SATELLITE IMAGE ORTHORECTIFICATION OVER HDFS USING SPARK

IRT, OCE PROJECT

- Mathias ORTNER <mathias.ortner@irt-saintexupery.com>
- Gregory FLANDIN <gregory.flandin@irt-saintexupery.com>
- Marc SPIGAI <marc.spigai@irt-saintexupery.com>

INTRODUCTION

BIG DATA ?

- store, process, analyse **large** volumes of data
- maximal volume not know a priori
- `cost = c x Volume`
- nature of data or processing may change

SATELLITE IMAGE GROUND SEGMENT

We have / expect :

- large volume of images to store
- large volume of images to explore
- new competitors from “big data world”

OUR WORK !

- well known case : orthorectification of SPOT 6 images (L1 to L2)
- study HDFS + spark over google cloud
- which advantages / drawbacks
- how does it fit whith out usual High Performance approaches

WHAT IS HADOOP ?

IN SHORT

- **Distributed data storage**
  1. store large volumes
  2. ... over low cost hardware
- scalable
- MapReduce
- Not Posix

DATA REPLICATION

- Data is stored on several computers (a cluster)
- Data is split into pieces (typically 64M Bytes)
- HDFS is not POSIX (set of commands)
- One namenode, several datanodes
- FS aware of data locallity

SPLIT AND CLONE

AND DISTRIBUTE

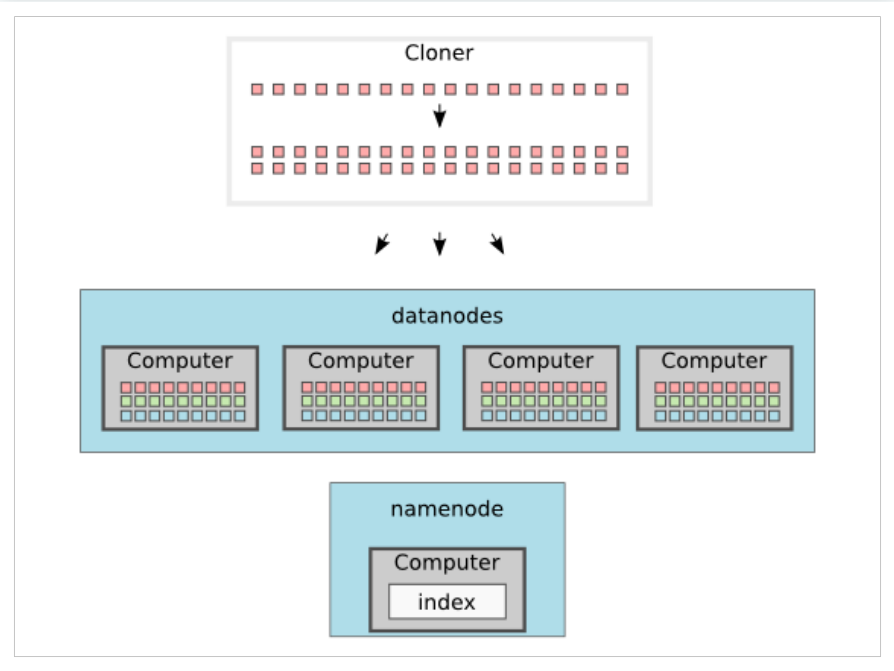

MAPREDUCE (MAP)

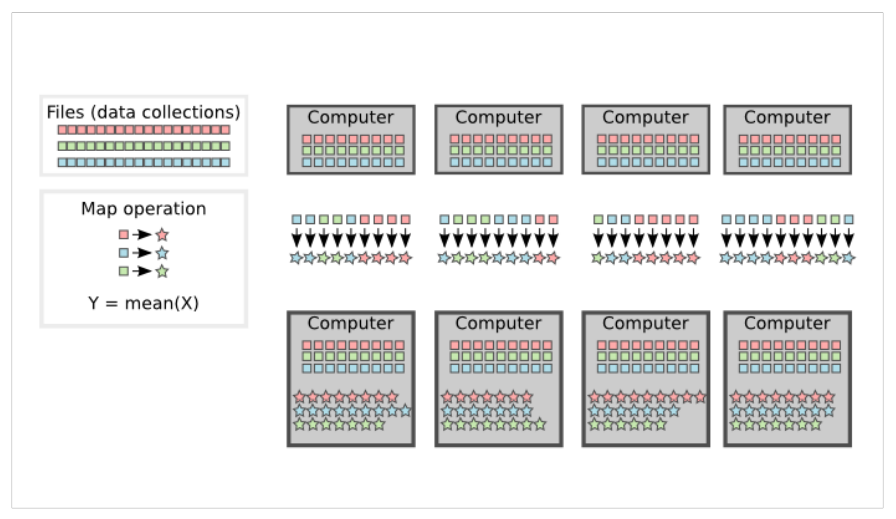

MAPREDUCE (REDUCE)

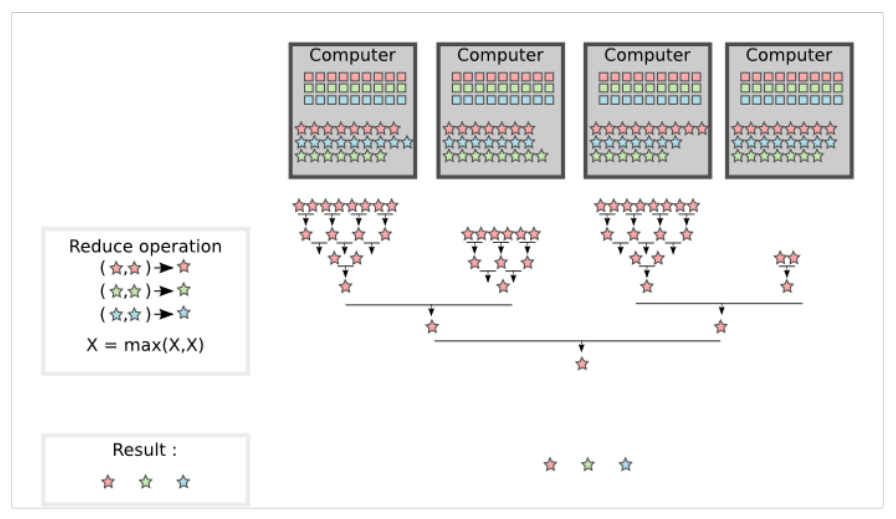

EXAMPLE (LS)

```
[ortner@cluster-data-master ~]$ hdfs dfs -ls /
Found 2 items
-rw-r--r--   2 ortner supergroup 1030588144 2015-04-28 19:09 /brisbane1-T40
```

EXAMPLE (PUT)

```
[ortner@cluster-data-master ~]$ hdfs dfs -put temp/brisbane1-T1000-BB.avro
[ortner@cluster-data-master ~]$ hdfs dfs -ls /
Found 2 items
-rw-r--r--   2 ortner supergroup  281398027 2015-04-30 14:57 /brisbane1-T10
-rw-r--r--   2 ortner supergroup 1030588144 2015-04-28 19:09 /brisbane1-T40
```

EXAMPLE (STATUS OF FILE)

```
[ortner@cluster-data-master ~]$  hdfs fsck /brisbane1-T4000-BP.avro -blocks
Connecting to namenode via http://cluster-data-master:50070
FSCK started by ortner (auth:SIMPLE) from /10.240.11.148 for path /brisbane
/brisbane1-T4000-BP.avro 1030588144 bytes, 8 block(s):  OK
0. BP-1139583743-10.240.69.224-1430248096224:blk_1073741825_1001 len=134217
1. BP-1139583743-10.240.69.224-1430248096224:blk_1073741826_1002 len=134217
2. BP-1139583743-10.240.69.224-1430248096224:blk_1073741827_1003 len=134217
3. BP-1139583743-10.240.69.224-1430248096224:blk_1073741828_1004 len=134217
4. BP-1139583743-10.240.69.224-1430248096224:blk_1073741829_1005 len=134217
5. BP-1139583743-10.240.69.224-1430248096224:blk_1073741830_1006 len=134217
6. BP-1139583743-10.240.69.224-1430248096224:blk_1073741831_1007 len=134217
7. BP-1139583743-10.240.69.224-1430248096224:blk_1073741832_1008 len=910640

Status: HEALTHY
 Total size:    1030588144 B
 Total dirs:    0
 Total files:   1
 Total symlinks:                0
 Total blocks (validated):      8 (avg. block size 128823518 B)
 Minimally replicated blocks:   8 (100.0 %)
 Over-replicated blocks:        0 (0.0 %)
 Under-replicated blocks:       0 (0.0 %)
 Mis-replicated blocks:         0 (0.0 %)
```

```
FSCK ended at Thu Apr 30 15:04:22 UTC 2015 in 1 milliseconds


The filesystem under path '/brisbane1-T4000-BP.avro' is HEALTHY
```

WHAT IS ORTHORECTIFICATION ?

ON GROUND PROJECTION

- image is acquired in sensor geometry (L1)
- but the user needs a projection on ground (L2)
- Projection accounts for :
  - Time measurements
  - Line of sight calibration
  - Satellite attitude measurements
  - Satellite orbit measurements
  - Terrain model

SPOT6 SPECIFICITY : MULTIPLE SENSORS

The system swath is 60km, native resolution is 2.2m.

- two cameras (1 and 2),
- and in each camera, two retinas (A and B),
- in each retina, 5 bands (1 Pan and 4 Multispectral),
- Panchromatic band : 7000 pixels, and in the Multispectral bands 1500 pixels.

The L1 product is therefore actually made of 20 images.

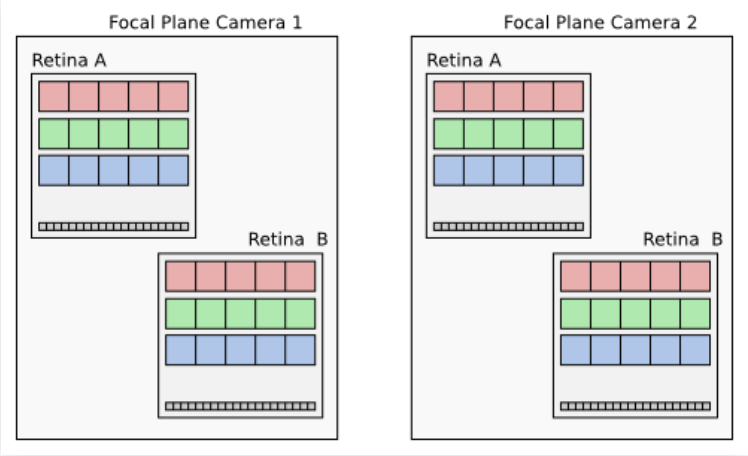

INPUT IS L1

We start from *a Level 1 image* which is the native image in the focal plane geometry

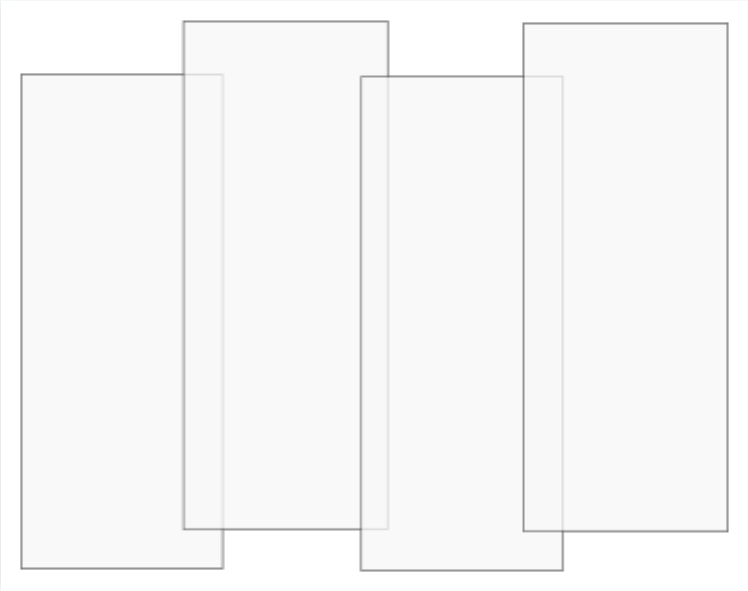

OUTPUT IS L2

We produce *a level 2 image*, i.e. an image that is projected on the ground using a Digital Terrain model and a cartographic frame.

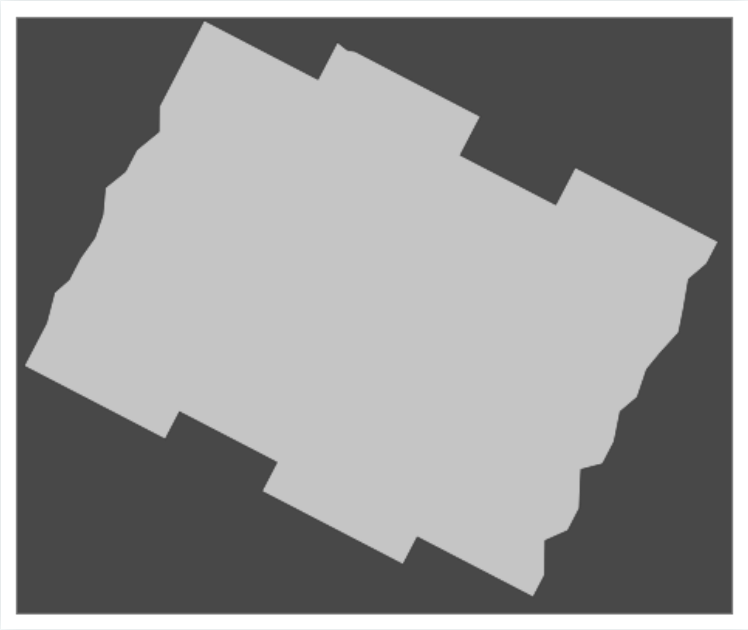

IN BETWEEN :

The production relies on four steps :

- Dtm interpolation
- Inverse localization (from ground to focal plane)
- Image interpolation (pixel lookup, *B spline resampler*)
- Fusion for overlapping parts

IN BETWEEN :

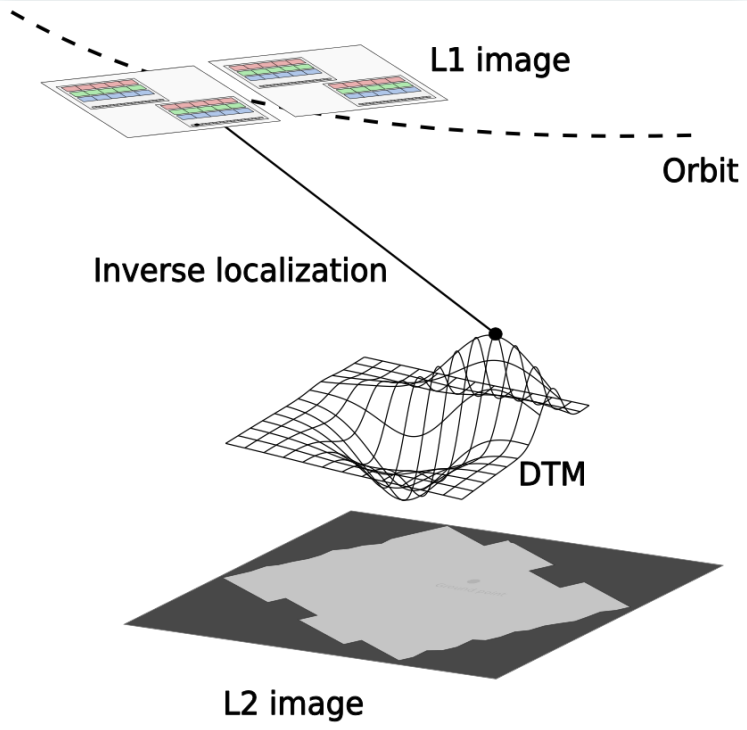

TERRAIN MODEL IS A SHARED RESSOURCE (SRTM)

SRTM terrain model is a standard, freely available, and widely used terrain model.

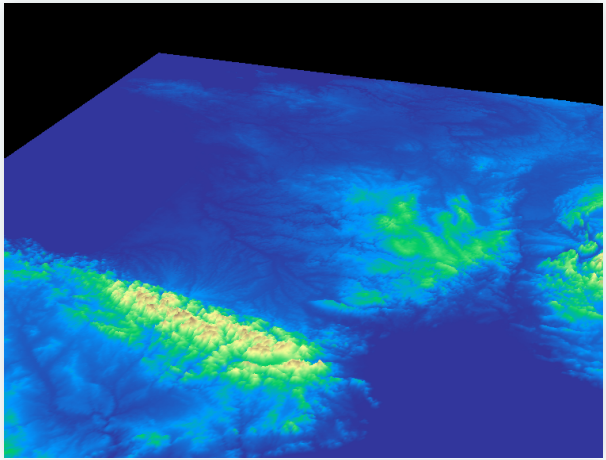

LARGE VOLUMES OF DATA TO BE HANDLED

We focus in this study on a *full* PAN image production, with typical dimensions of

- 44754 x 49135 pixels (1X)
- 110858 x 81327 pixels (4X)
- 352829 x 42817 pixels (7X)

WHAT IS SPARK ?

APACHE PROJECT

Use distributed computation and map reduce easily

It is written in

- akka, on top of
- scala, on top of
- java

It has bindings in python, scala, java.

RESILIENT DISTRIBUTED DATASETS (RDDS)

All data is stored in *collections* of objects called *resilient distributed datasets* (RDDs)

Collections are distributed on the network.

ACTIONS ON RDDS

We have different possible actions :

- Creation
- Mapping (transform a RDD into another one)
- Reductions
- Writing, collection

LAZY EVALUATION

Spark is a functional language

- Define rules...
- ... only needed rules are applied

PAY (MODERATELY) FOR (HIGH) USAGE !

WHAT IS IT ?

Create and use virtual machines....

EXAMPLE : CREATE A MACHINE !

```
#!/bin/bash

gcloud compute instances create cluster-data-master
            --image centos7-image-java-xvfb
            --disk name=data-disk device-name=sdb mode=rw
            --local-ssd interface=SCSI
            --metadata-from-file startup-script=startup.sh
            --machine-type n1-highmem-16
```

```
Created [https://www.googleapis.com/compute/v1/projects/XXXX/zones/europe-w

NAME                ZONE           MACHINE_TYPE  INTERNAL_IP EXTERNAL_IP  S
cluster-data-master europe-west1-b n1-highmem-16 10.240.5.82 104.155.0.44 R
```

HOW MUCH DOES IT COST ?

Machine type	Virtual CPUs	Memory	Typical price (USD) per hour
n1-standard-1	1	3.75GB	$0.038
n1-standard-2	2	7.5GB	$0.076
n1-standard-4	4	15GB	$0.152
n1-standard-8	8	30GB	$0.304
n1-standard-16	16	60GB	$0.608
n1-highmem-2	2	13GB	$0.096
n1-highmem-4	4	26GB	$0.192
n1-highmem-8	8	52GB	$0.384
n1-highmem-16	16	104GB	$0.768

☺ ERGONOMY !

Distribution of algorithm *incredibly easy to write*

WHY USE HDFS / MAPREDUCE ?

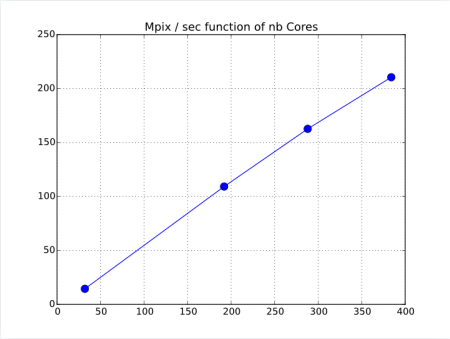

FIRST AND MAIN RESULT

- input is on HDFS
- output is on HDFS
- production is made using Spark

PRODUCE 4 UNITS ON 4 SLAVES...

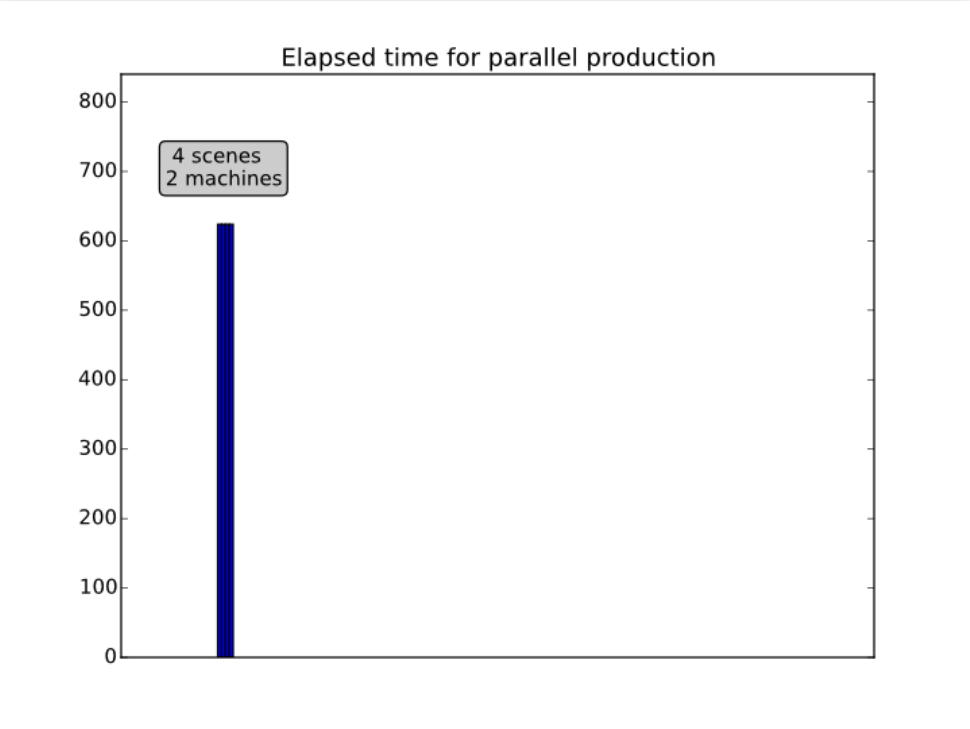

... OR 24 UNITS ON 12 SLAVES ...

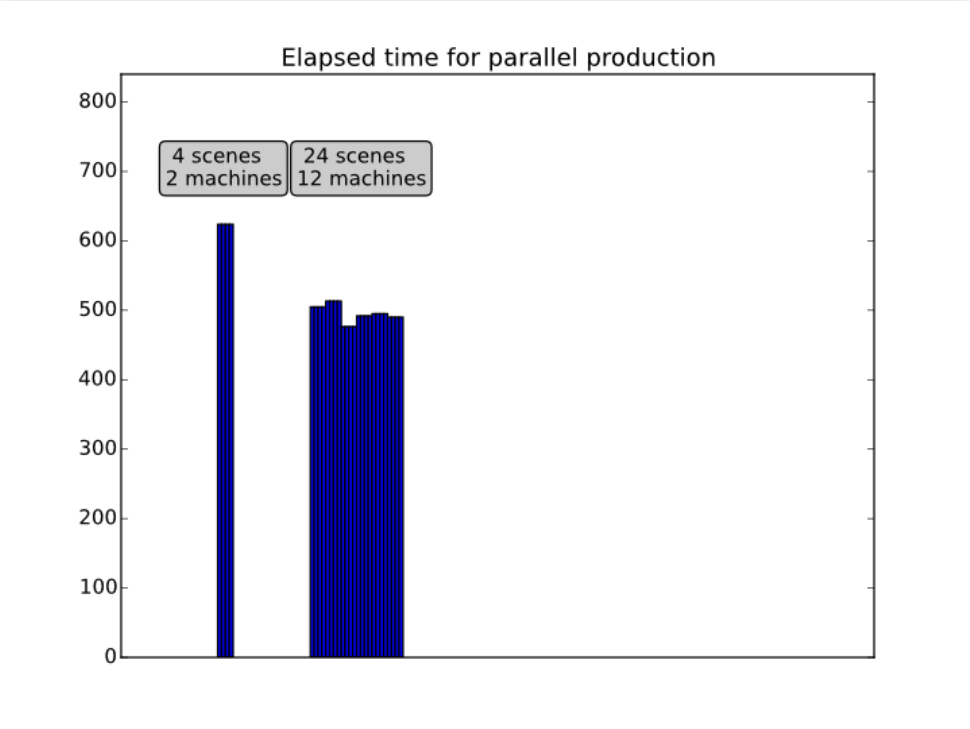

... 36 UNITS ON 18 SLAVES ...

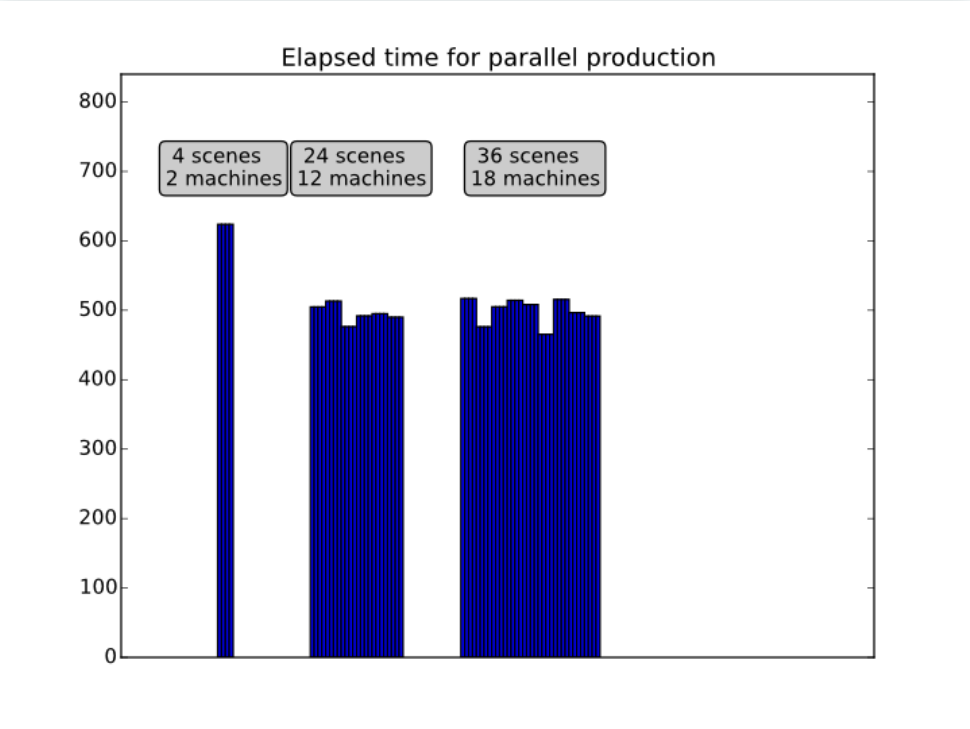

OR 48 ON 24 SLAVES ...

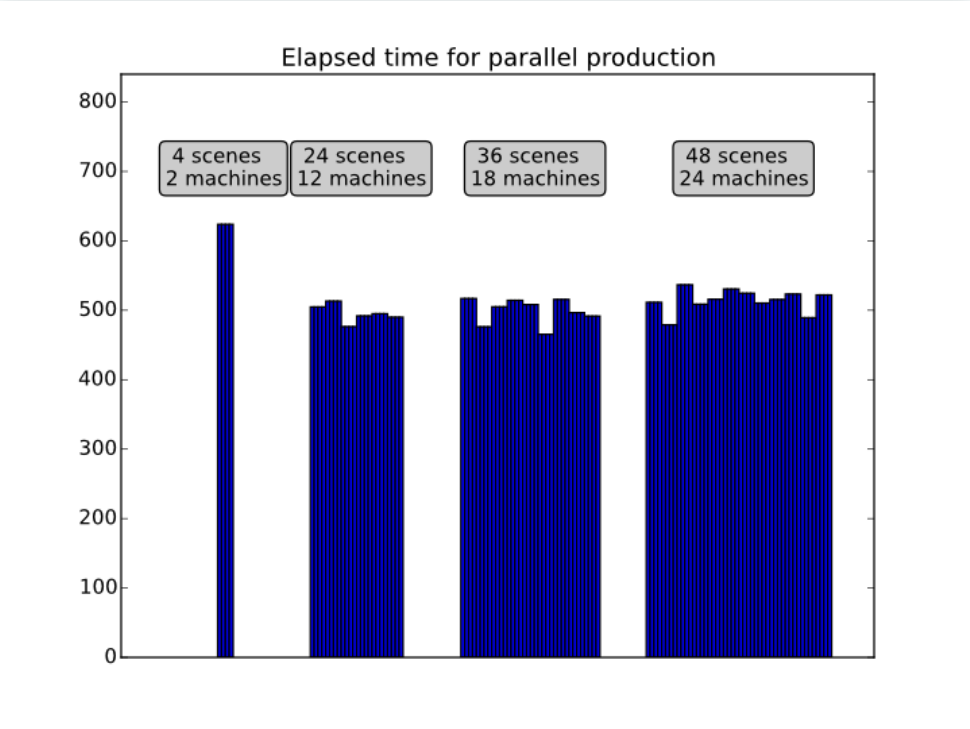

IN SAME AMOUNT OF TIME !

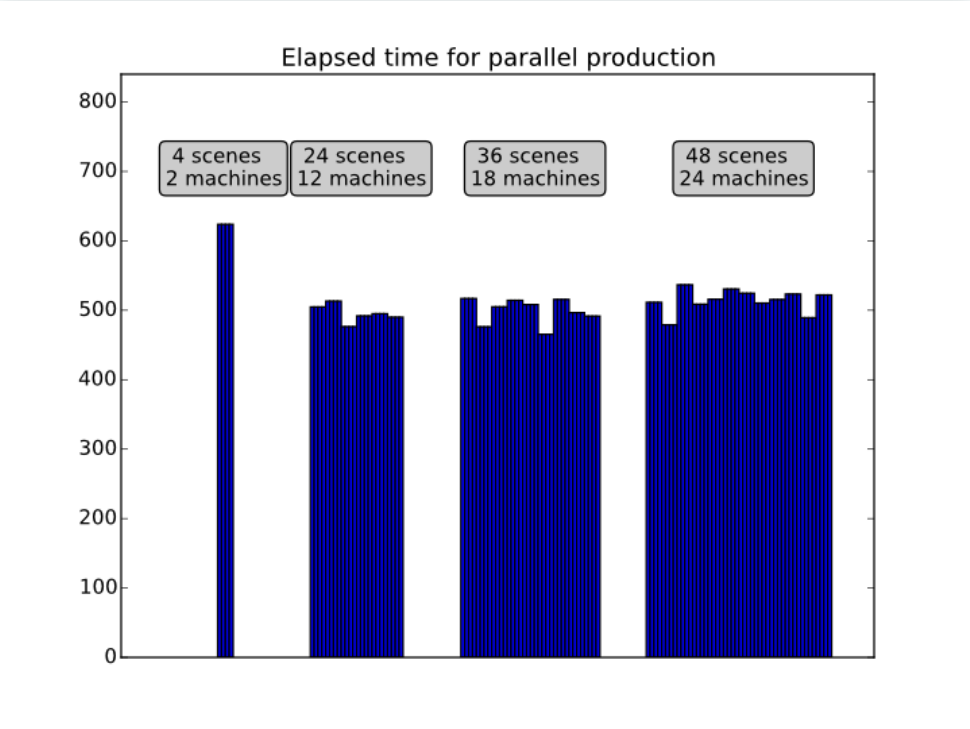

INCREASE PRODUCTION BY ADDING CORES

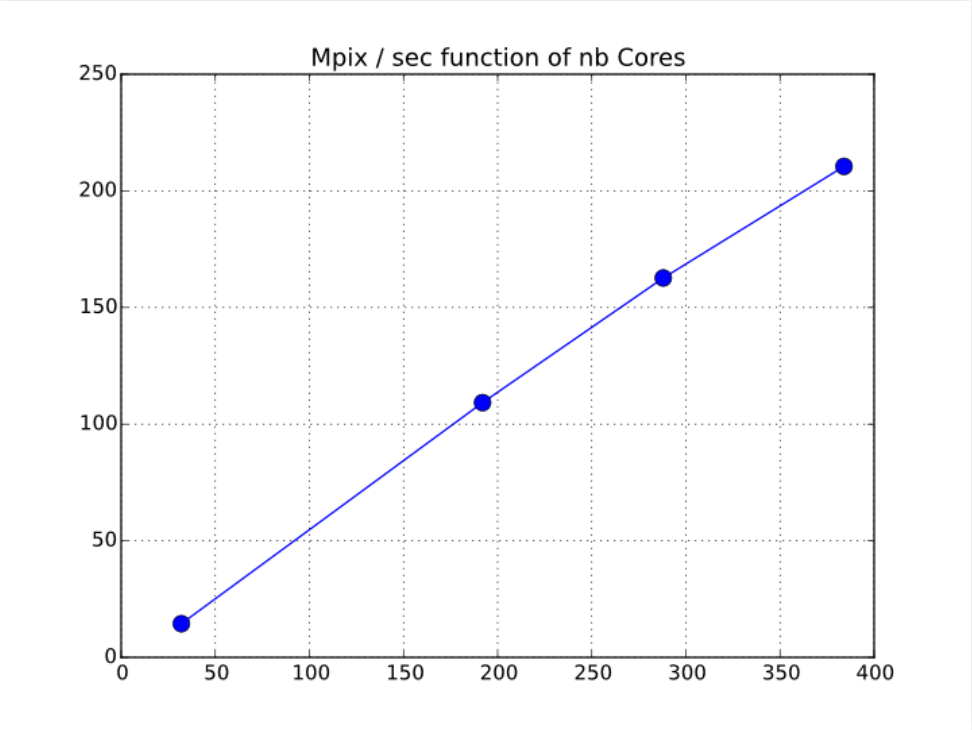

WHAT WE HAVE DONE :

- Analysis of *High Performance mono computer* orthorectification algorithm
- Full implementation of *scalable* orthorectification algorithm
- Deployment on *google cloud engine*
- Analysis of *distribution performances*