


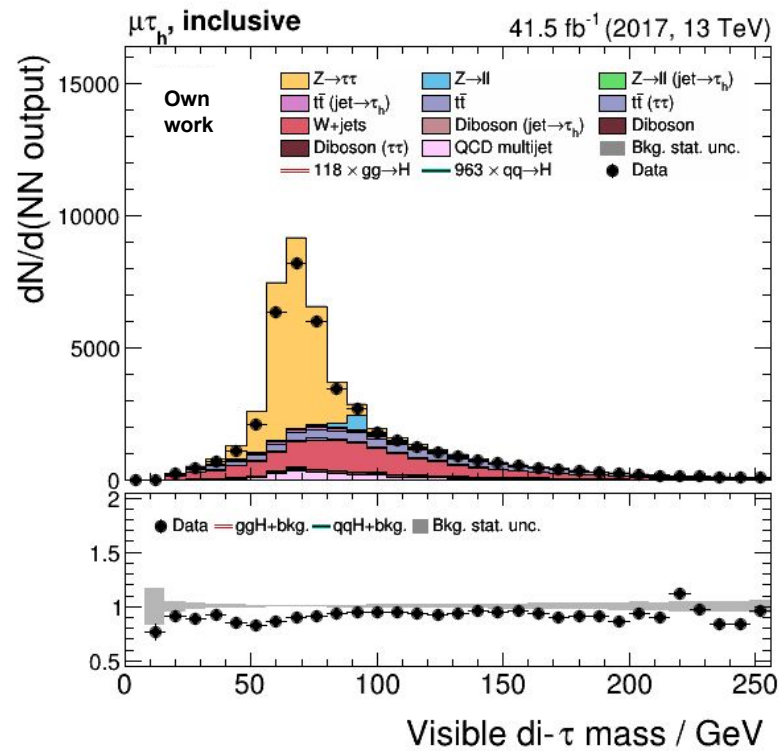
# Benchmarking an RDataFrame Complex Analysis

Massimiliano Galli

[massimiliano.galli@cern.ch](mailto:massimiliano.galli@cern.ch)




# Recap - Motivation




- CMS search for Higgs decay into tau-tau final state
- Full Run 2 analysis
- Production of many plots with data and simulations of signal and various background contributions
- Obtain the same results in a faster and more efficient way using modern ROOT facilities (RDataFrame vs TTree::Draw())

# Recap - Orders of Magnitude

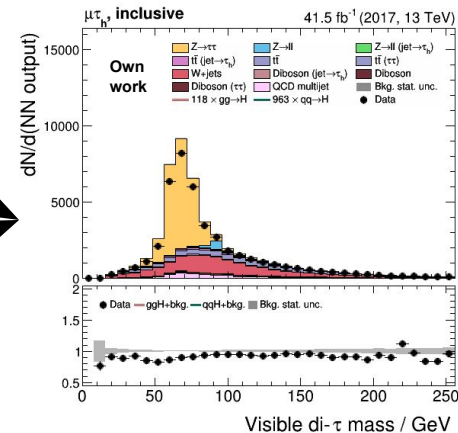
## NTuples



## Histograms



## Stack Plots




**1.1 Tb** (main  
NTuples)  
**120 Gb** (friends)

~ Mb

~ Mb


# Recap - Orders of Magnitude

## NTuples



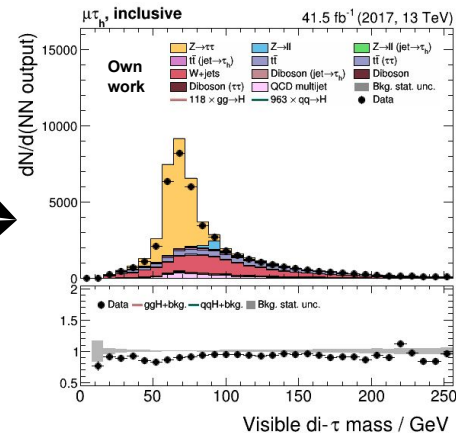
**~100** ROOT files,  
**100-150** TTrees  
each

## Histograms



**$\sim 3.6 \times 10^5$**   
histograms

## Stack Plots



**~ 100** stack plots

# Programming Model

## Book Results

For every histogram that want to produce we declare initial dataset, cuts, weights and systematic variations that we want to apply

## Optimize Computations

Datasets, selections and histogram productions are treated as nodes of a graph. The common ones are merged to perform every action only once

## Run Computations


The previous graphs are converted to the language of RDataFrame and the event loop is run

# Programming Model - 1. Book Results

## Book Results

For every histogram that want to produce we declare initial dataset, cuts, weights and systematic variations that we want to apply

```
Unit(dataset=dy, selections=[mt, ztt, vbf], histo)  
Unit(dataset=dy, selections=[mt, zl], histo)
```



# Programming Model - 1. Book Results

## Book Results

For every histogram that want to produce we declare initial dataset, cuts, weights and systematic variations that we want to apply

Types of systematic variations:

- ChangeDataset
- AddWeight
- ReplaceWeight
- SquareWeight
- RemoveWeight
- AddCut
- RemoveCut

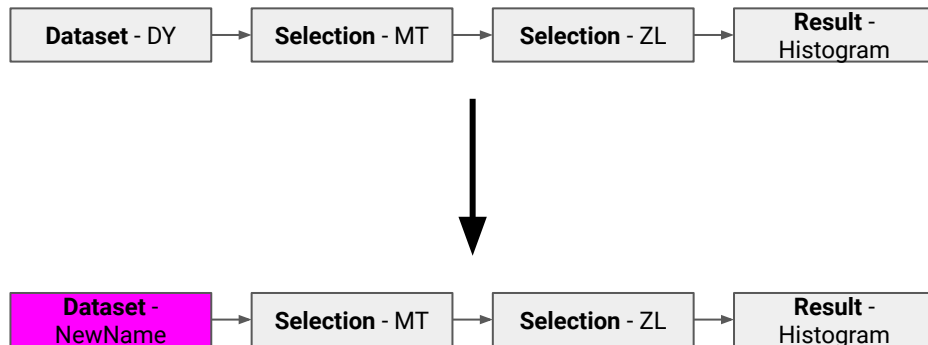
# Programming Model - 1. Book Results

## Book Results

For every histogram that want to produce we declare initial dataset, cuts, weights and systematic variations that we want to apply

Ex: ChangeDataset

```
var = ChangeDataset("NewName", "NewDirectory")  
um.book([z1_unit], [var])
```





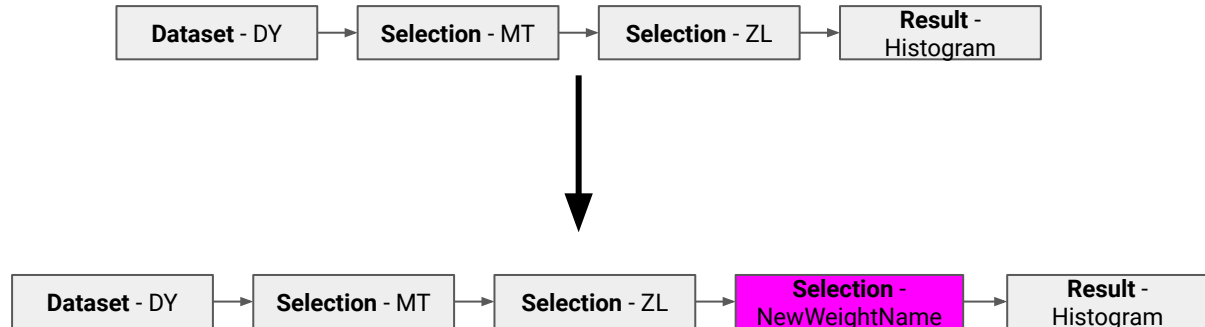
# Programming Model - 1. Book Results

## Book Results

For every histogram that want to produce we declare initial dataset, cuts, weights and systematic variations that we want to apply

Ex: AddWeight

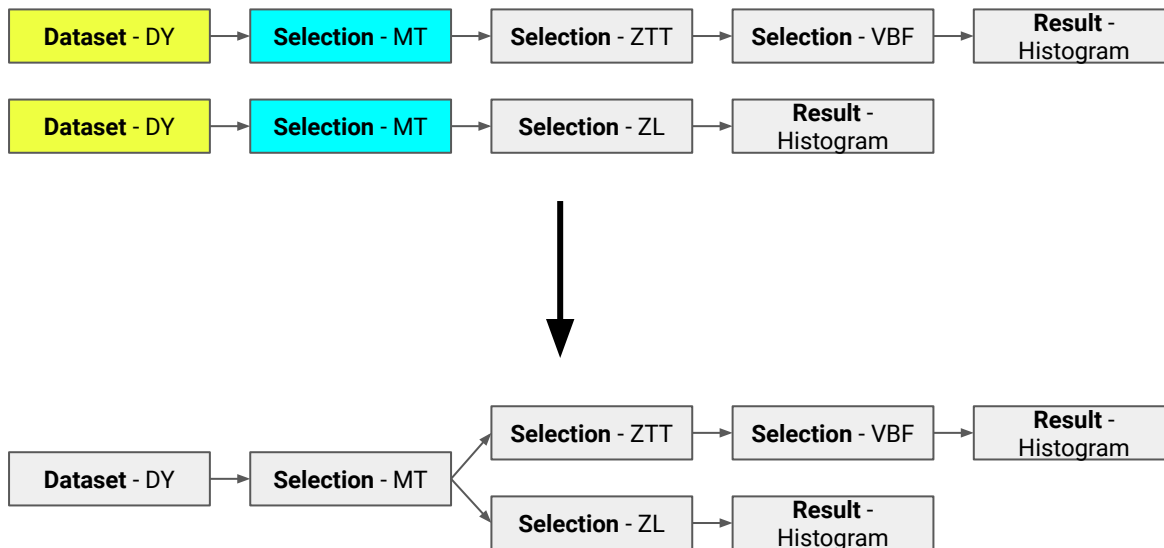
```
var = AddWeight("VariationName", Weight("NewWeightExp",  
"NewWeightName"))  
um.book([z1_unit], [var])
```



# Programming Model - 2. Optimize Computations

## Optimize Computations

Datasets, selections and histogram productions are treated as nodes of a graph. The common ones are merged to perform every action only once




# Programming Model - 3. Run Computations

## Run Computations

The previous graphs are converted to the language of RDataFrame and the event loop is run

- One RDataFrame for each node of type 'dataset'
- Support for splitting in jobs and sending them to different computing environments

# Full Systematics Analysis - Data Size



**50 ROOT Files**


**1174 of 5521 TTrees**

**~20 of ~550 branches**

**$\sim 4 \times 10^9$  events**

**18 Gb of 595 Gb**

**no friends**




**436 histograms for 1 variable**

# Benchmark Scenarios

---

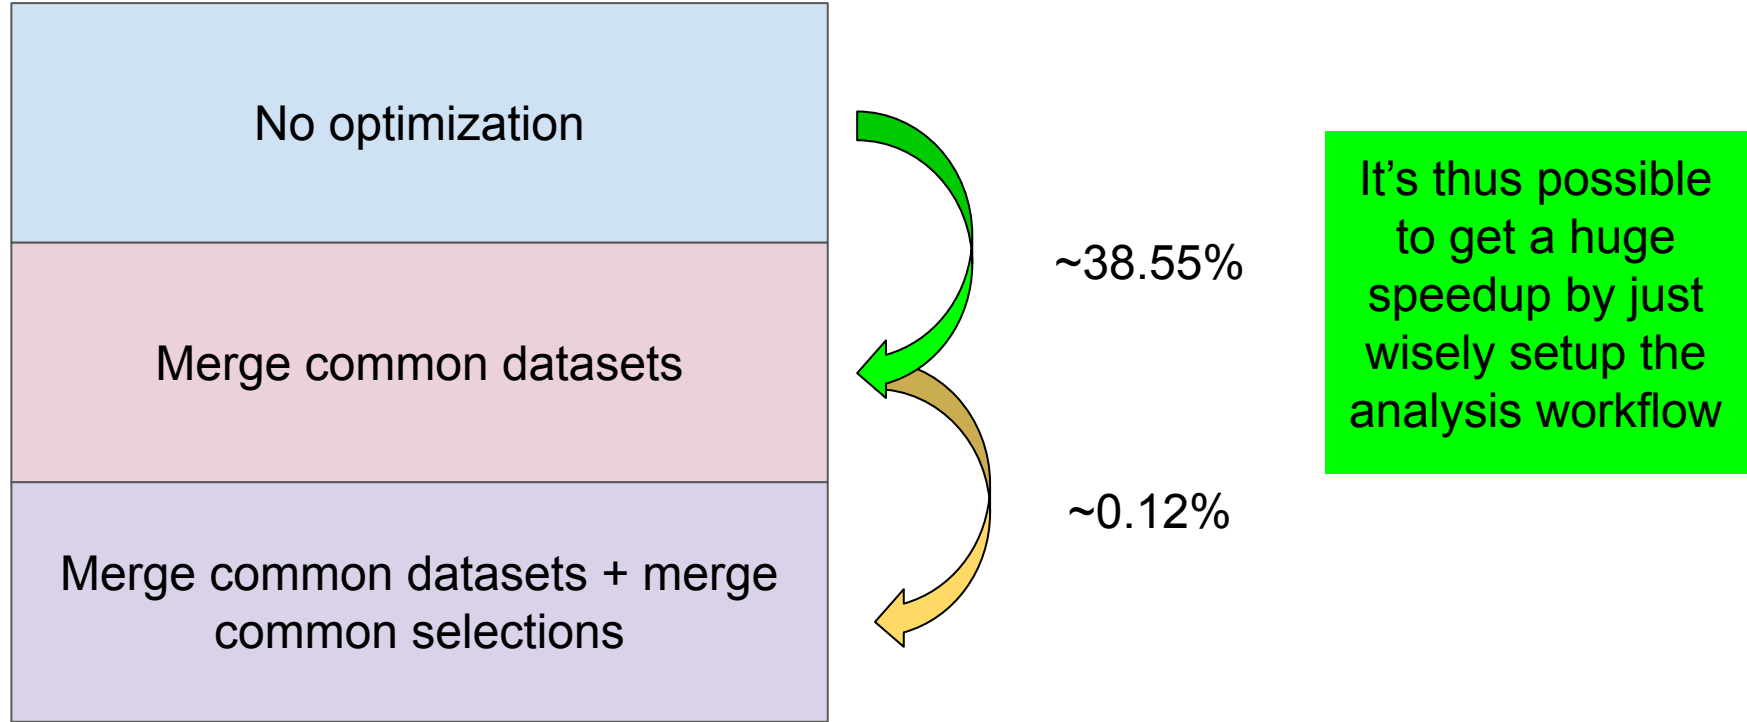
1. Merge common datasets and nodes
2. Multiprocessing
3. Multithreading
4. Many variables scaling

# 1. Merge Common Datasets and Nodes



(cfr. In the control plots analysis presented last time we merged 22 into 7)

# 1. Merge Common Datasets and Nodes




## 2. Multiprocessing

---

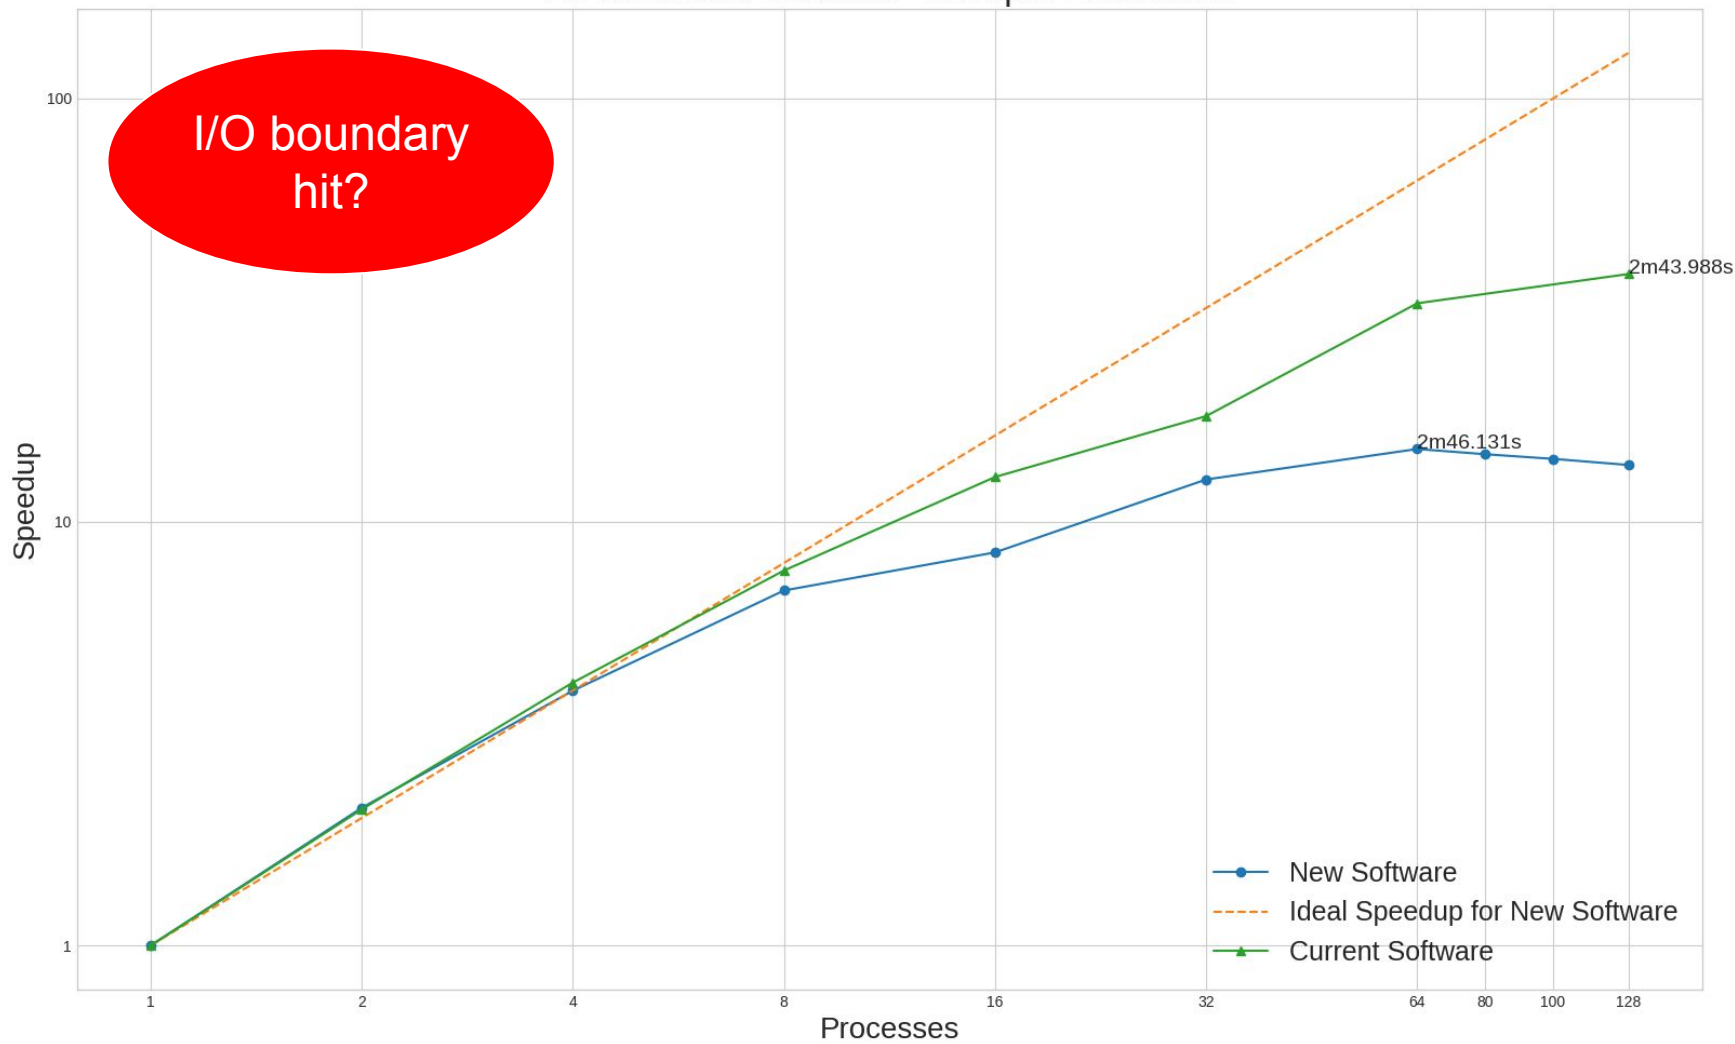
- 1 channel, 1 variable
- All systematic variations included
- Production of 436 histograms
- Scale from 1 to 128 processes
- Test on machine with 128 (64) logical (physical) cores
- Comparison with the current software (`TTree::Draw()`)



All variations included - Multiple Processes



# All variations included - Multiple Processes

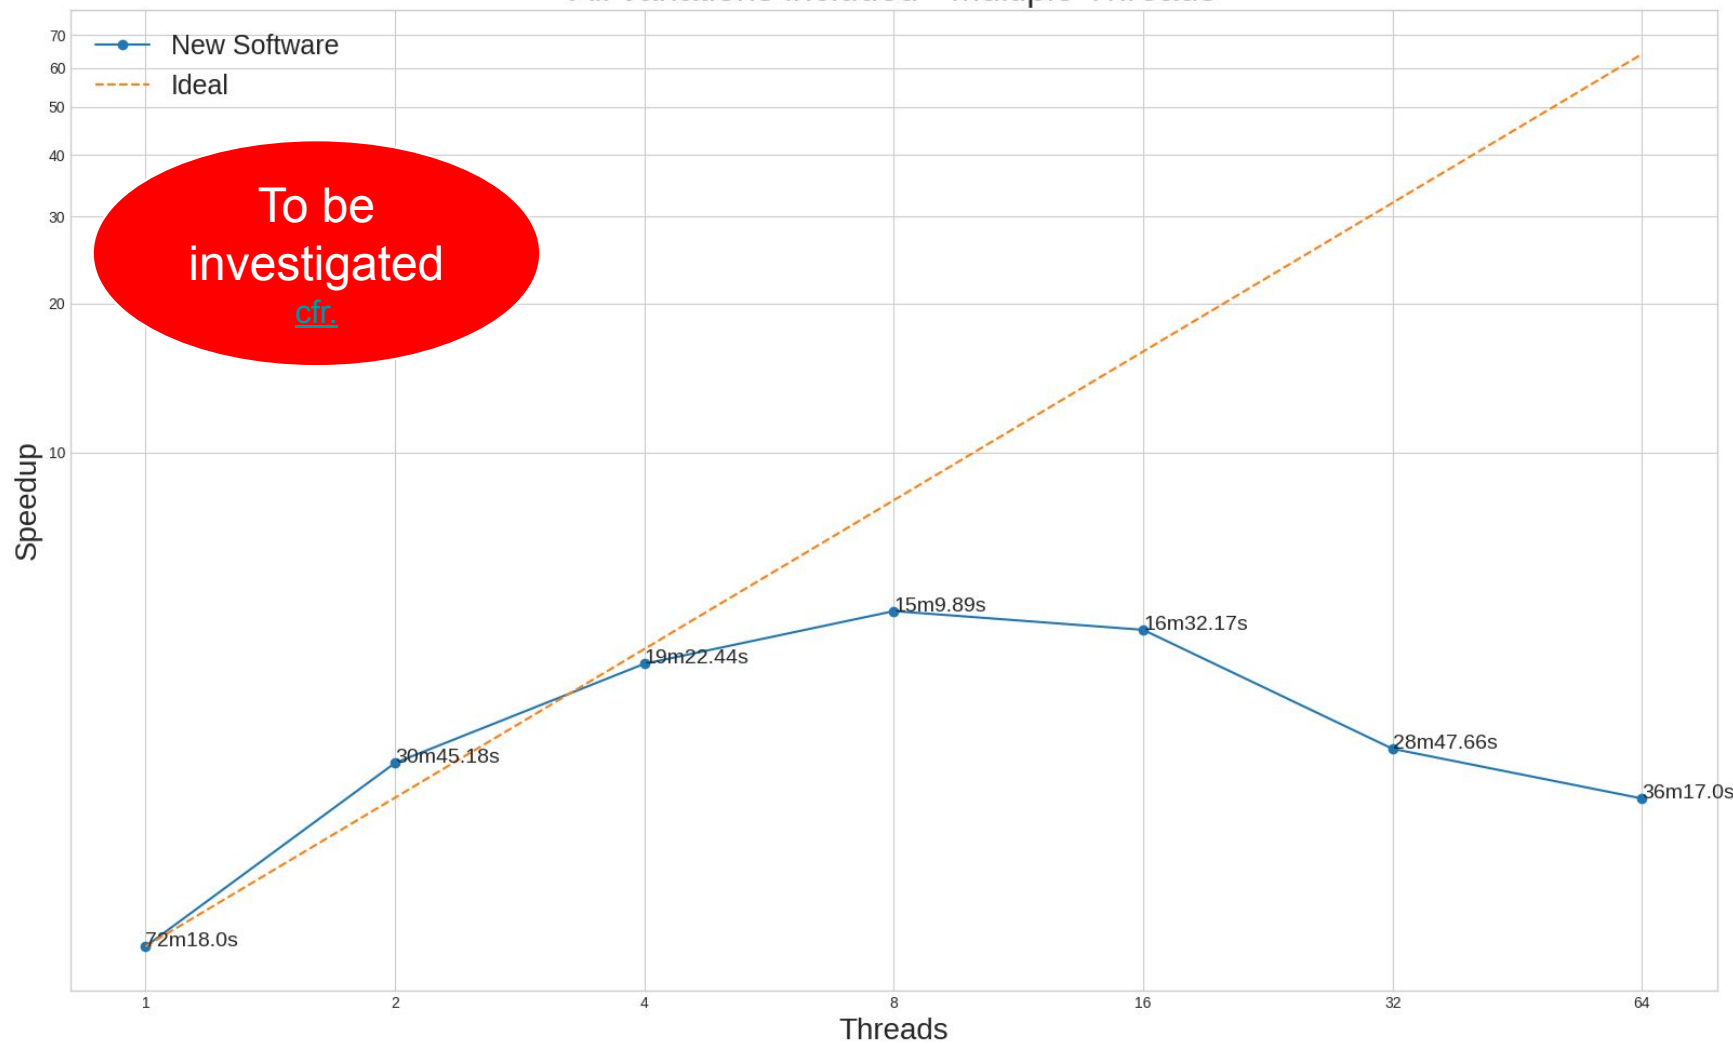


### 3. Multithreading

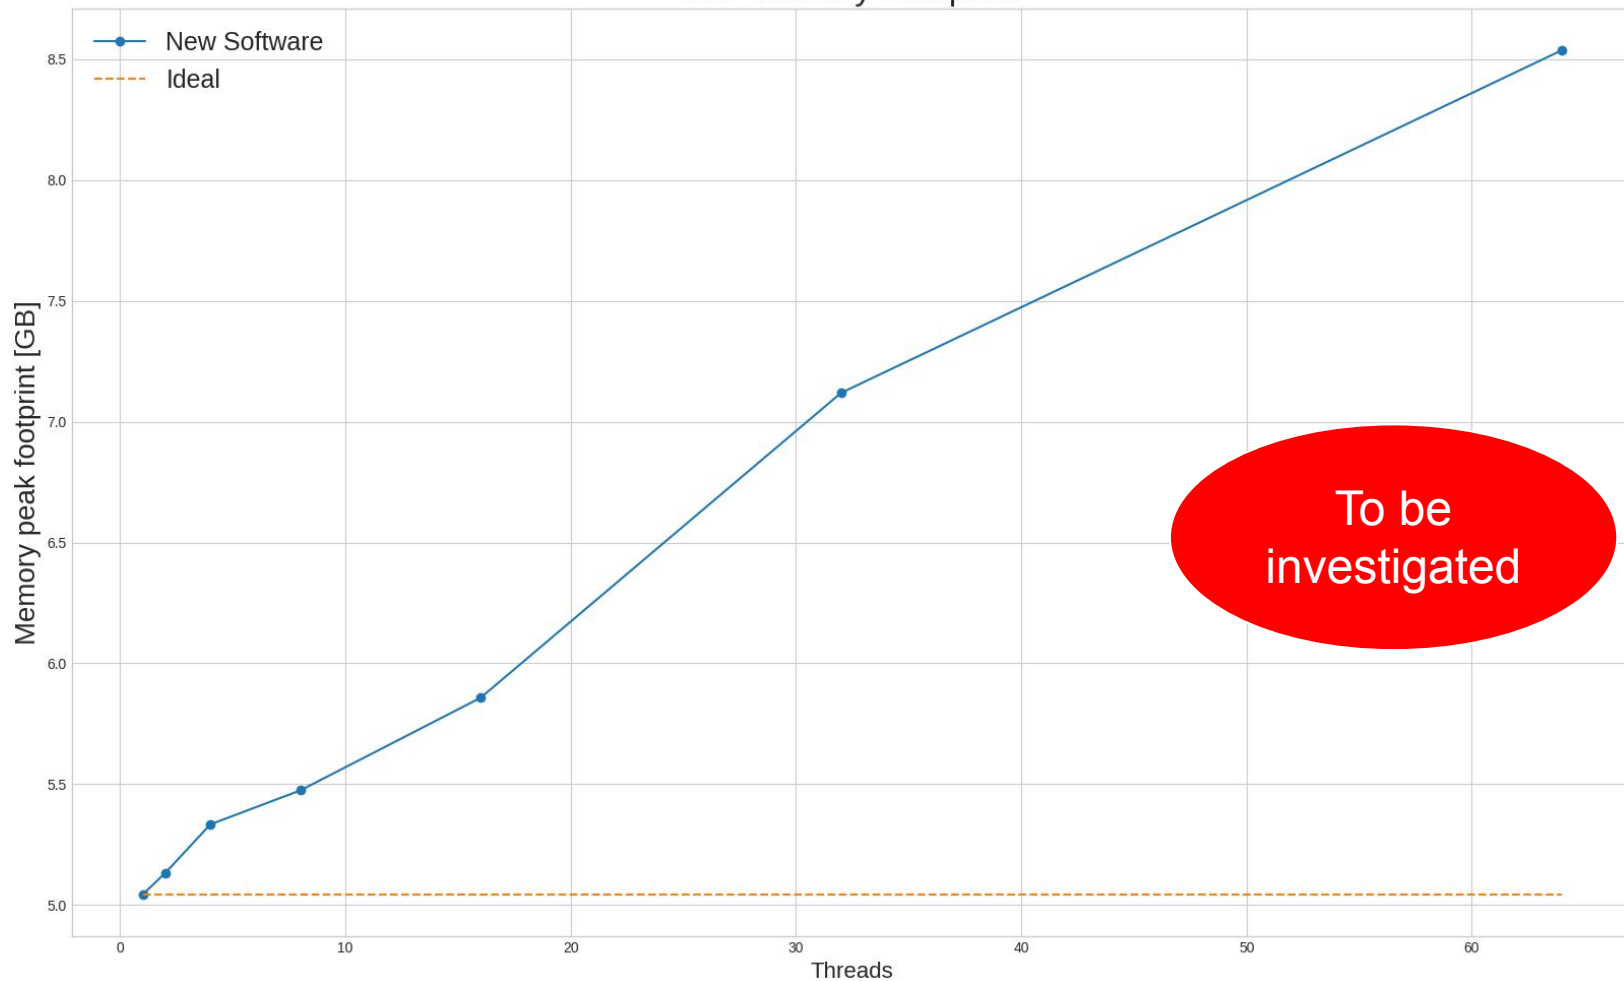
---

- Enabled with `RDataFrame.EnableImplicitMT(N)`
- 1 channel, 1 variable, single process
- Production of 436 histograms
- Scale from 1 to 64 threads
- Benchmark speedup
- Benchmark memory footprint and compare to multiprocessing


# All variations included - Multiple Threads



## MT-Memory Footprint



MP/MT - Memory Footprint

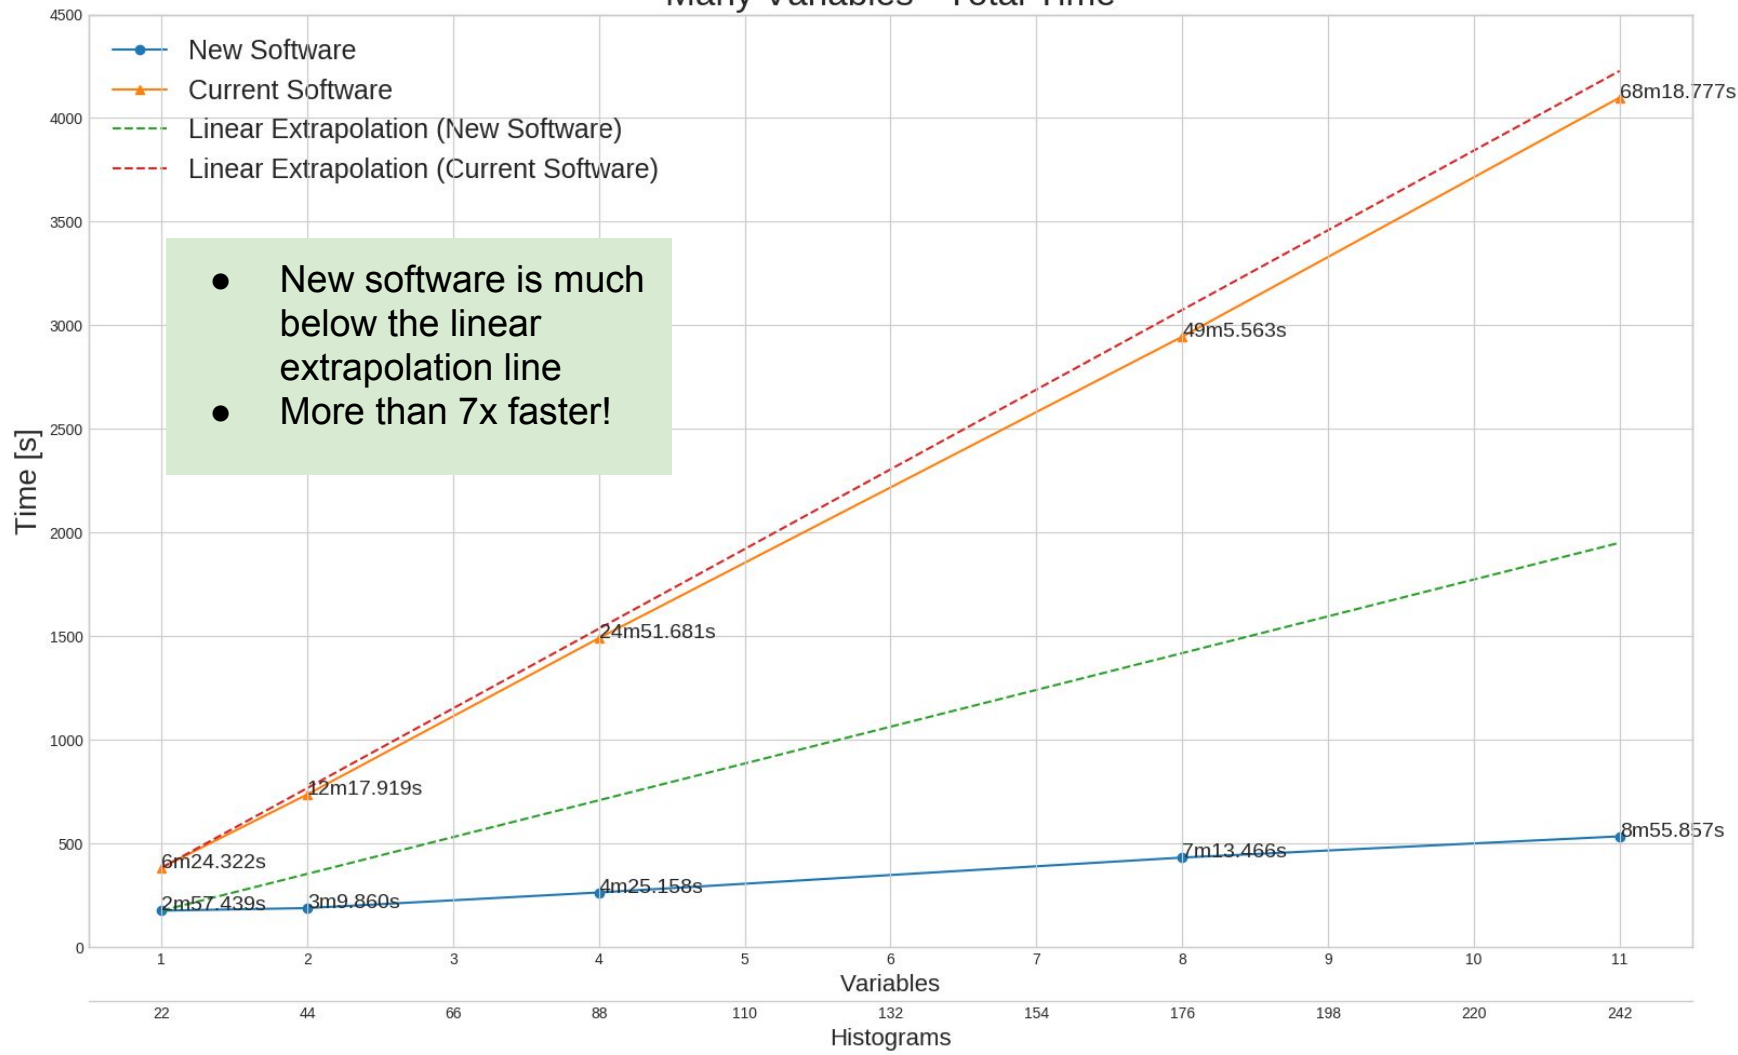


## 4. Many Variables Scaling

---


- Single process, single thread
- Inclusive analysis (22 histograms produced for each variable)
- Different number of variables each time
- Benchmarked event loop time scaling, time per single histogram, memory footprint
- Comparison with the current software

## Many Variables - Total Time



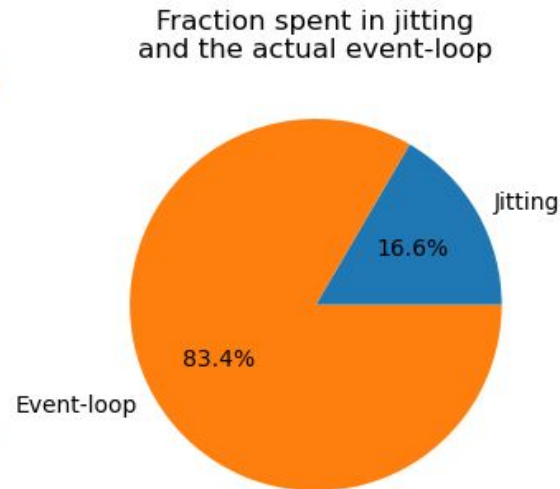
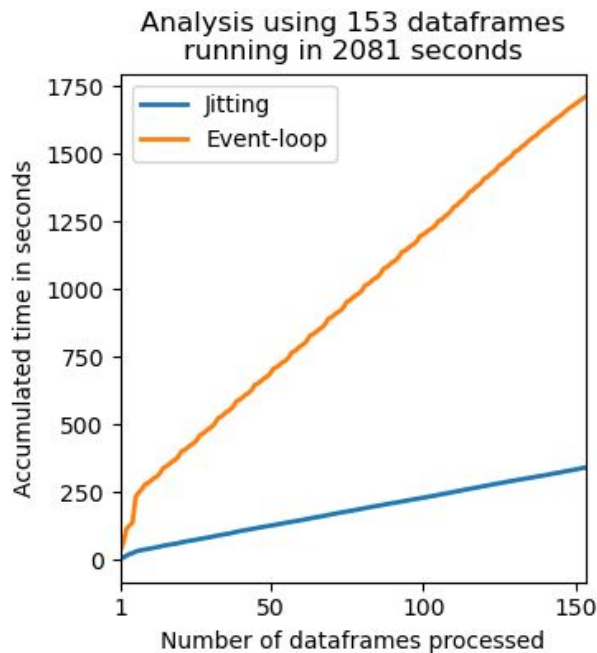


## Many Variables - Memory Footprint



# Time spent in RDF event-loop

- Follow up to last weeks PPP
- Analysis with all systematics uses 153 dataframes
- Measuring time spent in RDF event-loop
  - Jitting
  - Actual event-loop



# Discussion

---

- Drop in scaling with MP: I/O boundary hit?
- Very bad scaling with MT: why?
- Further benchmarking: suggestions?
  - Is it worth combining ROOT MT with Python MP?




# Backup

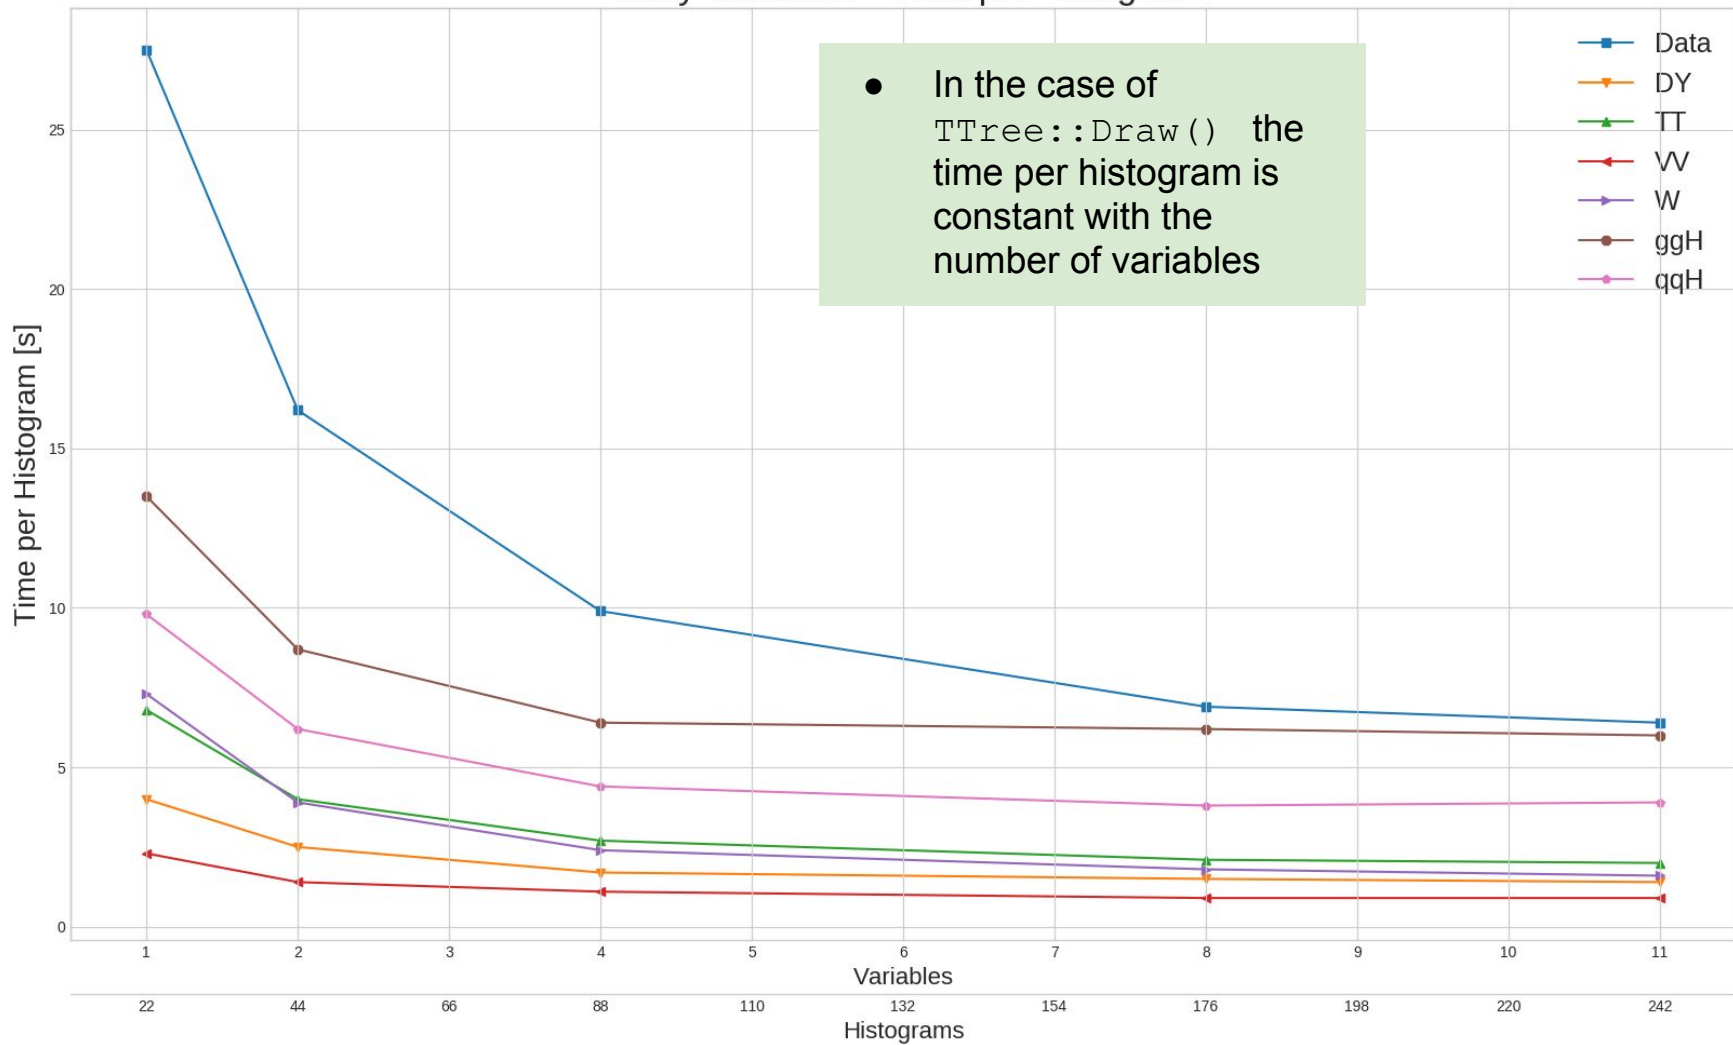
```
def get_nominal_datasets(channel):  
    datasets = dict()  
    for key, names in nominal_files.items():  
        datasets[key] = dataset_from_nameset(  
            key, names, channel + '_nominal', base_file, base_friends)  
    return datasets
```

```
def get_nominal_units(channel, datasets):  
    return {  
        (...)  
        'ztt' : Unit(datasets['DY'], [channel_selection(channel), DY_process_selection(channel),  
                                     ZTT_process_selection(channel)][histos]  
        ),  
        (...)  
    }
```

# All variations included - Multiple Threads



# Many Variables - Time per histogram



# Book Results - Concise and Structured

```
channels = ['mt', ...]

# Book nominal Units

nominals = {}
nominals['2017'] = {}
nominals['2017']['datasets'] = {}

# E.g. DY dataset

dy_dataset = dataset_from_nameset('DY', nominal_files['DY'], 'mt_nominal',
                                   base_file, base_friends)

# nominal files is placed inside ntuple_config
nominal_files = {
    (...)
    'DY': [
        'DY2JetsToLLM50_RunIIFall17MiniAODv2_PU2017_13TeV_MINIAOD_madgraph-pythia8_ext1-v2',
        'DY2JetsToLLM50_RunIIFall17MiniAODv2_PU2017_13TeV_MINIAOD_madgraph-pythia8_v1',
        'DY3JetsToLLM50_RunIIFall17MiniAODv2_PU2017_13TeV_MINIAOD_madgraph-pythia8_ext1-v1',
        (...)
    ],
    (...)
}
```