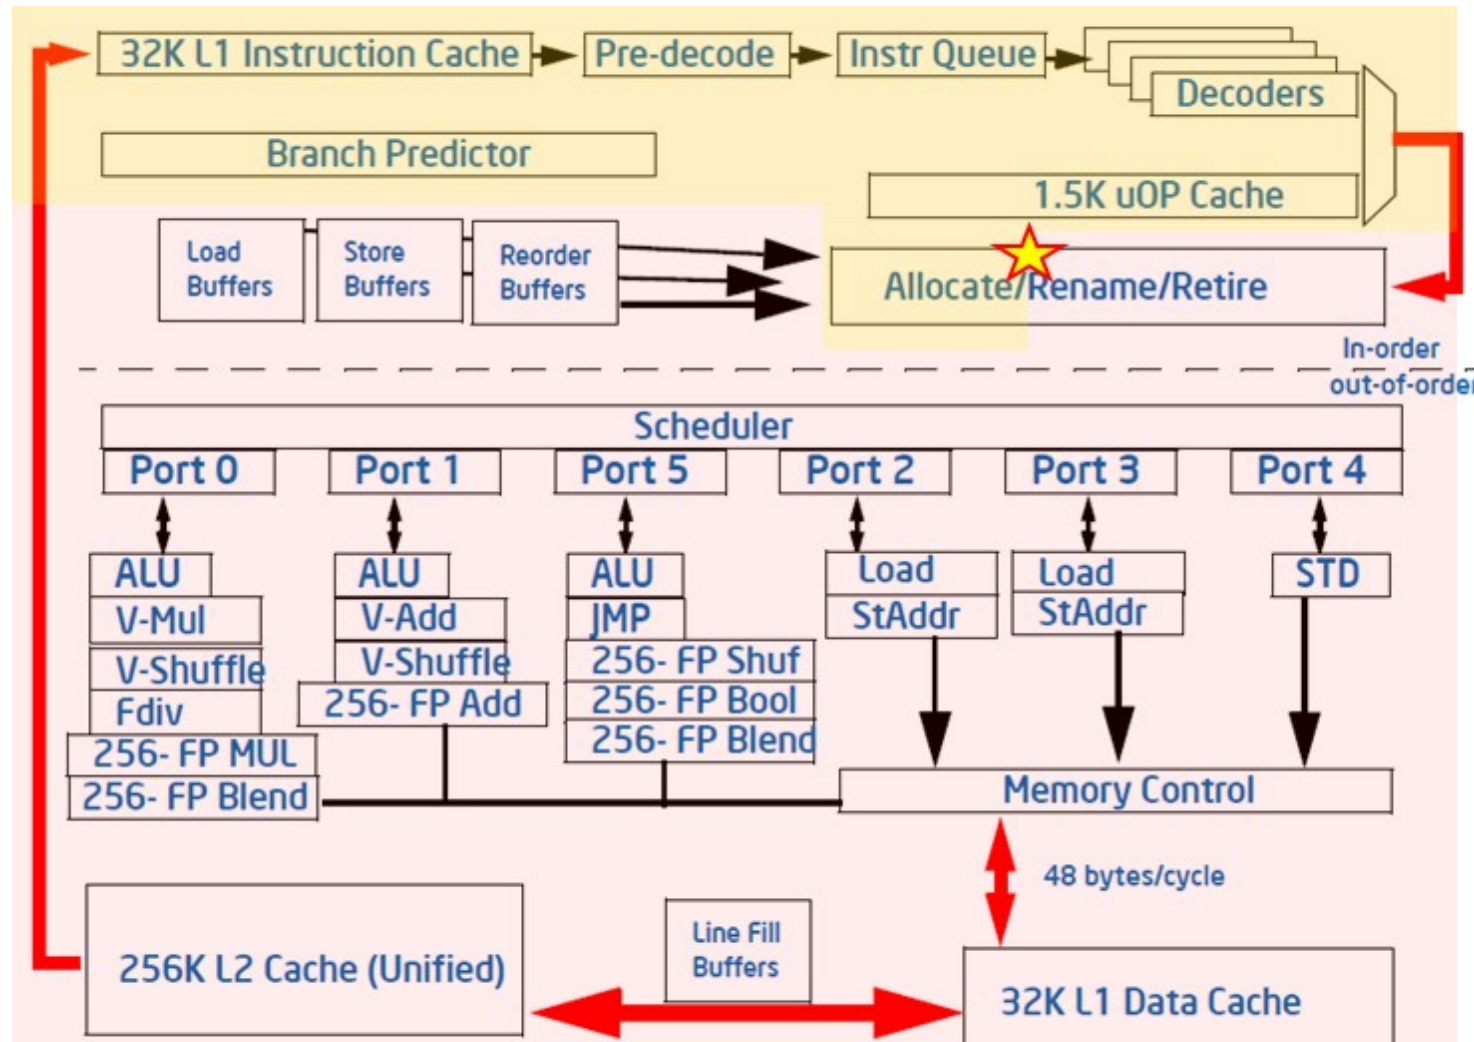


**PROFILING SHOWS THE
PERFORMANCE OF
HARDWARE AND SOFTWARE COUPLED**



Modern Processor Pipeline

Front END

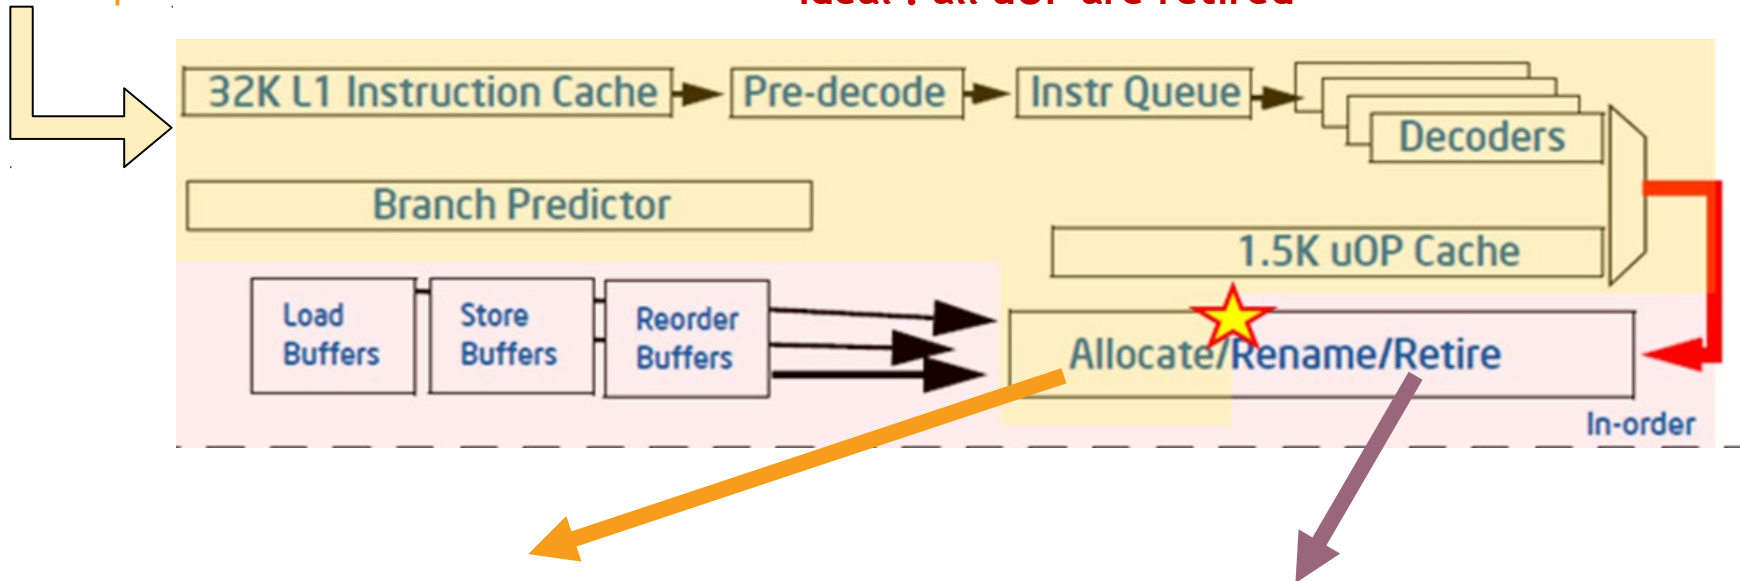




Front End

Compiled CODE

Ideal : all uOP are retired



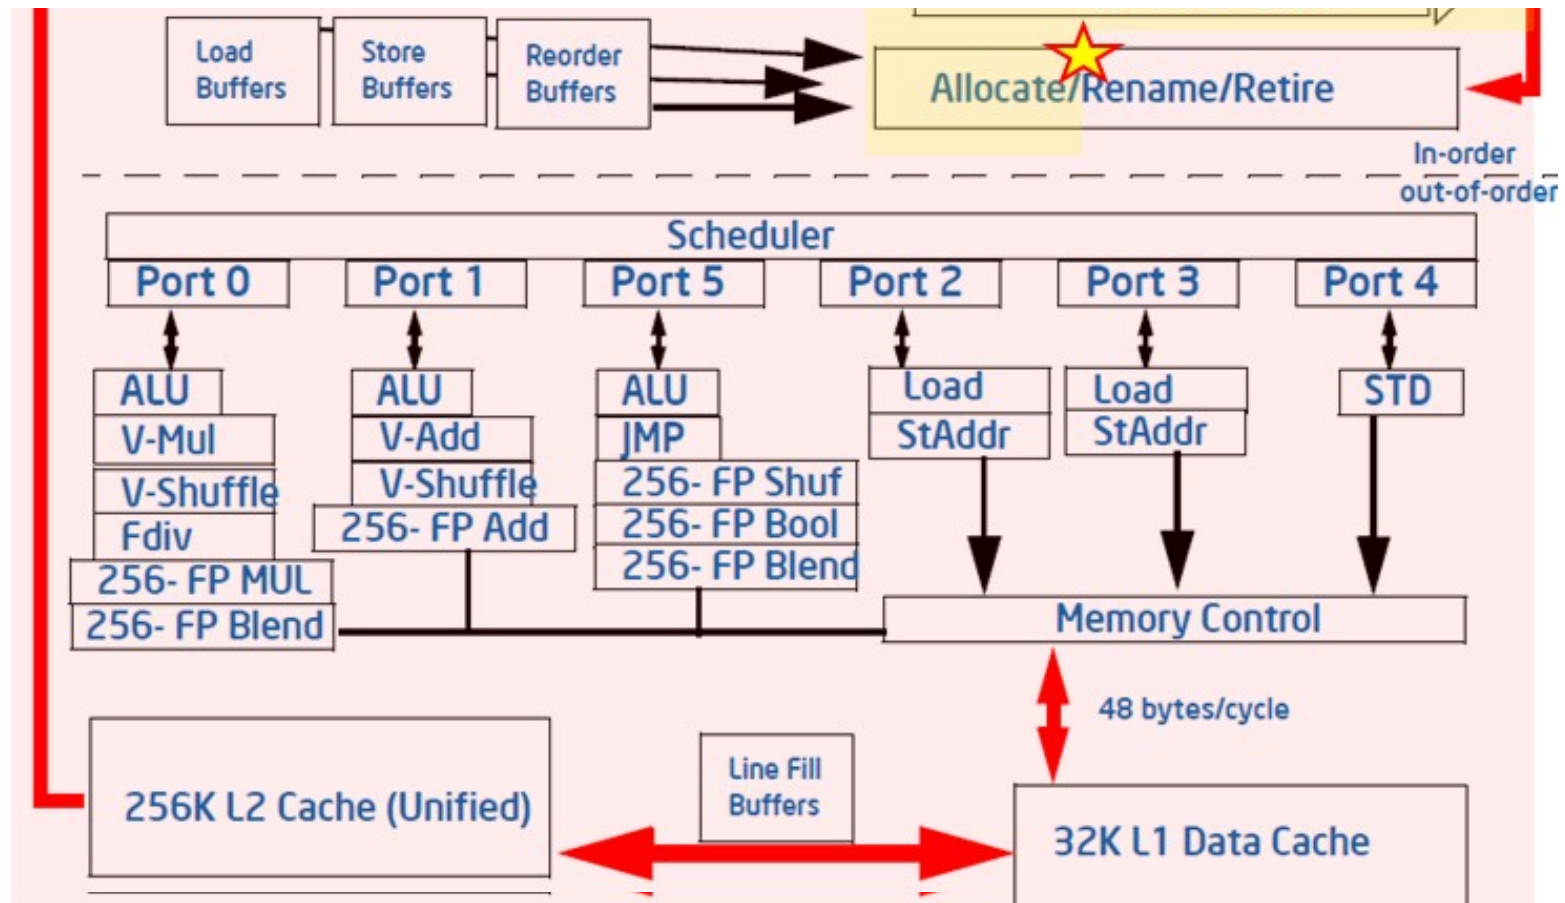
Allocation:
process that
fed the BE of uOP

Retirement:

The completion of a uOp's execution. The results are
committed to the architectural state.

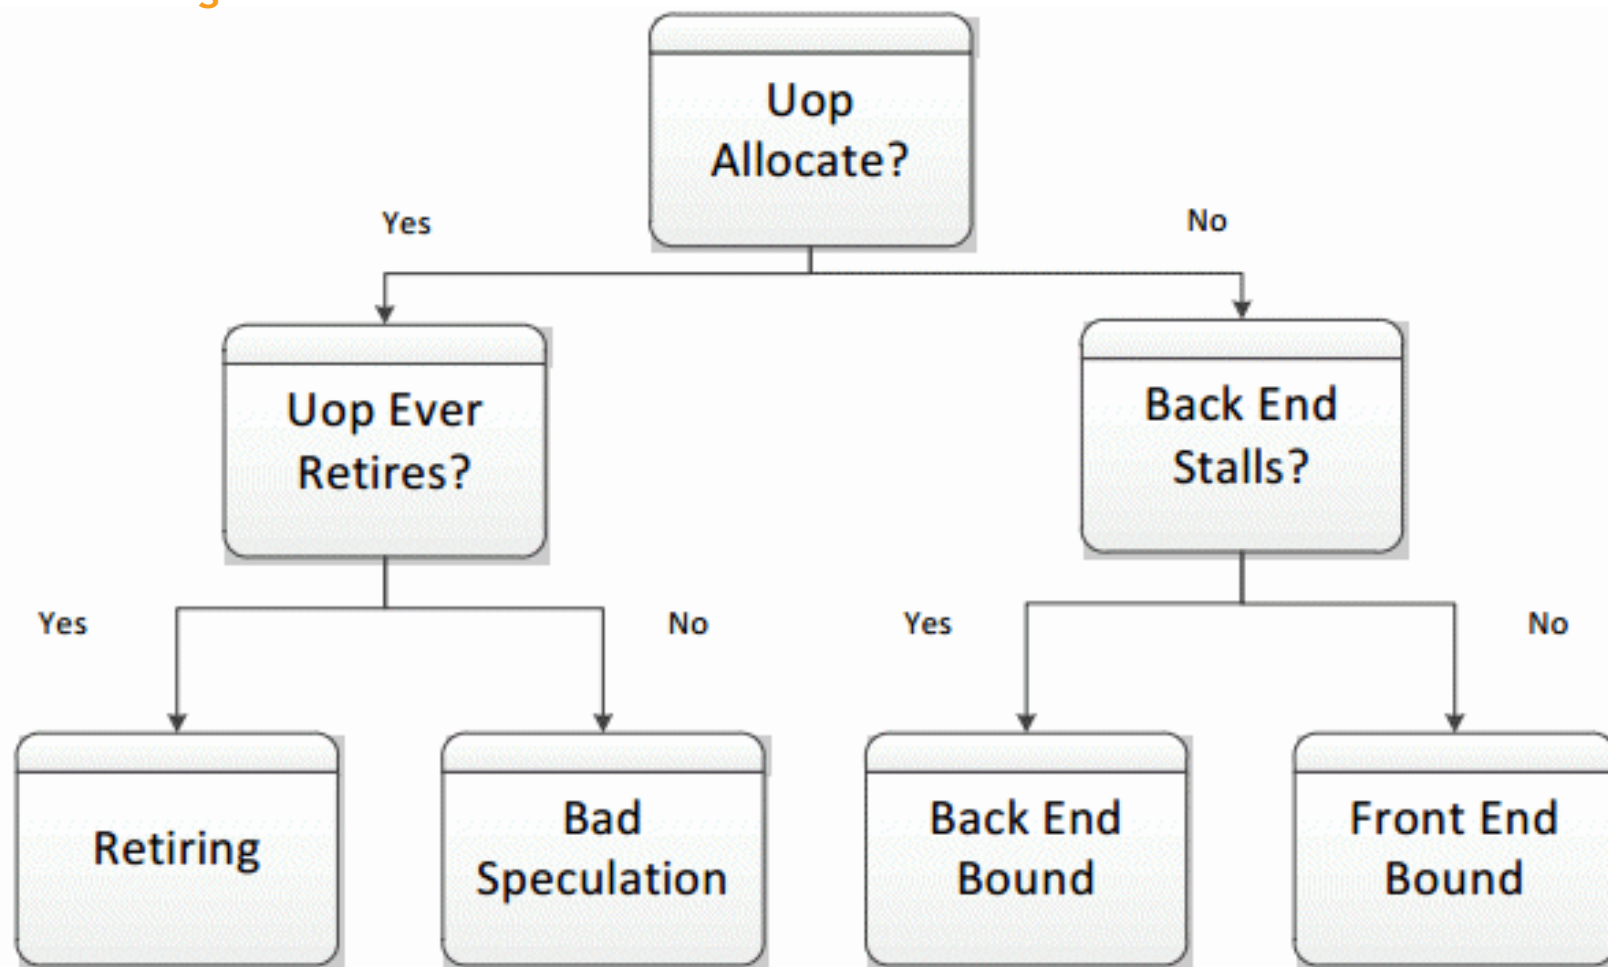


Back End





Intel's Diagram





FE, Retiring and Bad Speculation

Retiring : the performance issues are probably due to an heavy use of micro sequencer
(assistant generator of long stream of uOPs)

Bad Speculation : when the pipeline is busy fetching and executing non-useful operations due to incorrect speculation.

Front End: problem related to code layout, try to solve using PGO(Profiling guided optimization).



PROFILE GUIDED OPTIMIZATION

`icc -prof_gen code.c`

`icc -prof_use code.c`




`Launch code.xx`

On typical dataset



What is the difference?

Matrix Mul




AXPI

$a * X + Y$



What is the difference?

Matrix Mul



AXPI

$a * X + Y$

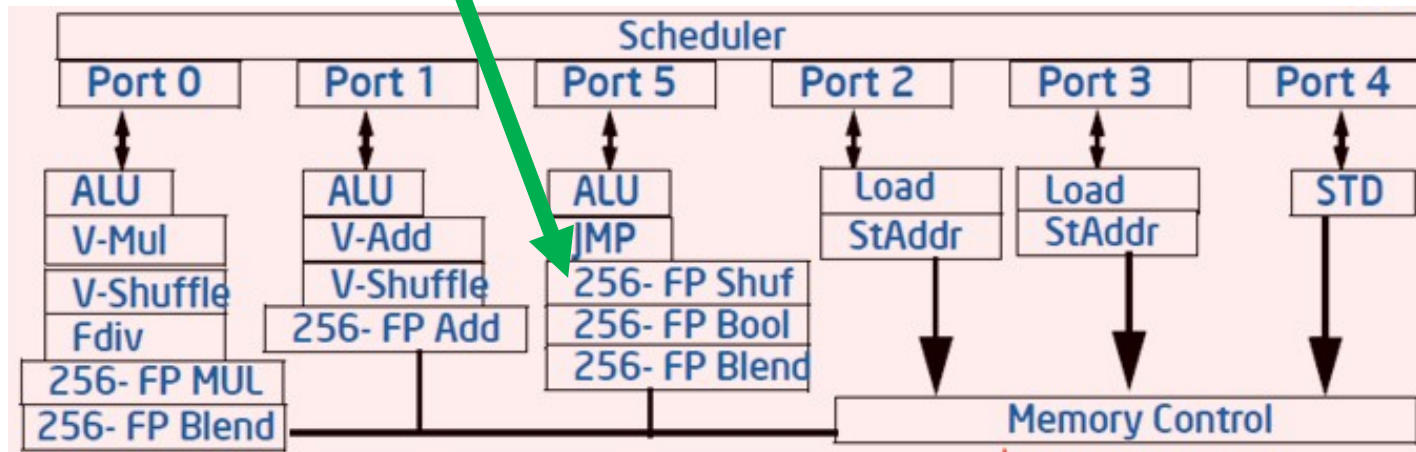
CORE BOUND

MEM BOUND



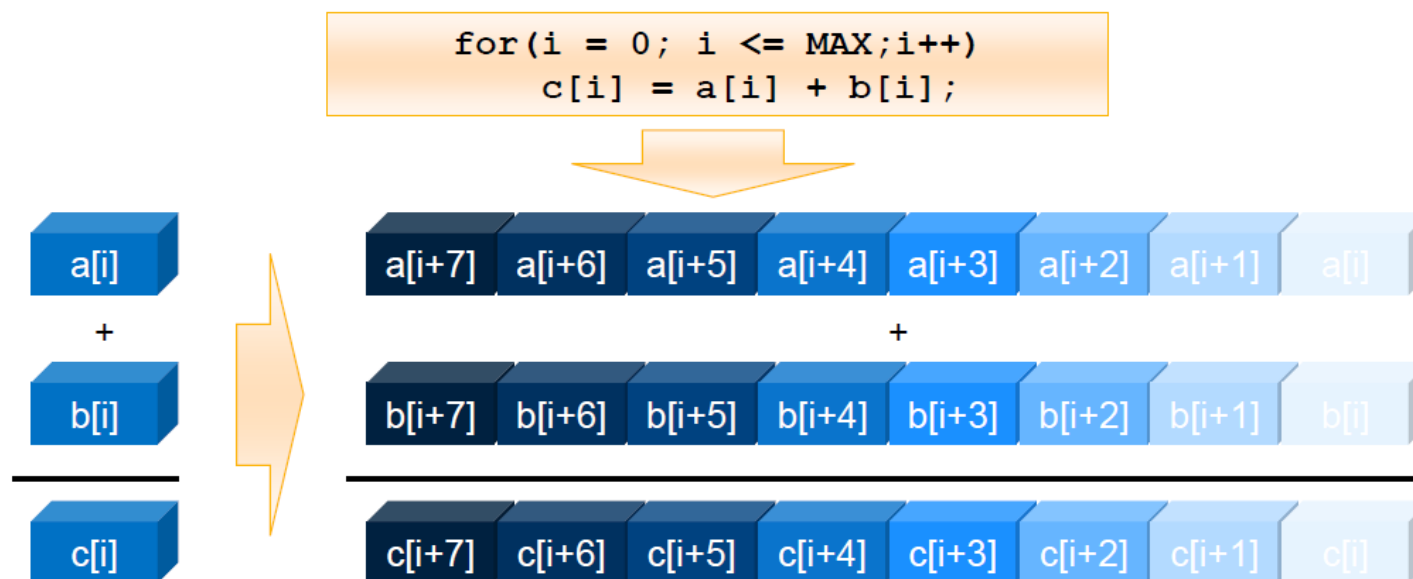
Increase Performance

256 Vector Register





Single Instruction Multiple Data





Many Ways

Compiler Auto
Vectorization

Compiler Assist
Vectorization

Intrinsic Function

Assembler code

Easy

Difficult



Compiler AutoVectorization

`icc -x<ARCH>`

`icc -no-vec`

`icc -ax<ARCH1>,<ARCH2>`

CINECA MARCONI ARCH = avx2

To see code optimization report

`icc -qopt-report 5`



Reason for Vectorization Fails

Data dependencies

Wrong Alignement

Non-unit stride access

Non Vec Math functions

Function calls

Loop body

Too complex



Data Alignment

Data alignment means putting the data at
a memory address equal to some multiple of
the word size (AVX 256 bit).

```
void * _mm_malloc(int size, int word)
```

Remember to compiler that an array is aligned `__assume_aligned(var,word);`



Data Dependency

Flow dependence

```
for(int i;..){  
    X[i] =  
    ... = X[i]  
}
```

Anti dependence

```
for(int i;..){  
    ... = X[i]  
    X[i] = ...  
}
```

Output dependence

```
for(int i;..){  
    X[i] = ...  
    X[i] = ...  
}
```



DOWNLOAD RCM : <https://hpc-forge.cineca.it/svn/RemoteGraph/branch/multivnc/build/dist/Releases/?p=817>

Exercise 1 : VTune Profiling

`module load autoload vtune`

`qsub -l -l select=1:ncpus=1:mem=20GB -l walltime=00:30:00 -q R457635 -A train_scA2017 -X`

Exercise 2 : Vectorization

`module load autoload vtune python/3.5.2`

`source /marconi_scratch/userinternal/epascal1/adv_py_school/bin/activate`