



MPI Learn

Distributed Learning over MPI

Filipe Magalhães
Supervisors: Maurizio Pierini, Jean-Roch Vlimant

16/08/2018

The problem

Speed vs Complexity

Machine Learning




Parallelism



The project

Supercomputers ♥ Machine learning




The project

Supercomputers ♥ Machine learning

MPI Learn

Efficient Framework for Distributed Learning




The project

Supercomputers ♥ Machine learning

MPI Learn

Efficient Framework for Distributed Learning




My work

Make things faster

Threaded validation

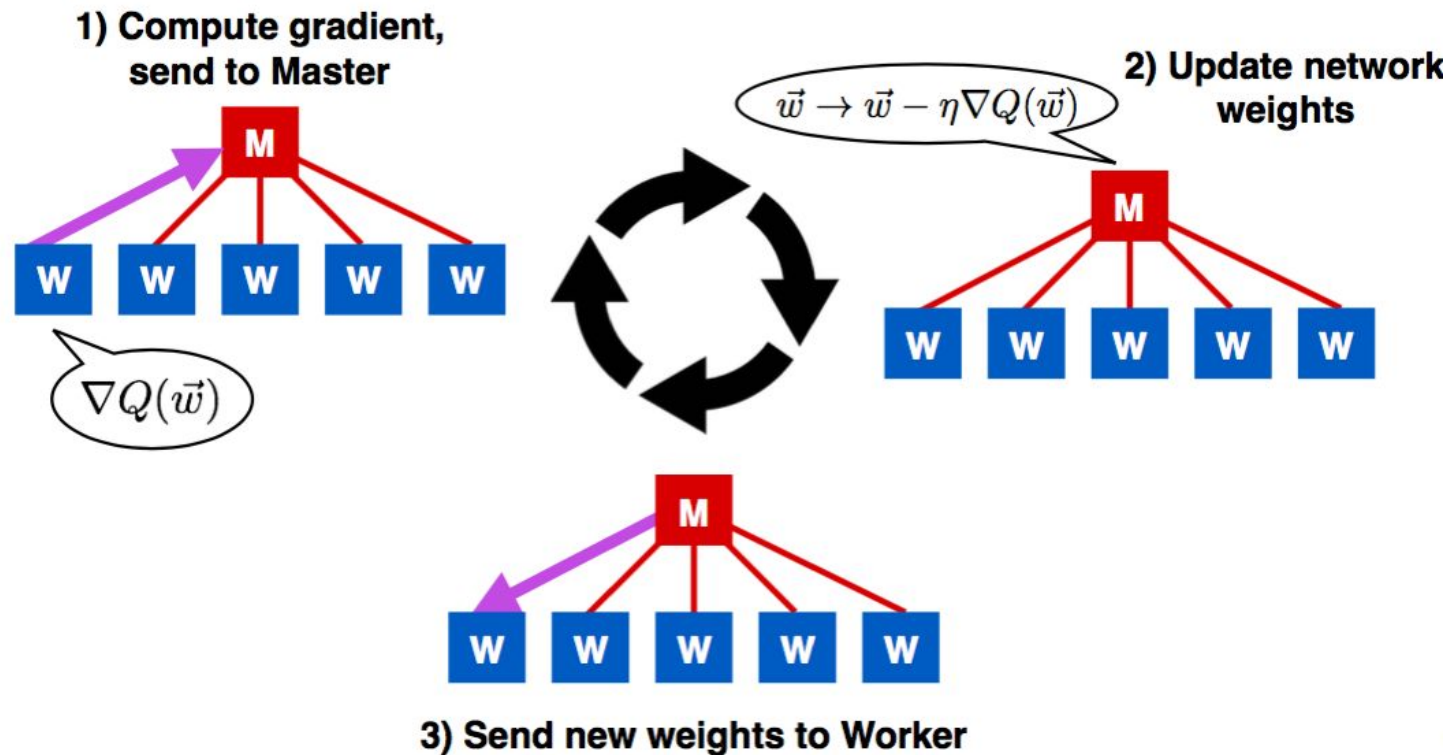
Single process mode

Profiling



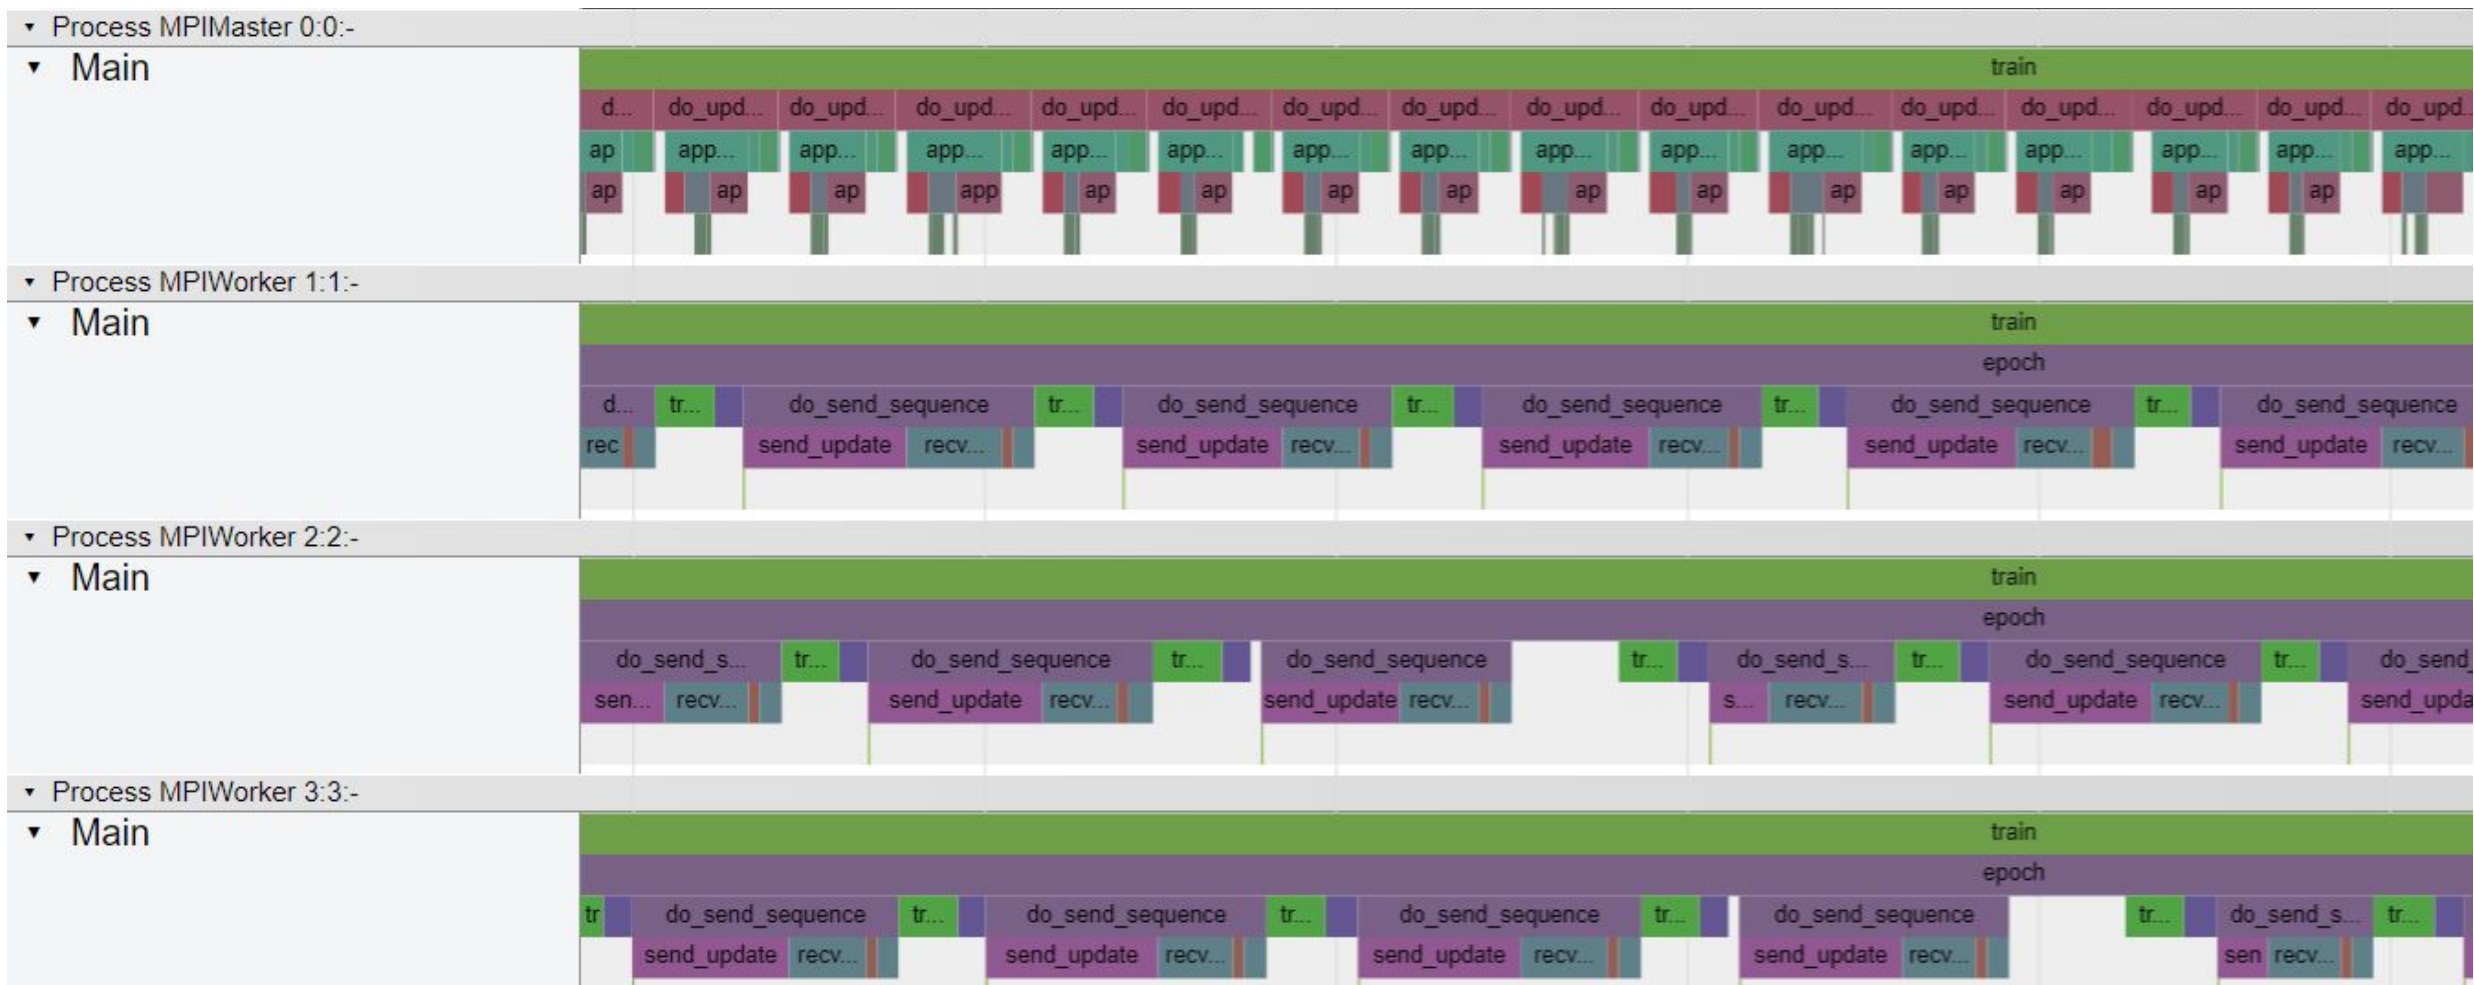
Profiling

Normal execution



Profiling

The problem



Improving

The solution?

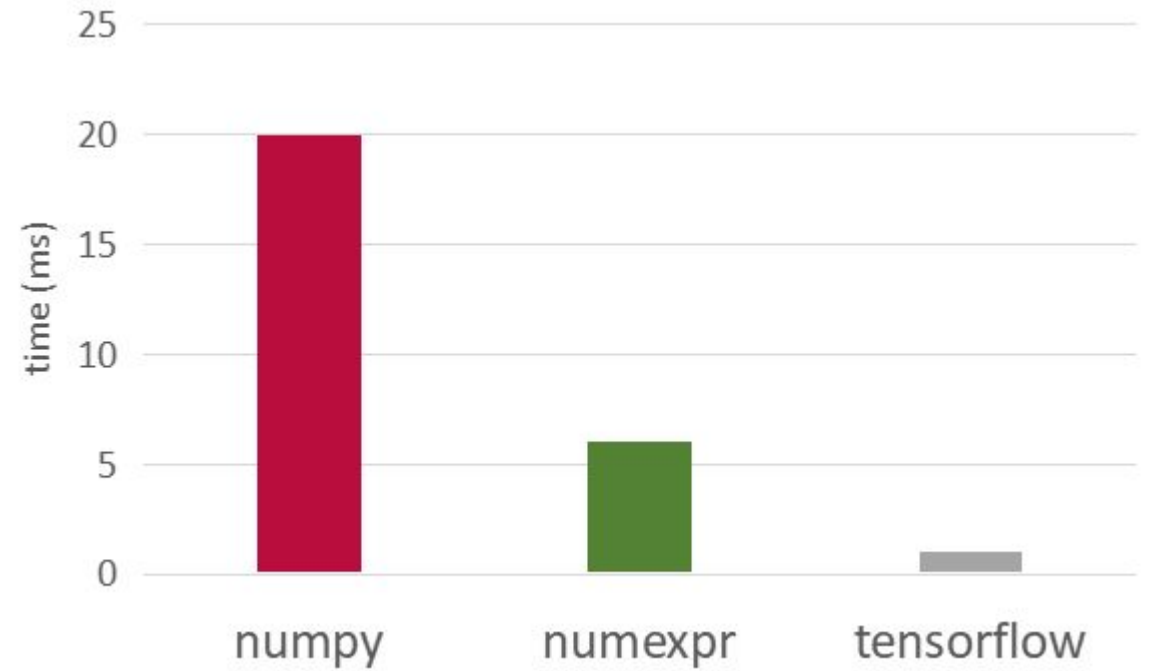
Numpy



numexpr



Tensorflow + GPU



Future

What's next

MPI Learn in active use

Continuous performance improvement

Add features as needed



Thank you

filipepgmagalhaes@tecnico.ulisboa.pt

