




advestis
one step ahead

JCAD2022



Execution of a feature
engineering algorithm in Python
on a supercomputer with Dask

Context

Introduction to the problematic : derivation

Dask

Derivation toy model

Chaining futures

Dask on Myria

CRIANN Presentation

Supercomputer *Myria*

Dask on Myria

Conclusion

Raw Data Selection



Feature engineering



Model Training



Recommandation



Visual Restitution



Example : ESG scores

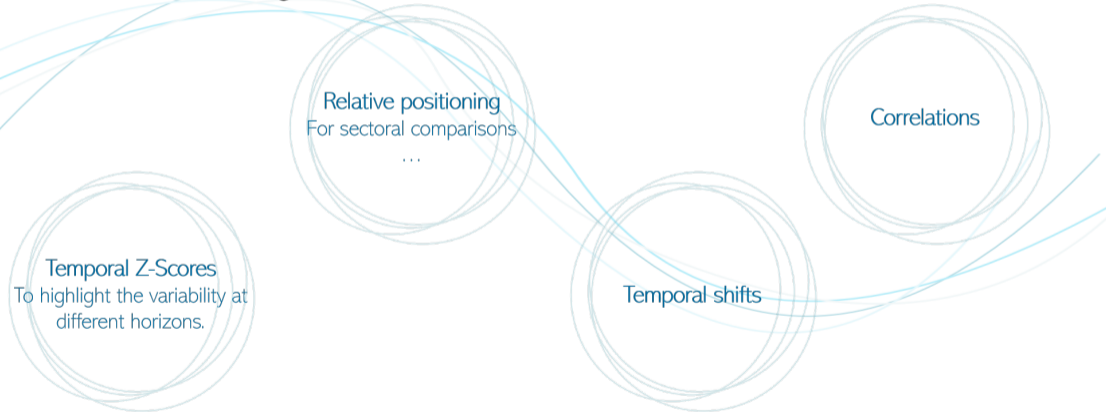
1000 shares

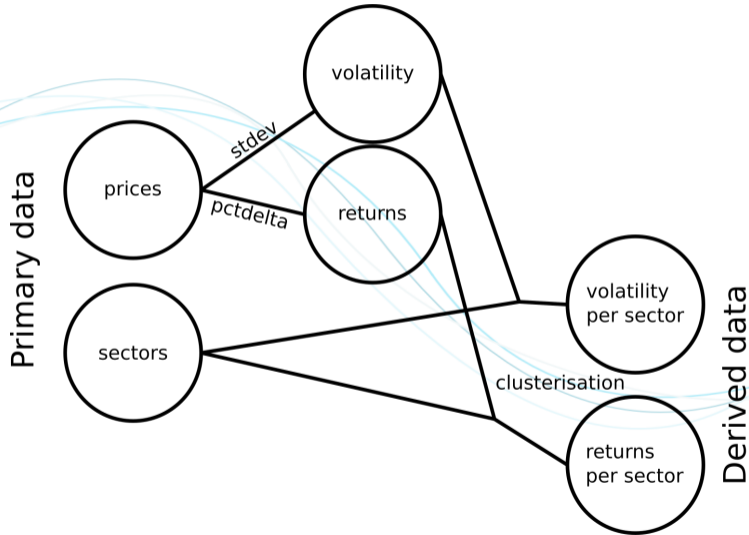
10 years (2600 business days)

	Share 1	Share 2	Share 3	Share 4	Share 5	Share 6	Share 7	Share 8
2010-09-30	1	2	1	1	0	0	1	
2010-10-01	1	2	1	1	0	0	1	
2010-10-04	1	2	1	1	0	0	1	
2010-10-05	1	2	1	1	0	0	1	
2010-10-06	1	2	1	1	0	0	1	
2010-10-07	1	2	1	1	0	0	1	
2010-10-08	1	2	1	1	0	0	1	
2010-10-11	1	2	1	1	0	0	1	
2010-10-12	1	2	1	1	0	0	1	
2010-10-13	1	2	1	1	0	0	1	
2010-10-14	1	2	1	1	0	0	1	
2010-10-15	1	2	1	1	0	0	1	
2010-10-18	1	2	1	1	0	0	1	
2010-10-19	1	2	1	1	0	0	1	
2010-10-20	1	2	1	1	0	0	1	
2010-10-21	1	2	1	1	0	0	1	
2010-10-22	1	2	1	1	0	0	1	
2010-10-25	1	2	1	1	0	0	1	
2010-10-26	1	2	1	1	0	0	1	
2010-10-27	1	2	1	1	0	0	1	
2010-10-28	1	2	1	1	0	0	1	

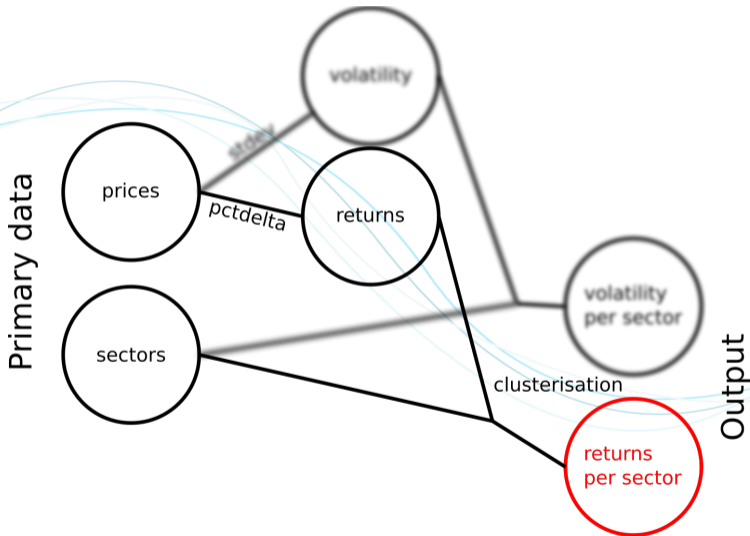
x276 of such dataframes (scores, prices, sectors, marketcap ...) >700x10⁶ cells

A set of transformations/combinations called "derivations" is applied to the primary data in order to feed the learnings.





Ex. : compute only the returns



Dask

- Python library for parallel data processing
- Fully integrated in the PyData Stack (NumPy, SciPy, Scikit-Learn, Pandas, ...)

Main parts

1. Dynamic task scheduling optimized for computation
2. Big Data collections for parallel data management

How dask work?

- Processing instructions on data generates a task graph
- The graph is scheduled by a cluster of workers
- Several solutions to manage the cluster: python multiprocessing (for single node only), k8s cluster, job scheduler like Slurm or PBS, MPI communicator, ...

Unary operator

moving
average



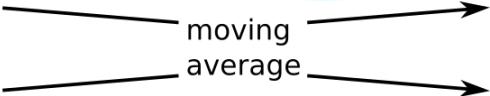
generic operator


correl.



Map

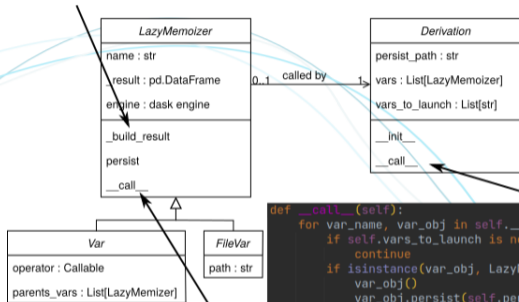
moving
average





```
1 class Basic(AbstractDerivation):
2     """The basic class to read data from data base"""
3
4     all_prices = FileVar('all_prices.csv')
5     marketcap = FileVar('marketcap.csv')
6     sectors_isin = FileVar('Sectors.csv')
7
8
9 class Temporary(AbstractDerivation):
10    """Temporary data for compute Y and X and others for the trading"""
11    fin_rdt_1d_all = Var(op.pctdelta, Basic.all_prices, {'period': 1})
12    fin_rdt_1d_std = Var(op.ewmstd, fin_rdt_1d_all, {'span': 260, 'min_periods': 260,
13                                                    'annulize': True})
14
15    clustervol = Var(op.get_clusters, fin_rdt_1d_std, {'levels': 5,
16                                                    'dend': DATE_END_CLUSTERIZATION})
17
18    fin_pricez1d = Var(op.zscore, Basic.all_prices, {'h': 1, 'minp': 60, 'inpct': True})
19    fin_pricedevz1d_byvol = Var(op.calc_dev_by, [fin_pricez1d, clustervol],
20                               {'do_reduce': False, 'clusterishito': True})
21    fin_std260d_Z1d_bysectors = Var(op.calc_dev_by, [fin_pricez1d, Basic.sectors_isin],
22                                   {'do_reduce': False})
23    fin_pricedevz1d_byuniver = Var(op.calc_dev_by, [fin_pricez1d], {'do_reduce': False})
24    a = Var(op.wait_for_dask, fin_pricez1d, {})
```

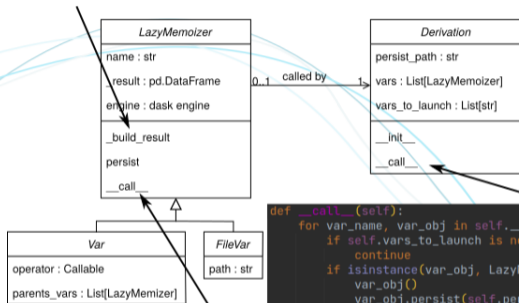
```
def _build_result(self):
    parent_inputs = list(map(self.engine.var_to_input, self.parent_vars))
    self._result.obj = self.engine.submit(self.name, self.operator, parent_inputs, self.params)
```



```
def __call__(self):
    for var_name, var_obj in self.__get_vars():
        if self.vars_to_launch is not None and var_name not in self.vars_to_launch:
            continue
        if isinstance(var_obj, LazyMemoizer):
            var_obj()
            var_obj.persist(self.persist_path)
```

```
def __call__(self):
    assert self.name, 'name should be set first from the derivation class'
    logger.info(f'Calling {self.name}')
    assert self._result is not None, 'self._result has not been built'
    if self._result.is_empty():
        self._build_result()
    return self._result
```

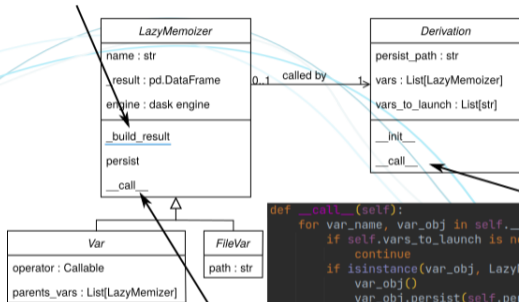
```
def _build_result(self):  
    parent_inputs = list(map(self.engine.var_to_input, self.parent_vars))  
    self._result.obj = self.engine.submit(self.name, self.operator, parent_inputs, self.params)
```



```
def __call__(self):  
    for var_name, var_obj in self.__get_vars():  
        if self.vars_to_launch is not None and var_name not in self.vars_to_launch:  
            continue  
        if isinstance(var_obj, LazyMemoizer):  
            var_obj()  
            var_obj.persist(self.persist_path)
```

```
def __call__(self):  
    assert self.name, 'name should be set first from the derivation class'  
    logger.info(f'Calling {self.name}')  
    assert self._result is not None, 'self._result has not been built'  
    if self._result.is_empty():  
        self._build_result()  
    return self._result
```

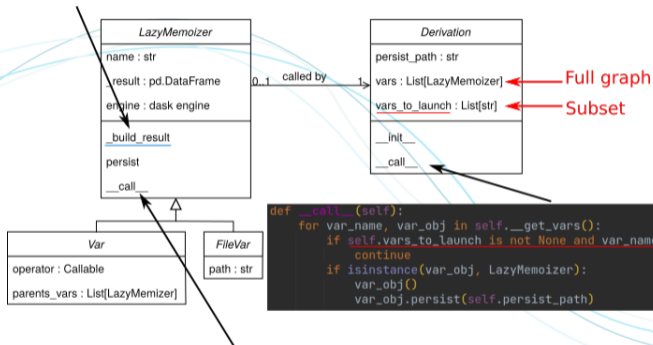
```
def _build_result(self):  
    parent_inputs = list(map(self.engine.var_to_input, self.parent_vars))  
    self._result.obj = self.engine.submit(self.name, self.operator, parent_inputs, self.params)
```



```
def _call_(self):  
    for var_name, var_obj in self._get_vars():  
        if self.vars_to_launch is not None and var_name not in self.vars_to_launch:  
            continue  
        if isinstance(var_obj, LazyMemoizer):  
            var_obj()  
            var_obj.persist(self.persist_path)
```

```
def __call__(self):  
    assert self.name, 'name should be set first from the derivation class'  
    logger.info(f'Calling {self.name}')  
    assert self._result is not None, 'self._result has not been built'  
    if self._result.is_empty():  
        self._build_result()  
    return self._result
```

```
def _build_result(self):  
    parent_inputs = list(map(self.engine.var_to_input, self.parent_vars))  
    self._result.obj = self.engine.submit(self.name, self.operator, parent_inputs, self.params)
```

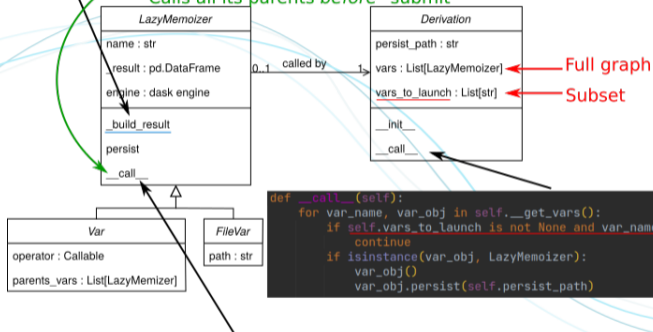


```
def __call__(self):  
    for var_name, var_obj in self.__get_vars():  
        if self.vars_to_launch is not None and var_name not in self.vars_to_launch:  
            continue  
        if isinstance(var_obj, LazyMemoizer):  
            var_obj()  
            var_obj.persist(self.persist_path)
```

```
def __call__(self):  
    assert self.name, 'name should be set first from the derivation class'  
    logger.info(f'Calling {self.name}')  
    assert self._result is not None, 'self._result has not been built'  
    if self._result.is_empty():  
        self._build_result()  
    return self._result
```

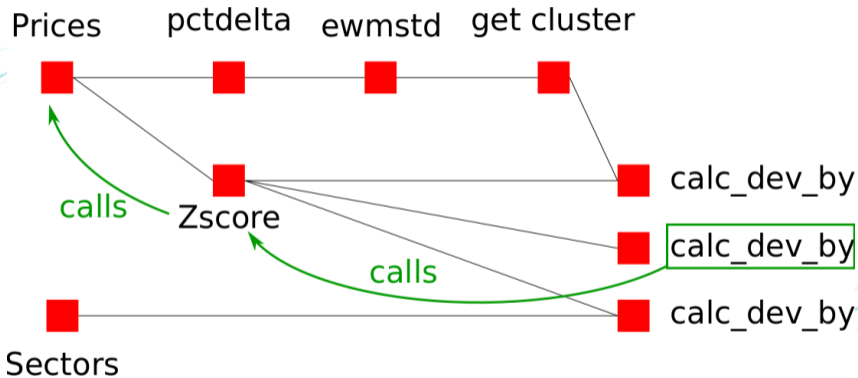
```
def _build_result(self):  
    parent_inputs = list(map(self.engine.var_to_input, self.parent_vars))  
    self._result_obj = self.engine.submit(self.name, self.operator, parent_inputs, self.params)
```

Calls all its parents before "submit"



```
def __call__(self):  
    for var_name, var_obj in self.__get_vars():  
        if self.vars_to_launch is not None and var_name not in self.vars_to_launch:  
            continue  
        if isinstance(var_obj, LazyMemoizer):  
            var_obj()  
            var_obj.persist(self.persist_path)
```

```
def __call__(self):  
    assert self.name, 'name should be set first from the derivation class'  
    logger.info(f'Calling {self.name}')  
    assert self._result is not None, 'self._result has not been built'  
    if self._result.is_empty():  
        self._build_result()  
    return self._result
```

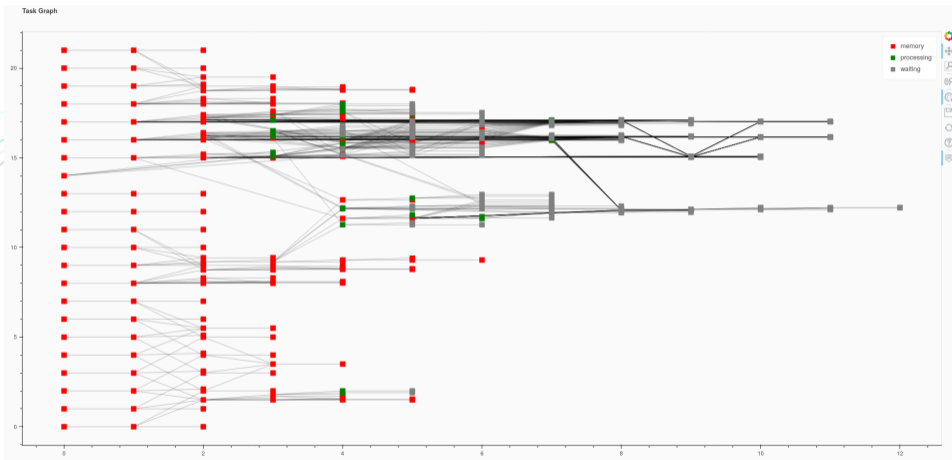
Prices



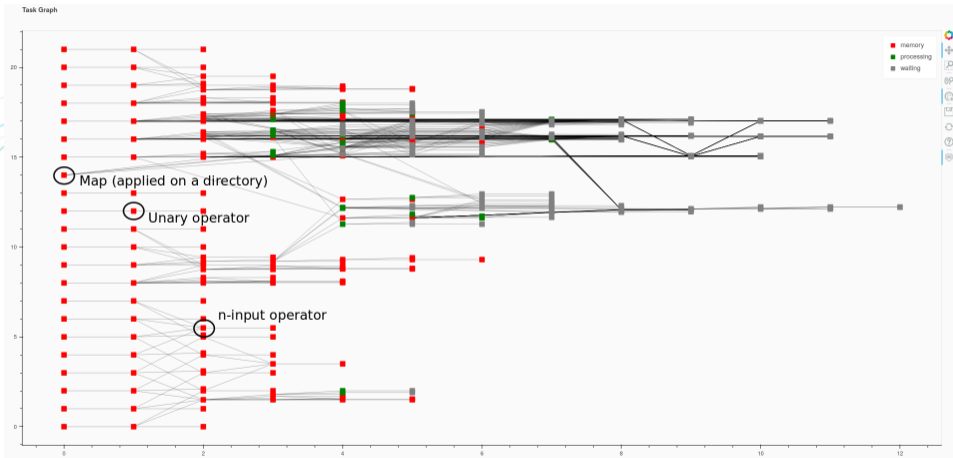
Zscore



calc_dev_by




View execution





- Created in 1991 as a Non Profit Organisation by Higher Educations Institutions & Universities
 - HPC center and regional network
- Located in Rouen Engineering Campus
- 13 employees (9 engineers or PhD)
- Funding is mainly public
 - Running costs supported by Normandy Region
 - Projects / investments funded by projects with EU, French State & Normandy Region
 - A small self-financing part
- HPC for public research, also open to private



773 TFlop/s (Peak)
(419 Xeon + 327 GPU + 27 KNL)

366 Broadwell Nodes (10248 cores)
28 cores@2.4 GHz - 128 GB RAM

Specialised nodes
- 26 GPUs (48 K80, 17 P100, 20 V100)
- 13/ I/O

SMP node Haswell
256 cores@2.2 GHz
4 To RAM DDR4

10 Xeon Phi KNL (640 cores)

Intel OmniPath
100 Gbit/s

DDN Storage
2,5 Po (HDD)

Access 4x10Gbit/s & 1x40Gbit/s
5 front-end & 3 visu. nodes

Cent OS - Slurm - GPFS

Since 2020

- Criann context: experimental works for growing competence
- Goal: service available for the next production machine (2023)

2020: 1 user project

- Techno watch for project users (workflow full redesign)
- Distributed data processing
- Mainly dask.dataframe features
- Results: efficient on a single node; dataframe partitioning inefficient in multinode

2022: Advestis

- Production code using Dask
- Distributed data processing
- Mainly dask.distributed features

Versions

- available only for privileged users
- specific version for Advestis
 - restricted access
 - in-house packages
- several versions for experimentation
 - Criann only

Dashboard unavailable
web access prohibited



LocalCluster

- Uses Python multiprocessing/multithreading to deploy dask workers

```
1 cluster = LocalCluster()  
2 client = Client(cluster)
```

- Pros
 - Easy to setup
 - In a slurm job, automatically adapt workers with the cgroups resources
- Cons
 - For single node only: no scalability

K8S Cluster

- Uses K8S cluster to deploy dask workers

```
1 from dask_kubernetes import KubeCluster  
2 cluster = KubeCluster.from_dict(self.  
    get_dask_kube_worker_spec())  
3 cluster.adapt(minimum=1, maximum=self.  
    dask_kube_max_pods)  
4 client = Client(cluster)
```

- Pros
 - Easy to scale up
 - Standard approach for Advestis
- Cons
 - Unavailable on Myria (Slurm cluster without container)

Job scheduler

- Uses job scheduler (Slurm, PBS, ...) to deploy dask workers
- Configuration file (*jobqueue.yaml*) specific to the cluster
- Pros
 - Fully integrated with HPC clusters
 - Easy to scale up
- Cons
 - Dask workers run in jobs independent of the main Python process: if resources are busy the main process can wait a long time
 - Slurm jobs submitted by slurm jobs: beware side effects

jobqueue.yaml

```
1 jobqueue:  
2   slurm:  
3     cores: 28  
4     memory: '120GB'  
5     processes: 7  
6     queue: 'tcourt_intra'  
7     walltime: '02:00:00'
```

Main python process

```
1 from dask_jobqueue import SLURMCluster  
2 cluster = SLURMCluster()
```

MPI Cluster (dask_mpi)

- Uses MPI to deploy dask workers
- Workers are MPI processes
- Two ways:
 - MPI for main Python process and dask workers
 - MPI only for dask workers
- Pros
 - Easy to scale up
 - Adapted for HPC clusters
- Cons
 - Dysfunctions with real use case (import error, unknown variables)

same MPI context

```
1 from dask_mpi import initialize
2 initialize()
3 from dask.distributed import Client
4 client = Client()
```

```
1 mpirun -np 4 python my_client_script.py
```

Warning: n - 2 workers

MPI for dask workers only

```
1 mpirun -np 4 dask-mpi --scheduler-file
   scheduler.json
```

```
1 from dask.distributed import Client
2 client = Client(scheduler_file='./scheduler.
   json')
```

Can run in a heterogenous job

Computation graph ran on

- Local machine (Intel Core i9-9900K CPU @ 3.60GHz × 16, 32 GB RAM)
- Kubernetes on a Google Cloud Platform cluster (up to 50 workers, one worker is a 16-core VM)
- Myria

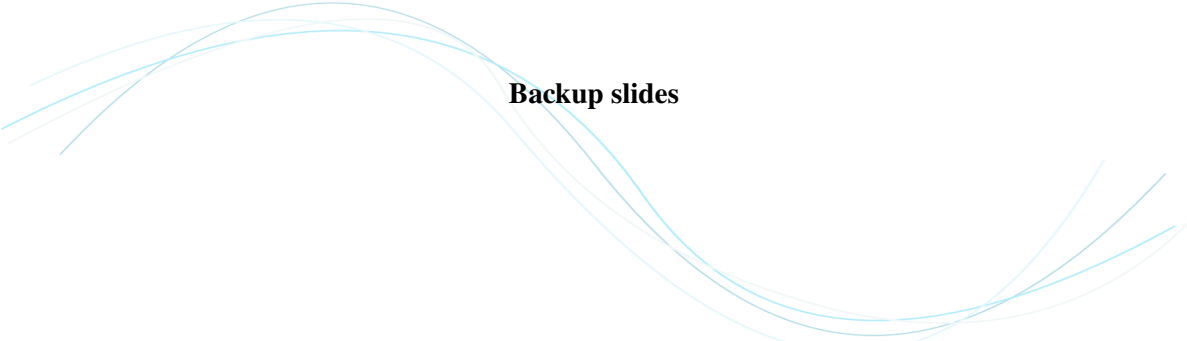
	full derivation (5327)	medium derivation (1383)
Local machine - Memory	120 minutes	20 minutes
Local machine - Dask	30 minutes	5 minutes
K8s on GCP - Dask	9 minutes	4 minutes
Myria - Memory	155 minutes	7 minutes
Myria - Dask	37 minutes	3 minutes

Computation time

Thank you !

nboumlaik@advestis.com
pcotte@advestis.com
benoist.gaston@criann.fr



The slide features several overlapping, light blue wavy lines that create a sense of motion and depth. These lines are centered around the text and extend across the width of the slide.

Backup slides

With a graph of +1000 nodes:

- How to automatically cache intermediate results to avoid running same nodes multiple times?
- How to run a subset of graph by selecting a subset of final nodes?

⇒ Dask future or Dask delayed?

Pros:

- Visualize is pretty cool
- Delayed are lazy objects. It means the complete graph may be configured but not executed. Thus it is possible to trigger a subset of this graph.

Cons:

- Persistence has to be implemented explicitly. For complex graphs it is not sufficient.

Pros:

- persistence happens as long as there is at least one python object referencing the future.

Cons:

- Future is triggered without delay
- There is no visualize method

Conclusion:

- Dask primary components are not sufficient. Thus a thin wrapper is mandatory in order to manage a lazy behavior
- We opted for **Dask Futures** because they automatically persist

```
1 class AbstractDerivation(ABC):
2     def __init__(self, persist_folder, vars_to_launch_file=None, fs=None):
3         self.persist_path = settings.DERIVED_PATH / persist_folder
4         self.vars_to_launch = None
5         if isinstance(vars_to_launch_file, Path):
6             self.vars_to_launch = vars_to_launch_file.read()
7         elif vars_to_launch_file is not None:
8             if not isinstance(vars_to_launch_file, list):
9                 vars_to_launch_file = Path(vars_to_launch_file, fs=fs)
10                self.vars_to_launch = vars_to_launch_file.read()
11            else:
12                self.vars_to_launch = vars_to_launch_file
13        if isinstance(self.vars_to_launch, str):
14            self.vars_to_launch = self.vars_to_launch.split("\n")
15        self.__add_name()
```

```
1
2 def __get_vars(self):
3     for var_name in filter(lambda x: not x.startswith('_'), dir(self)):
4         var_obj = getattr(self, var_name)
5         yield var_name, var_obj
6
7 def __add_name(self):
8     for var_name, var_obj in self.__get_vars():
9         if any(map(lambda x: isinstance(var_obj, x), [Var, Map, FileVar])):
10            logger.debug(f'adding name to {var_name}')
11            var_obj.name = var_name
12
13 def __call__(self):
14     for var_name, var_obj in self.__get_vars():
15         if self.vars_to_launch is not None and var_name not in self.vars_to_launch:
16             continue
17         if isinstance(var_obj, LazyMemoizer):
18             var_obj()
19             var_obj.persist(self.persist_path)
```

```
1 class LazyMemoizer(ABC):
2     """
3     Build a callable object and store the result through an engine
4     """
5
6     def __init__(self, *args, **kwargs):
7         self._result = None
8         self._writer = None
```

```
1     def __call__(self):
2         assert self.name, 'name should be set first from the derivation class'
3         logger.info(f'Calling {self.name}')
4         assert self._result is not None, 'self._result has not been built'
5         if self._result.is_empty():
6             self._build_result()
7         return self._result
```

```
1 class FileVar(LazyMemoizer):
2     """
3     Setup a Var object from a dataframe stored as a csv file. It is inherited
4     from Var class
5     """
6     def __init__(self, filename, reader=Reader, folder=settings.BASE_PATH):
7         """
8
9         params: filename: used be the reader to get the file
10        """
11        super().__init__()
12        self._result = Result()
13        self.reader = reader(folder)
14        self.filename = filename
15        self.__name = None
16
17        @property
18        def name(self):
19            return self.__name
20
21        @name.setter
22        def name(self, value):
23            self.__name = value
24
25        def _build_result(self):
26            self._result.obj = self.engine.submit(self.name, self.reader, dfs=(),
27                                                  params={'name': self.filename})
```

```
1 class Var(LazyMemoizer):
2     """
3     Setup a var as a transformation from another var
4     """
5
6     def __init__(self, operator, parent_vars=(), params=None):
```

```
1         super().__init__()
2         self._result = Result()
3         assert callable(operator)
4         self.operator = operator
5         if isinstance(parent_vars, Iterable):
6             self.parent_vars = parent_vars
7         else:
8             self.parent_vars = [parent_vars]
9         self.params = params if params else {}
10        self.__name = None
```

```
1     def _build_result(self):
2         parent_inputs = list(map(self.engine.var_to_input, self.parent_vars))
3         self._result.obj = self.engine.submit(self.name, self.operator, parent_inputs, self.params)
4
5     def persist(self, path):
6         writer = Writer(path)
7         self._writer = Var(writer, self, params={'name': self.name})
8         self._writer.name = f'writer_{self.name}'
9         self._writer()
```

```
1 class Basic(AbstractDerivation):
2     """The basic class to read data from data base"""
3
4     all_prices = FileVar('all_prices.csv')
5     marketcap = FileVar('marketcap.csv')
6     sectors_isin = FileVar('Sectors.csv')
7
8
9 class Temporary(AbstractDerivation):
10     """Temporary data for compute Y and X and others for the trading"""
11     fin_rdt_1d_all = Var(op.pctdelta, Basic.all_prices, {'period': 1})
12     fin_rdt_1d_std = Var(op.ewmstd, fin_rdt_1d_all, {'span': 260, 'min_periods': 260,
13                                                     'annulize': True})
14
15     clustervol = Var(op.get_clusters, fin_rdt_1d_std, {'levels': 5,
16                                                     'dend': DATE_END_CLUSTERIZATION})
17
18     fin_pricez1d = Var(op.zscore, Basic.all_prices, {'h': 1, 'minp': 60, 'inpct': True})
19     fin_pricedevz1d_byvol = Var(op.calc_dev_by, [fin_pricez1d, clustervol],
20                               {'do_reduce': False, 'clusterishito': True})
21     fin_std260d_Z1d_bysectors = Var(op.calc_dev_by, [fin_pricez1d, Basic.sectors_isin],
22                                     {'do_reduce': False})
23     fin_pricedevz1d_byuniver = Var(op.calc_dev_by, [fin_pricez1d], {'do_reduce': False})
24     a = Var(op.wait_for_dask, fin_pricez1d, {})
```

```
1 def main(vars_to_launch=None, fs=None):
2     make_directories()
3     derivations = [Basic(settings.S_FOLDER),
4                   Temporary(settings.T_FOLDER)]
5     for d in derivations:
6         d()
7     get_engine().wait()
```



```
1 import click
2
3 from adoptim3.derivation.append import update_db
4 # from adoptim3.derivation import deriv_ia_stoxx600
5 from adoptim3.derivation import deriv_toy_model
6 # from adoptim3.derivation import deriv_darling
7
8
9 @click.group()
10 def cli():
11     # override me
12     pass
```

```
1 @click.command()
2 @click.option("-l", "--vars_to_launch", default=None)
3 @click.option("-f", "--fs", default=None)
4 def derivtoymodel(vars_to_launch, fs):
5     deriv_toy_model.main(vars_to_launch, fs)
```

```
1 cli.add_command(derivtoymodel)
2
3
4 if __name__ == '__main__':
5     cli()
```