
Quantitative Big Imaging - Scaling up

Anders Kaestner

May 11, 2023

CONTENTS

0.1	Scaling Up and Big Data	1
0.2	Programming paradigms	9
0.3	Organization	11
0.4	Databases	13
0.5	Big Data	15
0.6	Environments for distributed computing	22
0.7	Cloud computing	33
0.8	Summary	36

This is the lecture notes for the 10th lecture of the Quantitative big imaging class given during the spring semester 2021 at ETH Zurich, Switzerland.

0.1 Scaling Up and Big Data

```
%load_ext autoreload
%autoreload 2
import seaborn as sns
import matplotlib.pyplot as plt
# from bokeh.resources import CDN
# from bokeh.io import output_notebook

# output_notebook(CDN, hide_banner=True)
local_cluster = False
if local_cluster:
    from dask.distributed import Client, LocalCluster
    cluster = LocalCluster(n_workers=2, threads_per_worker=2)
    client = Client(cluster)
plt.rcParams["figure.figsize"] = (8, 8)
plt.rcParams["figure.dpi"] = 72
plt.rcParams["font.size"] = 16
plt.rcParams["figure.constrained_layout.use"] = True
plt.rcParams['font.family'] = ['sans-serif']
plt.rcParams['font.sans-serif'] = ['DejaVu Sans']
plt.style.use('ggplot')
sns.set_style("whitegrid", {'axes.grid': False})
```

0.1.1 Literature / Useful References

Big Data

- Google's Presentation on Distributed Computing
 - Slides
- J. Dean et al. 2008, MapReduce: Simplified Data Processing on Large Clusters
- Scalable Systems Course
- Tutorial in Hadoop
- Intro to Data Science @UCB

Cluster Computing

- Altintas, I. (2013). Workflow-driven programming paradigms for distributed analysis of biological big data. In 2013 IEEE 3rd International Conference on Computational Advances in Bio and medical Sciences (ICCABS)
- Condor High-throughput Computing
- Condor Setup at ITET
- Sun (now Oracle) Grid Engine

Databases

- Ollion, J., Cochenec, J., Loll, F., Escud, C., & Boudier, T. (2013). TANGO: a generic tool for high-throughput 3D image analysis for studying nuclear organization. *Bioinformatics* (Oxford, England), 29(14)

Cloud Computing

- Amazon S3
- Sitaram, D., & Manjunath, G., 2012. *Moving To The Cloud*. null (Vol. null). Elsevier.
- Duan, P., Wang, W., Zhang, W., Gong, F., Zhang, P., & Rao, Y. (2013). Food Image Recognition Using Pervasive Cloud Computing. In *2013 IEEE International Conference on Green Computing and Communications and IEEE Internet of Things and IEEE Cyber, Physical and Social Computing* (pp. 1631-1637). IEEE.

0.1.2 Outline

- Computer Science Principles
 - Parallelism
 - Distributed Computing
 - Imperative Programming
 - Declarative Programming
- Organization
 - Queue Systems / Cluster Computing
 - Parameterization
 - Databases
- Big Data
 - MapReduce
 - Spark
 - Streaming
- Cloud Computing
- Beyond / The future

0.1.3 Motivation

There are three different types of problems that we will run into.

- Really big data sets
- Many data set
- Explorative studies

Really big data sets

- Several copies of the dataset need to be in memory for processing
 - Computers with more 256GB are expensive and difficult to find
 - Even they have 16 cores so still 16GB per CPU
 - Drive speed / network file access becomes a limiting factor
- If it crashes you **lose** everything
 - or you have to manually write a bunch of messy check-pointing code

Many datasets

- For genome-scale studies 1000s of samples need to be analyzed identically
- Dynamic experiments can have hundreds of measurements
- Animal phenotyping can have many huge data-sets (1000s of 328GB datasets)
- Radiologists in Switzerland alone make 1 Petabyte of scans per year

Exploratory Studies

Not sure what we are looking for

You need it to be

- Easy to develop new analyses
- Quick to test hypotheses

0.1.4 Example Projects

Zebra fish Full Animal Phenotyping

Full adult animal at cellular resolution 1000's of samples of full adult animals

- Imaged at 0.74 μm resolution:
- Images 11500 x 2800 x 628
- 20-40GVx / sample!

Objectives

- Identification of single cells (no down-sampling)
- Cell networks and connectivity
- Classification of cell type
- Registration with histology

0.1.5 Brain Project

Whole brain with cellular resolution 1 cm^3 scanned at $1 \mu\text{m}$ resolution Images \rightarrow 1000 GVx / sample

- Registration separate scans together
- Blood vessel structure and networks
- Registration with fMRI, histology

0.1.6 What is wrong with usual approaches?

Initial workflow

Normally when problems are approached they are solved for a single task as quickly as possible

- I need to filter my image with a median filter with a neighborhood of 5×5 and a square kernel
- then make a threshold of 10
- label the components
- then count how many voxels are in each component
- save it to a file

```
im_in      = imread('test.jpg');
im_filter  = medfilt2(im_in,[5,5]);
cl_img     = bwlabel(im_filter>10);
cl_count   = hist(cl_img,1:100);
dlmwrite(cl_count,'out.txt')
```

You want changes in the workflow

What if want to ...

- compare Gaussian and Median?
- look at 3D instead of 2D images?
- run the same analysis for a folder of images?

You have to rewrite everything, everytime

If you start with a bad approach, it is very difficult to fix,

- big data and
- reproducibility

must be considered from the beginning

0.1.7 Computer Science Principles

Disclosure : There are entire courses / PhD thesis's / Companies about this, so this is just a quick introduction

- Parallelism
- Distributed Computing
- Resource Contention
- Shared-Memory
- Race Conditions
- Synchronization
- Dead lock
- Imperative
- Declarative

What is parallelism?

Parallelism is when you can:

- divide a task into separate pieces
- which can then be worked on at the same time.

An example

If you have to walk 5 minutes and talk on the phone for 5 minutes



Fig. 1: Walking and talking as serial and parallel tasks.

- you can perform the tasks serially which then takes 10 minutes
- you can perform the tasks in parallel which then takes 5 minutes

Some tasks are easy to parallelize while others are very difficult.

Rather than focusing on programming, real-life examples are good indicators of difficulty.

0.1.8 What is distributed computing?

Distributed computing is very similar to parallel computing, but a bit more particular.

- Parallel means you process many tasks at the same time,
- Distributed means you are no longer on the same CPU, process, or even on the same machine.

The distributed has some important implications since once you are no longer on the same machine the number of variables like

- network delay,
- file system issues,
- and other users on the system

becomes a major problem.

Distributed Computing Examples

1. You have 10 friends who collectively know all the capital cities of the world.

- To find the capital of a single country you just yell the country and wait for someone to respond (+++)
- To find who knows the most countries, each, in turn, yells out how many countries they know and you select the highest (++)

2. Each friend has some money with them

- To find the total amount of money you tell each person to tell you how much money they have and you add it together (+)
- To find the **median** coin value, you ask each friend to tell you all the coins they have and you make one master list and then find the median coin (-)

A more parallel median

Assume you have limited number of items like coins.

- Each friend makes a histogram of their coins
- Merge the histograms from each friend to get the median.

0.1.9 Resource Contention

The largest issue with parallel / distributed tasks is the need to access the same resources at the same time

- memory / files
- pieces of information
- network resources

Dead-lock

Dining Philosopher's Problem

- 5 philosophers at the table
- 5 forks
- Everyone needs two forks to eat
- Each philosopher takes the fork on his left...

Challenges in parallel processing

1. Coordination

Parallel computing requires a significant of coordinating between computers for non-easily parallelizable tasks.

2. Mutability

The second major issue is mutability, if you have two cores / computers trying to write the same information at the same it is no longer deterministic (not good)

3. Blocking

The simple act of taking turns and waiting for every independent process to take its turn can completely negate the benefits of parallel computing

0.1.10 Parallel speedup and slowdown

Depend on:

- Available resources
- Implementation
- Scheduling
- Overhead

```
nCPU      = np.linspace(1,8,8)
speedup   = np.array([1,2,3,4,5,6,7,8])
realspeedup = 0.9 * speedup + 0.1
realspeedup[5:] = np.array([5.2, 5.5, 5])

fig, ax1 = plt.subplots(figsize=(10,6))

ax1.set_xlabel('# CPUs')
ax1.set_ylabel('Processing time')
ax1.plot(nCPU, 1/speedup,marker="o", color="limegreen",label = 'Ideal processing time
↵')
ax1.plot(nCPU, 1/realspeedup,marker="o",linewidth=3,color="seagreen", label = 'Real
↵processing time')
```

(continues on next page)

(continued from previous page)

```
ax2 = ax1.twinx() # instantiate a second axes that shares the same x-axis

ax2.set_ylabel('Speedup') # we already handled the x-label with ax1
ax2.plot(nCPU, speedup, marker="o",color="cornflowerblue", label = 'Ideal speedup')
ax2.plot(nCPU, realspeedup, marker="o", linewidth=3,color="navy",label = 'Real speedup
↪')
ax1.axvline(x=5,linewidth=1, color='darkorange', ymax=0.9)
ax1.axvline(x=7,linewidth=1, color='crimson',ymax=0.9)
fig.legend();
ax1.set_title("Performance of parallel computing");
```

```
-----
NameError                                Traceback (most recent call last)
Cell In[2], line 1
----> 1 nCPU      = np.linspace(1,8,8)
      2 speedup  = np.array([1,2,3,4,5,6,7,8])
      3 realspeedup = 0.9 * speedup + 0.1

NameError: name 'np' is not defined
```

Real performance test run - Spotting hotspots

A time consuming component

Introduced parallel code

The two graphs show the execution time of a tomographic reconstruction as a system task including data IO and preprocessing. Splitting the timing per task in this way makes it possible to identify which part of the processing that consumes much processing time. In this example it appears that the single threaded implementation of the ring cleaning algorithm is takes a dominant fraction of the total processing time. This is therefore the target of tuning. There is a remarkable speed-up by distributing the task on the available cores.

This example, there are components that are easy parallelize by altering the algorithm while others are harder. The yellow block labelled “other” mainly represents the time needed to read input data and to store the results. This brings us to take a closer look at Amdahls law.

Theoretical speed-up - Amdahl's law

The code can be divided into two parts:

- sequential
- parallelizable

$$S_{latency}(s) = \frac{1}{(1-p) + \frac{p}{s}}$$

p - time consumed by parallelizable codes - number of cores

0.1.11 Challenges in distributed processing

Inherits all of the problems of parallel programming with a whole variety of new issues.

Sending Instructions / Data Afar

Fault Tolerance

If you have 1000 computers working on solving a problem and one fails, you do not want your whole job to crash

Data Storage

How can you access and process data from many different computers quickly without very expensive infrastructure

0.2 Programming paradigms

0.2.1 Imperative Programming

Directly coordinating tasks on a computer.

- Languages like C, C++, Java, Matlab
- Exact orders are given (implicit time ordering)
- Data management is manually controlled
- Job and task scheduling is manual
- Potential to tweak and optimize performance

Making a soup (from lecture 1)

1. Buy vegetables at market
2. *then* Buy meat at butcher
3. *then* Chop carrots into pieces
4. *then* Chop potatoes into pieces
5. *then* Heat water
6. *then* Wait until boiling then add chopped vegetables
7. *then* Wait 5 minutes and add meat

0.2.2 Declarative

- Languages like SQL, Erlang, Haskell, Scala, Python, R can be declarative
- Goals are stated rather than specific details
- Data is automatically managed and copied
- Scheduling is automatic but not always *efficient*

Making a soup (from lecture 1)

- Buy vegetables at market $\rightarrow shop_{veggies}$
- Buy meat at butcher $\rightarrow shop_{meat}$
- Wait for $shop_{veggies}$: Chop carrots into pieces $\rightarrow chopped_{carrots}$
- Wait for $shop_{veggies}$: Chop potatoes into pieces $\rightarrow chopped_{potatos}$
- Heat water
- Wait for $boiling_{water}, chopped_{carrots}, chopped_{potatos}$: Add chopped vegetables
- Wait 5 minutes and add meat

0.2.3 Comparison

They look fairly similar, so what is the difference?

Imperative soup

Declarative soup

1. Buy {carrots, peas, tomatoes} at market
 2. *then* Buy meat at butcher
 3. *then* Chop carrots into pieces
 4. *then* Chop potatoes into pieces
 5. *then* Heat water
 6. *then* Wait until boiling then add chopped vegetables
 7. *then* Wait 5 minutes and add meat
- Buy {carrots, peas, tomatoes} at market $\rightarrow shop_{veggies}$
 - Buy meat at butcher $\rightarrow shop_{meat}$
 - Wait for $shop_{veggies}$: Chop carrots into pieces $\rightarrow chopped_{carrots}$
 - Wait for $shop_{veggies}$: Chop potatoes into pieces $\rightarrow chopped_{potatos}$
 - Heat water
 - Wait for $boiling_{water}, chopped_{carrots}, chopped_{potatos}$: Add chopped vegetables
 - Wait 5 minutes and $,shop_{meat}$: add meat

The second is needlessly complicated for one person, but what if you have a team:

- how can several people make an imperative soup faster (chopping vegetables together?)
- How can many people make a declarative soup faster? Give everyone a different task (not completely efficient since some tasks have to wait on others)

0.2.4 Results

Imperative

- optimize specific tasks (chopping vegetables, mixing) so that many people can do it faster
 - Matlab/Python do this with fast-fourier-transforms (automatically uses many cores to compute faster)
- make many soups at the same time (independent)
 - This leads us to cluster-based computing

Declarative

- run everything at once
- each core (computer) takes a task and runs it
- execution order does not matter
 - wait for portions to be available (dependency)

An alternative - Lazy Evaluation

- do not run anything at all
- until something needs to be exported or saved
- run only the tasks that are needed for the final result
 - never buy tomatoes since they are not in the final soup

0.3 Organization

One of the major challenges of scaling up experiments and analysis is keeping all of the results organized in a clear manner.

As we have seen in the last lectures, many of the results produced many text files

- many files are difficult to organize
- Matlab, R are designed for in-memory computation
- Datasets can have many parameters and be complicated
- Transitioning from Excel to Python, Matlab, or R means rewriting everything

0.3.1 Queue Computing

Queue processing systems (like Sun Grid Engine, Oracle Grid Engine, Apple XGrid, Condor, etc) are used to manage

Resources

Computers, memory, storage

- a collection of processors (CPU and GPU)
- memory, local storage
- access to bandwidth or special resource like a printer
- for a given period of time

Jobs

Tasks to be run

- specific task to run
- necessary (minimal/maximal) resources to run with
- including execution time

Users

People who submit jobs

- accounts submitting jobs
- it can be undesirable for one user to dominate all of the resources all the time

Based on a set of rules for how to share the resources to the users to run tasks.

0.3.2 Structure of Cluster

Master (or Name) Node(s)

The node with which every other node communicates, the main address.

Worker Nodes

The nodes where the computation is performed.

Scheduler

The actual process that decides which jobs will run using which resources (worker nodes, memory, bandwidth) at which time

0.4 Databases

A database is a collection of data stored in the format of tables:

- a number of columns (data categories in the table)
- and rows (stored items)

```
from IPython.display import display, Markdown
import pandas as pd
display(Markdown('### Animals\nHere we have an table of the animals measured in an
experiment and their weight'))
display(pd.DataFrame(dict(id=(1, 2, 3),
                          Weight=(100, 40, 80)
                          )))

display(Markdown(
    '### Cells\nThe cells is then an analysis looking at the cellular structures'))

display(pd.DataFrame(dict(
    Animal=( 1, 2, 3),
    Type=("Cancer", "Healthy", "Cancer"),
    Anisotropy=(0.5, 1.0, 0.5),
    Volume=(1, 2, 0.95))))
```

Non-consecutive header level increase; 0 to 3 [myst.header]

	id	Weight
0	1	100
1	2	40
2	3	80

Non-consecutive header level increase; 0 to 3 [myst.header]

	Animal	Type	Anisotropy	Volume
0	1	Cancer	0.5	1.00
1	2	Healthy	1.0	2.00
2	3	Cancer	0.5	0.95

0.4.1 SQL

SQL (pronounced Sequel) stands for Structured Query Language and is *nearly* universal for both

- searching (called querying)
- and adding (called inserting) data into databases.

SQL is used in various forms from

- Firefox storing its preferences locally (using SQLite)
- to Facebook storing some of its user information (MySQL and Hive).

So referring to the two tables we defined in the last entry, we can use SQL to get information about the tables independently of how they are stored (a single machine, a supercomputer, or in the cloud)

SQL - Basic queries

- Get the volume of all cells

```
SELECT Volume FROM Cells
```

→ [1, 2, 0.95]

- Get the average volume of all cancer cells

```
SELECT AVG(Volume) FROM Cells WHERE Type = "Cancer"
```

→ 0.975

We could have done these easily without SQL using e.g., Excel or a Pandas data frame in python.

More Advanced SQL

- Get the volume of all cells in heavy mice

```
SELECT Volume FROM Cells WHERE Animal IN  
(SELECT id FROM Animal WHERE Weight>80)
```

- Get weight and average cell volume for all mice

```
SELECT ATable.Weight, CTable.Volume FROM Animals as ATable  
INNER JOIN Cells as CTable on (ATable.id=CTable.Animal)
```

→ [1, 0.95]

0.4.2 Beyond SQL: NoSQL

Basic networks can be entered and queries using SQL but relatively simple sounding requests can get complicated very quickly

Network Analysis

If we try to store cells and their connections in a SQL database, we can handle millions of cells and connections easily in a structured manner. However trying to perform analysis is trickier

- *How many cells are within two connections of each cell*

```
SELECT id, COUNT(*) AS connection_count FROM Cells as CellsA  
INNER JOIN Network as NetA ON Where (id=NetA.id1)  
INNER JOIN Network as NetB ON Where (NetA.id2=NetB.id1)
```

This is *still* readable but becomes very cumbersome quickly and difficult to manage

NoSQL (Not Only SQL)

A new generation of database software which extends the functionality of SQL to allow for

- more scalability ([MongoDB](#))
- or specificity for problems like networks or graphs called generally **Graph Databases**

0.5 Big Data

0.5.1 Definition

Velocity, Volume, Variety

When a ton of heterogeneous is coming in fast. We need:

- **Performant**
- **Scalable**
- **Flexible**

You are ready when...

- **scaling isn't scary** 10X, 100X, 1000X is the same amount of effort
- **you are starving for enough data** Director of AMPLab said their rate limiting factor is always enough interesting data
- **no 'clicks' per sample** Everything is automated, no human interaction required during processing

0.5.2 A brief oversimplified story

Google ran into 'big data' and its associated problems years ago:

- Peta- and exabytes of websites to collect and make sense of.
- Google uses an algorithm called PageRank(tm) for evaluating the quality of websites.

They could have probably used existing tools if page rank were some magic program that could read and determine the quality of a site

```
for every_site_on_internet
  current_site.rank=secret_pagerank_function(current_site)
end
```

Just divide all the websites into a bunch of groups and have each computer run a group, **easy!**

PageRank

While the actual internals of PageRank are not public, the general idea is that sites are ranked based on how many sites link to them

```
for current_site in every_site_on_internet
  current_pagerank = new SecretPageRankObj(current_site);
  for other_site in every_site_on_internet
    if current_site is_linked_to other_site
      current_pagerank.add_site(other_site);
    end
  end
end
current_site.rank=current_pagerank.rank();
end
```

Complexity $O(N^2)$

How do you divide this task?

- Maybe try and divide the sites up: english_sites, chinese_sites, ...
 - Run pagerank and run them separately.
 - What happens when a chinese_site links to an english_site?
- Buy a really big, really fast computer?
 - On the most-powerful computer in the world, one loop would take months!

It gets better

- What happens if one computer / hard-drive crashes?
- Have a backup computer replace it (A backup computer for every single system)
- With a few computers ok, with hundreds of thousands of computers?
- What if there is an earthquake and all the computers go down?
- PageRank doesn't just count
- Uses the old rankings for that page
- Run pagerank many times until the ranks converge

Google's Solution: MapReduce (part of it)

some people claim to have had the idea before, Google is certainly the first to do it at scale

Several engineers at Google recognized common elements in many of the tasks being performed. They then proceeded to divide all tasks into two classes **Map** and **Reduce**

Map

Map is where a function is applied to every element in the list and the function depends only on exactly that element
 $\vec{L} = [1, 2, 3, 4, 5] f(x) = x^2 \text{map}(f \rightarrow \vec{L}) = [1, 4, 9, 16, 25]$

Reduce

Reduce is more complicated and involves aggregating a number of different elements and summarizing them. For example the Σ function can be written as a reduce function $\vec{L} = [1, 2, 3, 4, 5] g(a, b) = a + b$
Reduce then applies the function to the first two elements, and then to the result of the first two with the third and soon until all the
 $\vec{L}) = g(g(g(1, 2), 3), 4), 5)$

0.5.3 MapReduce

Google designed a framework for handling distributing and running such of jobs on clusters.

So for each job a dataset (\vec{L}), Map-task (f), a grouping, and Reduce-task (g) are specified

1. Partition input data (\vec{L}) into chunks across all machines in the cluster
2. Apply **Map** (f) to each element
3. Shuffle and Repartition or Group Data
4. Apply **Reduce** (g) to each group
5. Collect all of the results and write to disk

All of the steps in between can be written once in a robust, safe manner and then used for every task which can be described using this MapReduce paradigm.

These tasks $\langle \vec{L}, f(x), g(a, b) \rangle$ is referred to as a job.

0.5.4 Key-Value Pairs / Grouping

The initial job was very basic, for more complicated jobs, a new notion of Key-value (KV) pairs must be introduced.

A KV pair is made up of a key and value:

- A key must be comparable / hashable (a number, string, immutable list of numbers, etc) and is used for grouping data.
- The value is the associated information to this key.

0.5.5 Counting Words

Using MapReduce on a folder full of text-documents: $\vec{L} = ["Info \dots", "Expenses \dots", \dots]$

Map

is then a function f which takes in a long string and returns a list of all of the words (text separated by spaces) as key-value pairs with the value being the number of times that word appeared

```
f(x) = [(word,1) for word in x.split(" ")]
```

Grouping is then performed by keys (group all words together)

Reduce

adds up the values for each word

Workflow

```
L = ["cat dog car",  
     "dog car dog"]
```

↓ **Map** : $f(x)$

```
[("cat",1), ("dog",1), ("car",1), ("dog",1), ("car",1), ("dog",1)]
```

↓ Shuffle / Group

```
"dog": (1,1,1)  
"car": (1,1)````  
$$ \downarrow \text{Reduce } : g(a,b) $$  
````[("cat",1), ("dog",3), ("car",2)]````
```

### Word Count Example

Here we make a word count example using all the lines of Shakespeare's "A midsummer-night's dream"

```
import os
shake_path = 'data/shakespeare.txt'
with open(shake_path, 'r') as f:
 all_lines = f.readlines()

print(all_lines[:5])
```

```
["A MIDSUMMER-NIGHT'S DREAM\n", '\n', 'Now , fair Hippolyta , our nuptial hour \n',
↪ 'Draws on apace : four happy days bring in \n', 'Another moon ; but O !_\n',
↪ 'methinks how slow \n']
```



### MapReduce Approach

Here we use the Map Reduce approach to divide the function up into Map and Reduce components

First we need a function to convert lines to words:

```
import doctest
import copy
import functools
tests are very important for map reduce

def autotest(func):
 globs = copy.copy(globals())
 globs.update({func.__name__: func})
 doctest.run_docstring_examples(
 func, globs, verbose=True, name=func.__name__)
 return func
```

```
map function
@autotest
def line_to_words(in_line):
 """
 Takes a single line and returns the words and counts
 >>> line_to_words("hi i am. bob . ")
 ['hi', 'i', 'am', 'bob']
 """
 words = in_line.lower().strip().split(' ')
 v_words = [''.join([c for c in c_word if c in string.ascii_lowercase])
 for c_word in words]
 return [c_word for c_word in v_words if len(c_word) > 0]
```

```
Finding tests in line_to_words
Trying:
 line_to_words("hi i am. bob . ")
Expecting:
 ['hi', 'i', 'am', 'bob']
ok
```

### 0.5.6 MapReduce approach using Dask bags

Dask Bag implements operations like

- map,
- filter,
- fold,
- and groupby on collections of generic Python objects.

Execution on bags provide two benefits:

- Parallel: data is split up, allowing multiple cores or machines to execute in parallel
- Iterating: data processes lazily, allowing smooth execution of larger-than-memory data, even on a single machine within a single partition



## Build a bag for the word analysis

```
import dask.bag as dbag
line_bag = dbag.from_sequence(all_lines, partition_size=10000)
line_bag
```

```
dask.bag<from_sequence, npartitions=13>
```

```
map_output = line_bag.map(line_to_words).flatten()
map_output
```

```
dask.bag<flatten, npartitions=13>
```

```
we cheat a bit for the reduce step
reduce_output = map_output.frequencies()
top10 = reduce_output.topk(10, lambda x: x[1])
bot10 = reduce_output.topk(10, lambda x: -x[1])
```

## Run the analysis

```
import dask.diagnostics as diag

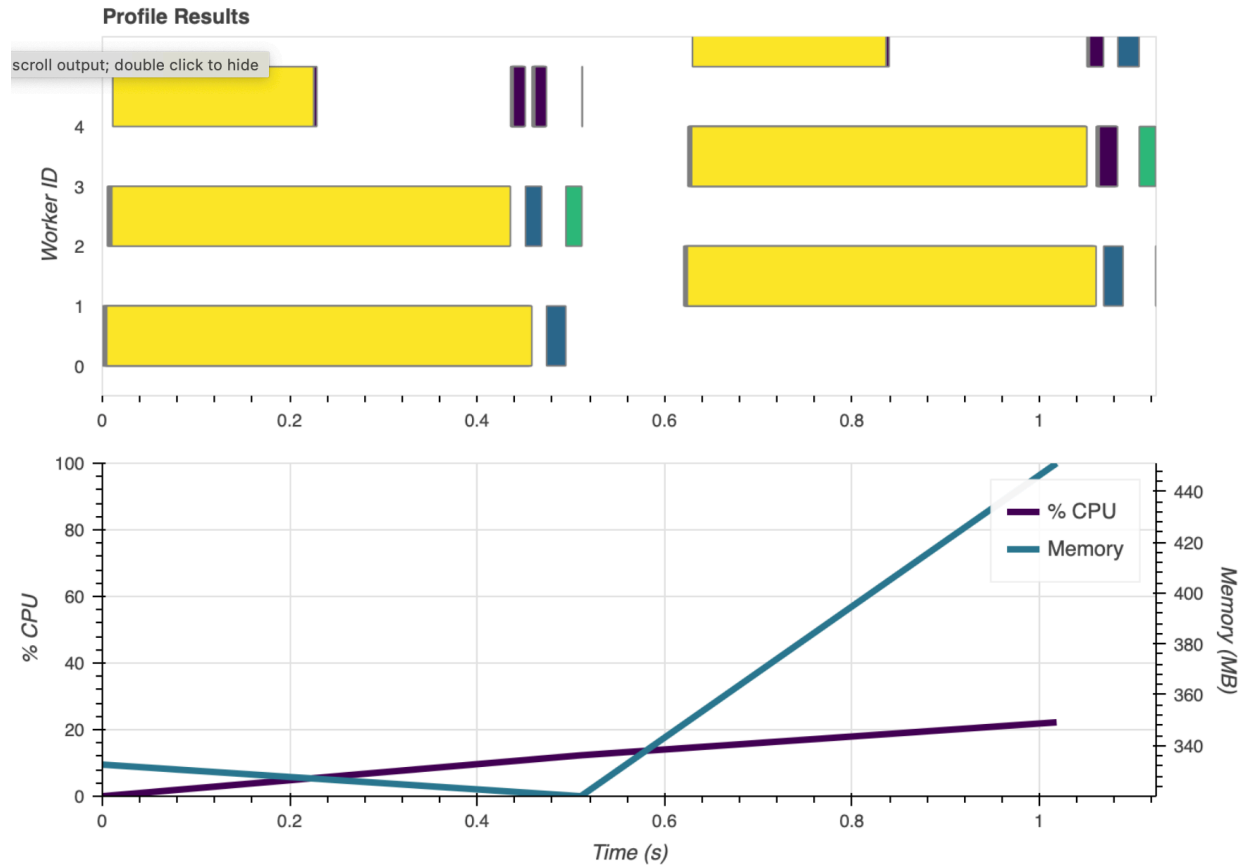
workers = 10

with diag.ProgressBar(), diag.Profiler() as prof, diag.ResourceProfiler(0.5) as rprof:
 print('Top 10\n', top10.compute(num_workers=workers))
 print('Bottom 10\n', bot10.compute(num_workers=workers))
```

```
[#####] | 100% Completed | 639.72 ms
Top 10
[('the', 26851), ('and', 24077), ('i', 20535), ('to', 18561), ('of', 16013), ('you
↵', 13856), ('a', 13840), ('my', 12282), ('that', 10761), ('in', 10537)]
[#####] | 100% Completed | 568.60 ms
Bottom 10
[('midsummernights', 1), ('wanes', 1), ('newbent', 1), ('solemnities', 1), (
↵ 'merriments', 1), ('interchangd', 1), ('lovetokens', 1), ('prevailment', 1), (
↵ 'unhardend', 1), ('filchd', 1)]
```

## Visualize parallel task distribution

```
diag.visualize([prof, rprof]);
```



```
diag.visualize([prof, rprof]);
```

## 0.6 Environments for distributed computing

### 0.6.1 Hadoop

Hadoop is the opensource version of MapReduce developed by Yahoo and released as an Apache project.

It provides underlying infrastructure and filesystem that handles storing and distributing data so each machine stores some of the data locally and processing jobs run where the data is stored.

- Non-local data is copied over the network.
- Storage is automatically expanded with processing power.
- It's how Amazon, Microsoft, Yahoo, Facebook, ... deal with exabytes of data

## 0.6.2 Spark / Resilient Distributed Datasets

### Technical Specifications

- Developed by the Algorithms, Machines, and People Lab at UC Berkeley in 2012
- General tool for all Directed Acyclical Graph (DAG) workflows
- Course-grained processing → simple operations applied to entire sets
  - Map, reduce, join, group by, fold, foreach, filter,...
- In-memory caching

Zaharia, M., et. al (2012). Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing

### Practical Specification

- Distributed, parallel computing without **logistics**, libraries, or compiling
- Declarative rather than imperative
  - Apply operation  $f$  to each image / block
  - **NOT** tell computer 3 to wait for an image from computer 2 to and perform operation  $f$  and send it to computer 1
  - Even scheduling is handled automatically
- Results can be stored in memory, on disk, redundant or not

## 0.6.3 Dask

In the pure python ecosystem, there has been a recent development called Dask.

Dask aims to bring

- the fault-tolerant,
- robust distributed computing

to numerical python codes.

In particular the focus has been on taking libraries like numpy and scipy and making them run as easily as possible in a distributed setting.

We will use these for the examples but they can be applied equally well to Spark and Hadoop-like problems.

## 0.6.4 Directed acyclical graphs - DAGs

More general than the MapReduce structure is the idea of making directed acyclical graphs.

These are used in

- Spark,
- Dask for distributed computing
- and in Tensorflow and PyTorch

for massively parallel computing since it allows complex operations to be defined in a declarative way.

This allows them to be optimized later depending on the actual resources available (and re-executed if some of those resources crash).

### Some DAG resources

- PyData Dask - <https://dask.pydata.org/en/latest/>
- Apache Spark - <https://spark.apache.org/>
- Spotify Luigi - <https://github.com/spotify/luigi>
- Airflow - <https://airflow.apache.org/>
- KNIME - <https://www.knime.com/>
- Google Tensorflow - <https://www.tensorflow.org/>
- Pytorch / Torch - <http://pytorch.org/>

### 0.6.5 Tensor Comprehensions

Facebook shows an example of why such representations are useful since they allow for the operations to be optimized later and massive performance improvements even for *fairly* basic operations.

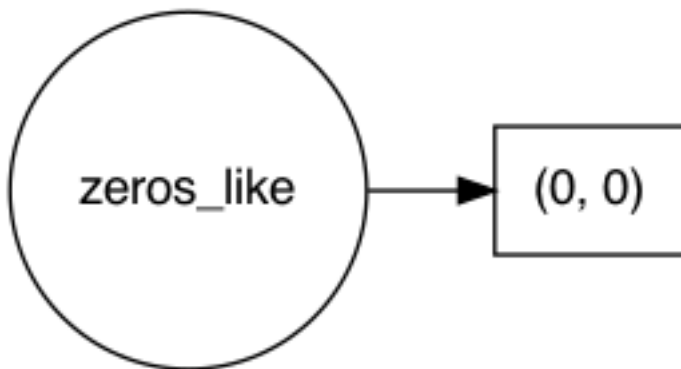
### 0.6.6 DAG examples

#### A basic DAG

Create two 5x5 images and use a single chunk:

- image1 all elements = 0
- image2 all elements = 1

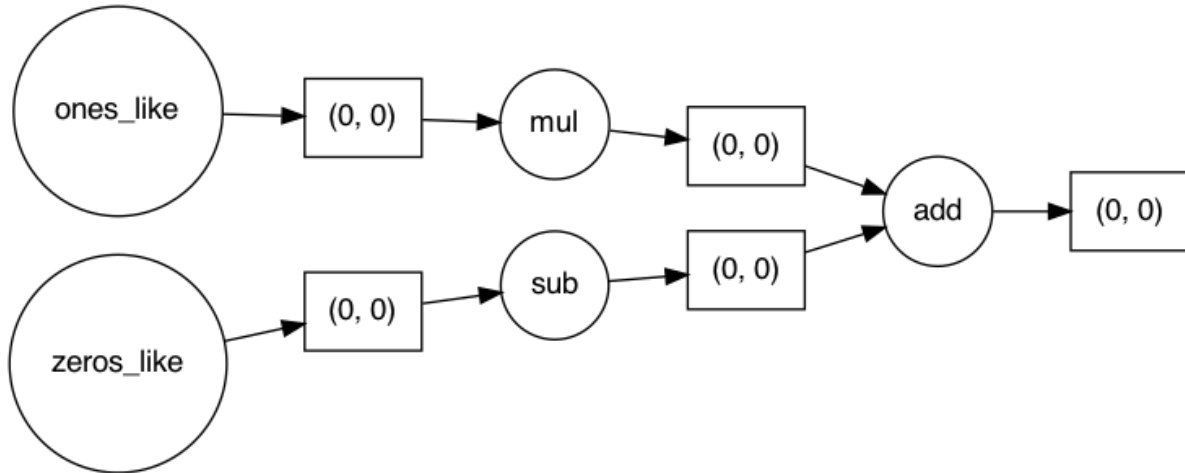
```
import dask.array as da
from dask.dot import dot_graph
image_1 = da.zeros((5, 5), chunks=(5, 5))
image_2 = da.ones((5, 5), chunks=(5, 5))
dot_graph(image_1.dask, rankdir="LR")
```



## Image arithmetics

We want to compute:  $image_4 = (image_1 - 10) + (image_2 * 50)$

```
image_4 = (image_1-10) + (image_2*50)
dot_graph(image_4.dask, rankdir="LR")
```



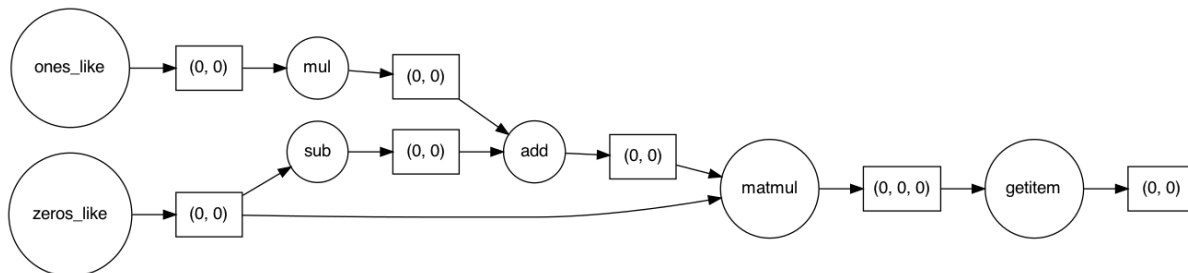
## More calculations

$$image_5 = image_1 * image_4$$

## Remember

$$image_4 = (image_1 - 10) + (image_2 * 50)$$

```
image_5 = da.matmul(image_1, image_4)
dot_graph(image_5.dask, rankdir="LR")
```



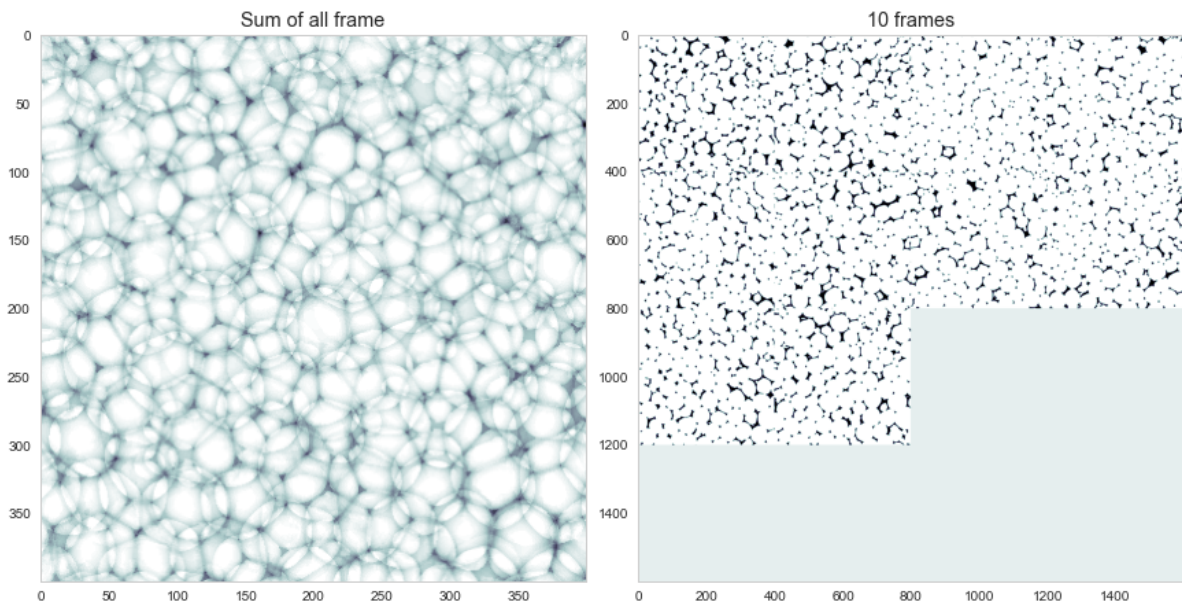
## 0.6.7 Image Processing using DAGs

the initial examples were shown on very simple image problems.

Here we can see how it looks for real imaging issues.

```
%matplotlib inline
import dask.array as da
from dask.dot import dot_graph
import numpy as np
from skimage.io import imread
from skimage.util import montage as montage2d
for showing results
import dask.diagnostics as diag

foam_stack = imread('data/plateau_border.tif')[:-54, 52:-52, 52:-52]
fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12,6))
ax1.imshow(np.sum(foam_stack, 0), cmap='bone_r', ax1.set_title('Sum of all frame'))
ax2.imshow(montage2d(foam_stack[::10]), cmap='bone_r', ax2.set_title('10 frames'));
```



### Visualizing the data as 3D rendering

```
foam_stack.dtype
```

```
dtype('uint8')
```

```
import pyvista as pv
img = foam_stack.astype(int)
img = np.random.normal(size=[100,100,100])
grid = pv.UniformGrid()

grid.dimensions = np.array(img.shape) + 1
```

(continues on next page)

(continued from previous page)

```

grid.origin = (0, 0, 0) # The bottom left corner of the data set
grid.spacing = (1, 1, 1) # These are the cell sizes along each axis, i.e., the pixel_
↳size

grid.cell_data["values"] = img.flatten(order="F") # Flatten the array!
slices = grid.slice_orthogonal(x=img.shape[0]//2, z=img.shape[2]//2)
cpos = [
 (540.9115516905358, -617.1912234499737, 180.5084853429126),
 (128.31920055083387, 126.4977720785509, 111.77682599082095),
 (-0.1065160140819035, 0.032750075477590124, 0.9937714884722322),
]
dargs = dict(cmap='plasma')

p = pv.Plotter()
p.add_mesh(slices, **dargs)

p.show()

```

```

AttributeError Traceback (most recent call last)
Cell In[18], line 11
 8 grid.origin = (0, 0, 0) # The bottom left corner of the data set
 9 grid.spacing = (1, 1, 1) # These are the cell sizes along each axis, i.e.,
↳ the pixel size
----> 11 grid.cell_data["values"] = img.flatten(order="F") # Flatten the array!
 12 slices = grid.slice_orthogonal(x=img.shape[0]//2, z=img.shape[2]//2)
 13 cpos = [
 14 (540.9115516905358, -617.1912234499737, 180.5084853429126),
 15 (128.31920055083387, 126.4977720785509, 111.77682599082095),
 16 (-0.1065160140819035, 0.032750075477590124, 0.9937714884722322),
 17]

File ~/miniconda3/lib/python3.9/site-packages/pyvista/core/datasetattributes.
↳py:223, in DataSetAttributes.__setitem__(self, key, value)
 220 raise TypeError('Only strings are valid keys for DataSetAttributes.')
 222 has_arr = key in self
--> 223 self.set_array(value, name=key)
 225 # do not make array active if it already exists. This covers
 226 # an inplace update like self.point_data[key] += 1
 227 if has_arr:

File ~/miniconda3/lib/python3.9/site-packages/pyvista/core/datasetattributes.
↳py:623, in DataSetAttributes.set_array(self, data, name, deep_copy)
 620 if not isinstance(name, str):
 621 raise TypeError('`name` must be a string')
--> 623 vtk_arr = self._prepare_array(data, name, deep_copy)
 624 self.VTKObject.AddArray(vtk_arr)
 625 self.VTKObject.Modified()

File ~/miniconda3/lib/python3.9/site-packages/pyvista/core/datasetattributes.
↳py:873, in DataSetAttributes._prepare_array(self, data, name, deep_copy)
 868 # this handles the case when an input array is directly added to the
 869 # output. We want to make sure that the array added to the output is not
 870 # referring to the input dataset.

```

(continues on next page)

(continued from previous page)

```

 871 copy = pyvista_ndarray(data)
--> 873 return helpers.convert_array(copy, name, deep=deep_copy)

File ~/miniconda3/lib/python3.9/site-packages/pyvista/utilities/helpers.py:176, in
↳ convert_array(arr, name, deep, array_type)
 173 else:
 174 # This will handle numerical data
 175 arr = np.ascontiguousarray(arr)
--> 176 vtk_data = _vtk.numpy_to_vtk(num_array=arr, deep=deep, array_
↳ type=array_type)
 177 if isinstance(name, str):
 178 vtk_data.SetName(name)

File ~/miniconda3/lib/python3.9/site-packages/vtkmodules/util/numpy_support.py:164,
↳ in numpy_to_vtk(num_array, deep, array_type)
 161 result_array.SetNumberOfTuples(shape[0])
 163 # Ravel the array appropriately.
--> 164 arr_dtype = get_numpy_array_type(vtk_typecode)
 165 if numpy.issubdtype(z.dtype, arr_dtype) or \
 166 z.dtype == numpy.dtype(arr_dtype):
 167 z_flat = numpy.ravel(z)

File ~/miniconda3/lib/python3.9/site-packages/vtkmodules/util/numpy_support.py:94,
↳ in get_numpy_array_type(vtk_array_type)
 92 def get_numpy_array_type(vtk_array_type):
 93 """Returns a numpy array typecode given a VTK array type."""
--> 94 return get_vtk_to_numpy_typemap()[vtk_array_type]

File ~/miniconda3/lib/python3.9/site-packages/vtkmodules/util/numpy_support.py:74,
↳ in get_vtk_to_numpy_typemap()
 72 def get_vtk_to_numpy_typemap():
 73 """Returns the VTK array type to numpy array type mapping."""
--> 74 _vtk_np = {vtkConstants.VTK_BIT:numpy.bool,
 75 vtkConstants.VTK_CHAR:numpy.int8,
 76 vtkConstants.VTK_SIGNED_CHAR:numpy.int8,
 77 vtkConstants.VTK_UNSIGNED_CHAR:numpy.uint8,
 78 vtkConstants.VTK_SHORT:numpy.int16,
 79 vtkConstants.VTK_UNSIGNED_SHORT:numpy.uint16,
 80 vtkConstants.VTK_INT:numpy.int32,
 81 vtkConstants.VTK_UNSIGNED_INT:numpy.uint32,
 82 vtkConstants.VTK_LONG:LONG_TYPE_CODE,
 83 vtkConstants.VTK_LONG_LONG:numpy.int64,
 84 vtkConstants.VTK_UNSIGNED_LONG:ULONG_TYPE_CODE,
 85 vtkConstants.VTK_UNSIGNED_LONG_LONG:numpy.uint64,
 86 vtkConstants.VTK_ID_TYPE:ID_TYPE_CODE,
 87 vtkConstants.VTK_FLOAT:numpy.float32,
 88 vtkConstants.VTK_DOUBLE:numpy.float64}
 89 return _vtk_np

File ~/miniconda3/lib/python3.9/site-packages/numpy/__init__.py:305, in __getattr__
↳ (attr)
 300 warnings.warn(
 301 f"In the future `np.{attr}` will be defined as the "
 302 "corresponding NumPy scalar.", FutureWarning, stacklevel=2)
 304 if attr in __former_attrs__:
--> 305 raise AttributeError(__former_attrs__[attr])

```

(continues on next page)



(continued from previous page)

```

307 # Importing Tester requires importing all of unittest which is not a
308 # cheap import since it is mainly used in test suits, we lazy import it
309 # here to save on the order of 10 ms of import time for most users
310 #
311 # The previous way Tester was imported also had a side effect of adding
312 # the full `numpy.testing` namespace
313 if attr == 'testing':

```

AttributeError: module 'numpy' has no attribute 'bool'.  
`np.bool` was a deprecated alias for the builtin `bool`. To avoid this error in existing code, use `bool` by itself. Doing this will not modify any behavior and is safe. If you specifically wanted the numpy scalar type, use `np.bool\_` here. The aliases were originally deprecated in NumPy 1.20; for more details and guidance see the original release note at:  
<https://numpy.org/devdocs/release/1.20.0-notes.html#deprecations>

```
import pyvista as pv
```

```

p1 = pv.Plotter()
va = p1.add_volume(foam_stack, cmap='viridis')
f = lambda val: va.GetProperty().SetScalarOpacityUnitDistance(val)
p1.add_slider_widget(f, [-0.1, 1], title="Opacity Distance")
p1.show()

```

```

from itkwidgets import view
from IPython.display import display
import itk

```

```

if True:
 viewer = view(itk.GetImageFromArray(foam_stack.copy()))
 display(viewer)

```


## 0.6.8 Create a DAG with the foam image

- The image is 100x400x400
- We want 20 slices per chunk
- Scale intensities to from [0,255] to [0,1]

```

da_foam = da.from_array(
 foam_stack/255.0, chunks=(20, 400, 400), name='FoamImage')
dot_graph(da_foam.dask)

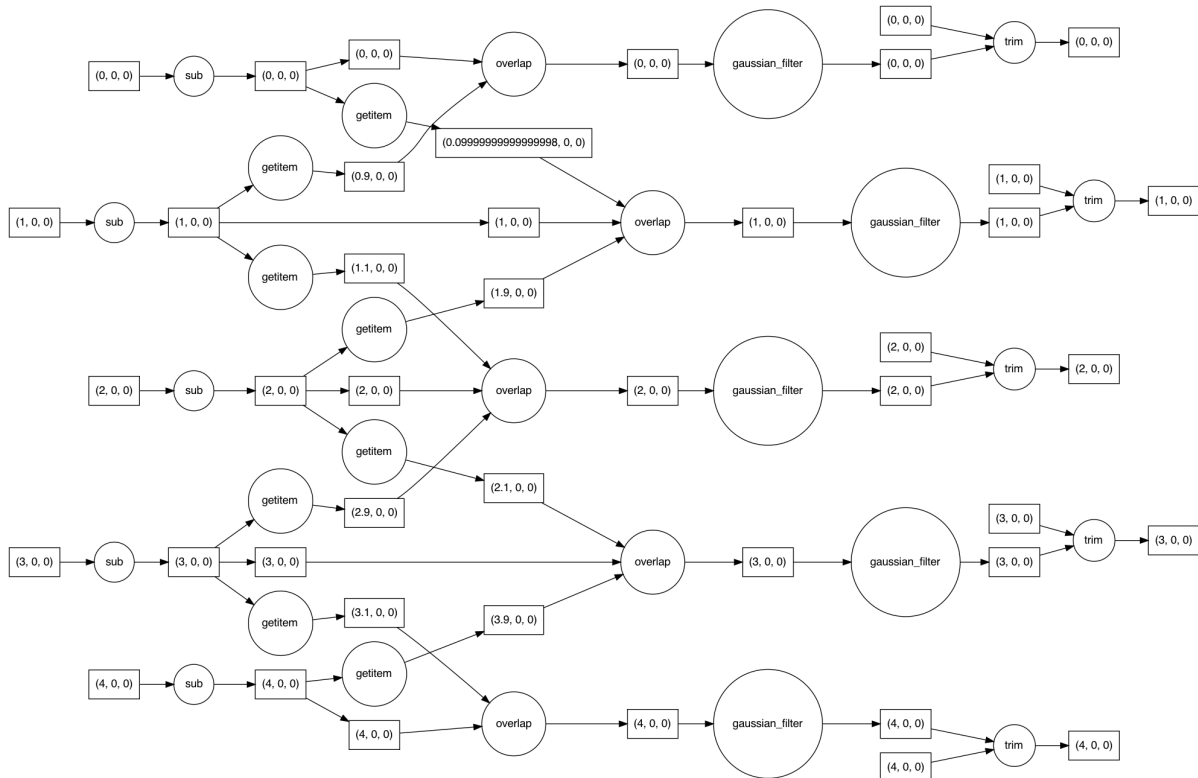
```



### Add filter operation

- We want to use the Dask version of `scipy.ndfilter.gaussian_filter`

```
import dask_ndfilters as da_ndfilt
image_filt = da_ndfilt.gaussian_filter(1-da_foam, sigma=(3, 6, 6))
dot_graph(image_filt.dask, rankdir="LR")
```



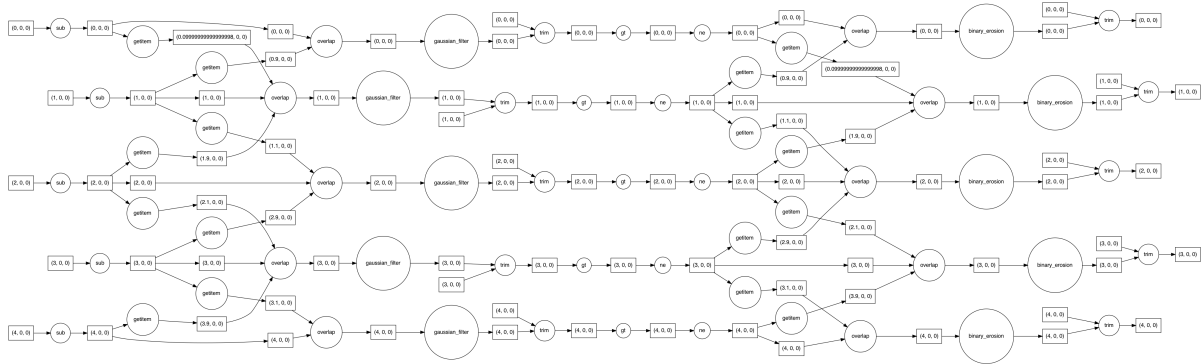
### Why so complicated?

The filter needs to process the boundaries correctly  
 This requires an exchange of boundary data between the blocks.

### Add segment and erode to the DAG

- Apply a threshold at 0.9
- Erode with a ball structure element with radius 12

```
#import scipy.ndimage.morphology as ndmorph
import dask_ndmorph as ndmorph
from skimage.morphology import ball
erode_foam = ndmorph.binary_erosion(image_filt > 0.9, ball(12))
dot_graph(erode_foam.dask, rankdir='LR')
```



## Label the items in the image

- Create a label function
- Labels must be globally unique

```
from scipy.ndimage import label

def block_label(in_block, block_id=None):
 slice_no = block_id[0]
 offset = (np.prod(in_block.shape)*slice_no).astype(np.int64)
 label_img = label(in_block)[0].astype(np.int64)
 label_img[label_img > 0] += offset
 return label_img

lab_bubbles = erode_foam.map_blocks(block_label, dtype='int64')
```

## Run the labelling DAG

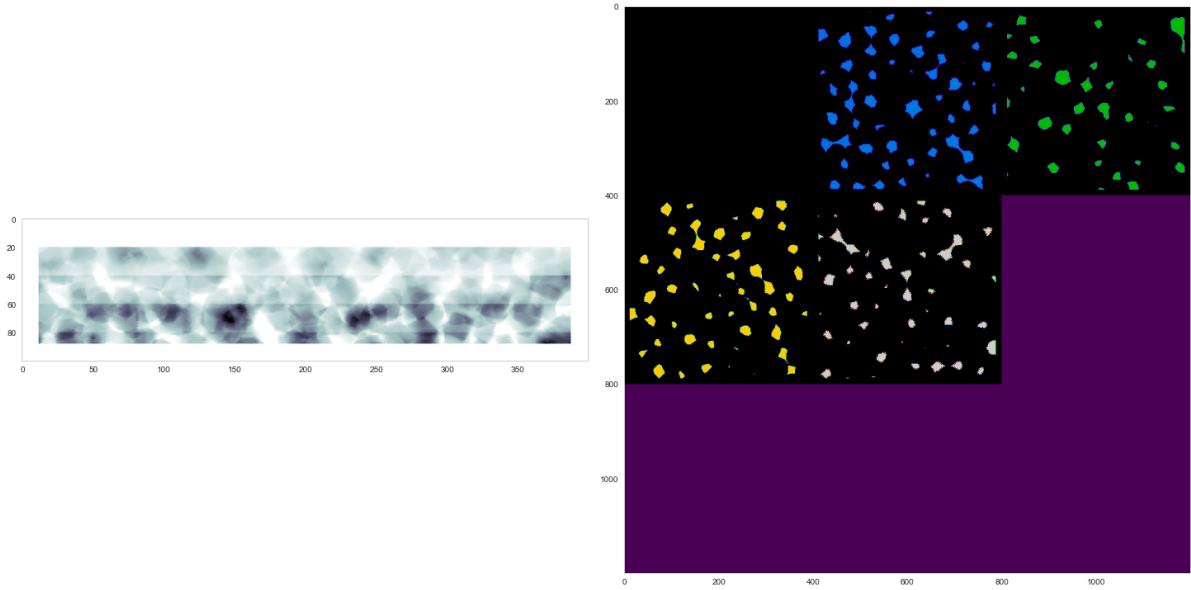
The labelling is executed with a profiler to check the timing

```
with diag.ProgressBar(), diag.Profiler() as prof, diag.ResourceProfiler(0.5) as rprof:
 processed_stack = lab_bubbles.compute(num_workers=4)
```

```
[#####] | 100% Completed | 8.70 sms
```

## Results of running the DAG

```
fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(20, 10))
ax1.imshow(np.sum(processed_stack, 1), cmap='bone_r')
ax2.imshow(montage2d(processed_stack[:, :20]), cmap='nipy_spectral');
```



### Profiler output

```
diag.visualize([prof, rprof]);
```

## 0.6.9 The importance of operation order

### Select a slice and filter

#### Filter and select a slice

```
foam_slices_da = da.from_array(
 foam_stack/255.0, chunks=(10, 500, 500), name='FoamSlices')


filt_slices = da_ndfilt.gaussian_filter(foam_slices_da, sigma=(1.0, 9, 9))
single_slice = filt_slices[50]

single_slice.visualize(filename="singleslice.svg", optimize_graph=True);
filt_slices.visualize(filename="filt_slices.svg" , optimize_graph=True);
```

### Compare performance

```
fig, (ax1,ax2) = plt.subplots(1,2,figsize=(12,5))
with diag.ProgressBar():
 ax1.imshow(single_slice.compute(num_workers = 4))
 ax2.imshow(filt_slices.compute(num_workers = 4)[50])
```

```
[#####] | 100% Completed | 207.44 ms
[#####] | 100% Completed | 417.62 ms
```



Same output but different timing!

What happened?

Select a slice and filter

Filter and select a slice

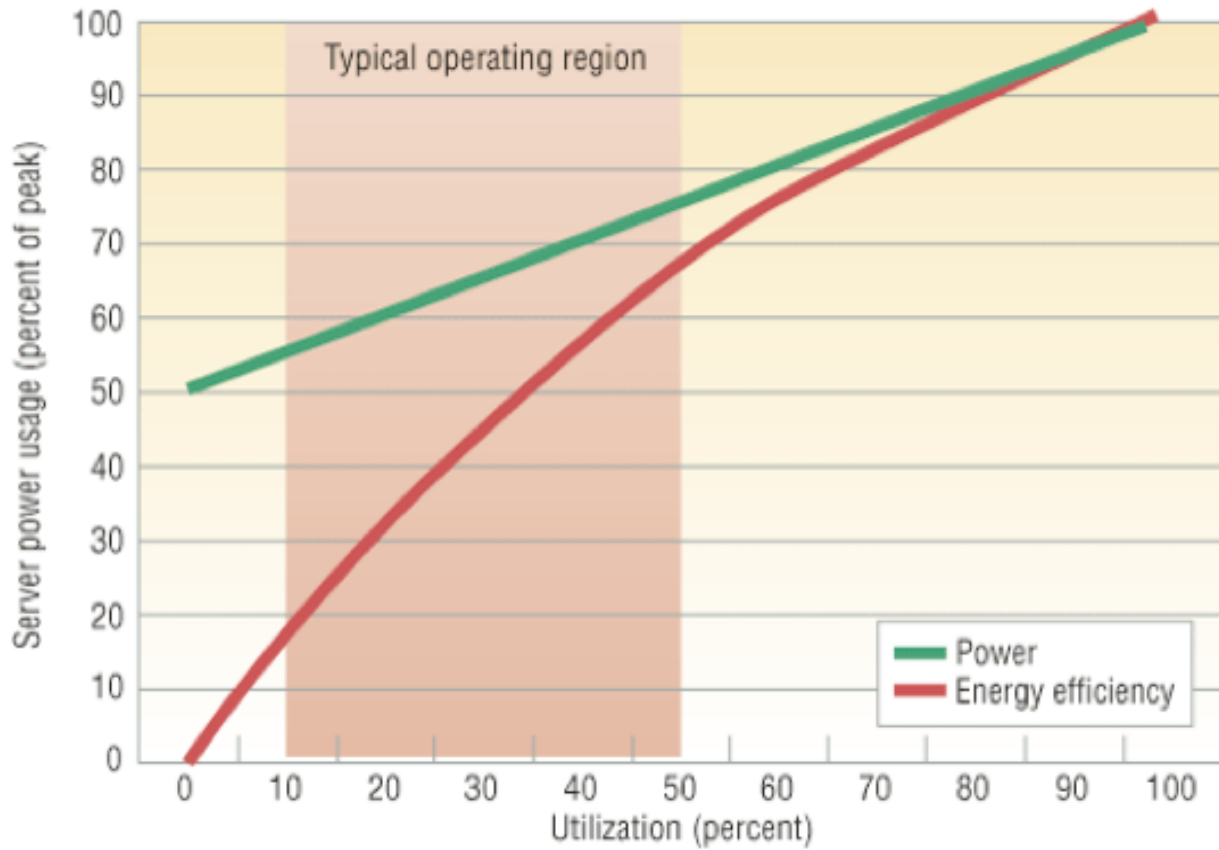
- A single slice only activates the chunks needed (lazy evaluation)
- 3 Chunks are processed
- The image is divided into 10 chunks
- All chunks are processed
- The slice is taken from the result

## 0.7 Cloud computing

### 0.7.1 Motivation for Cloud Computing

Local resources

- Local resources are expensive and underutilized
- Management and updates are expensive and require dedicated IT staff



### Cloud Resources

- Automatically setup
- “Unlimited” potential capacity and storage
- Cluster management already setup
- Common tools with many people using the same

### 0.7.2 Spark - A rich, heavily developed platform

#### Available Tools

Tools built for table-like data data structures and much better adapted to it.

- K-Means,
- Matrix Factorization, Genomics, Graph Analytics, Machine Learning

## Commercial Support

Dozens of major companies (Apple, Google, Facebook, Cisco, ...) donate over \$30M a year to development of Spark and the Berkeley Data Analytics Stack

- 2 startups in the last 6 months with seed-funding in excess of \$15M each

## Academic Support

- All source code is available on GitHub
- Elegant (20,000 lines vs my PhD of 75,000+)
- No patents or restrictions on usage
- Machine Learning Course in D-INFK next semester based on Spark

## 0.7.3 Beyond: Streaming

### Post-processing goals

- Analysis done in weeks instead of months
- Some real-time analysis and statistics

### Streaming

- Can handle static data
- or live data coming in from a 'streaming' device like a camera to do real-time analysis.

The exact same code can be used for real-time analysis and static code

### Scalability

- Connect more computers.
- Start workers on these computer.

## 0.7.4 Beyond: Approximate Results

Projects at AMPLab like Spark and BlinkDB are moving towards approximate results.

- Instead of `mean(volume)`
- `mean(volume).within_time(5)`
- `mean(volume).within_ci(0.95)`

For real-time image processing it might be the only feasible solution and could drastically reduce the amount of time spent on analysis.

## 0.8 Summary

- More can be done in parallel
- Limited resources, redundancy, race conditions
- Cloud computing
- DAGs