Exploring the Trade-offs between Energy and Performance of Federated Learning Algorithms

— A measurement —

Presenter

Mai Huong Do

Advisors

Georges Da Costa
Millian Poquet

IRIT - Université Toulouse III

February 21, 2025

Table of Contents

What, Why, How?

What is it?

Trade-offs between **Energy** and **Performance** of **Federated Learning**?

- **Energy**: Energy consumption when running a process (FL process).
- **Performance**: Performance of model, usually are evaluated by accuracy, loss, time processing.
- **Federated Learning (FL)**: One of the most growing research in ML, also requires huge of resources.

Federated Learning framework

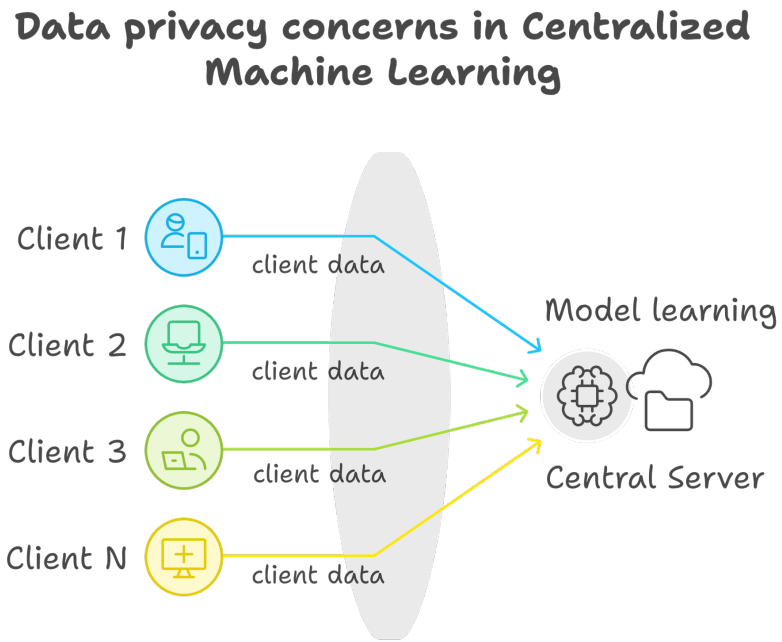

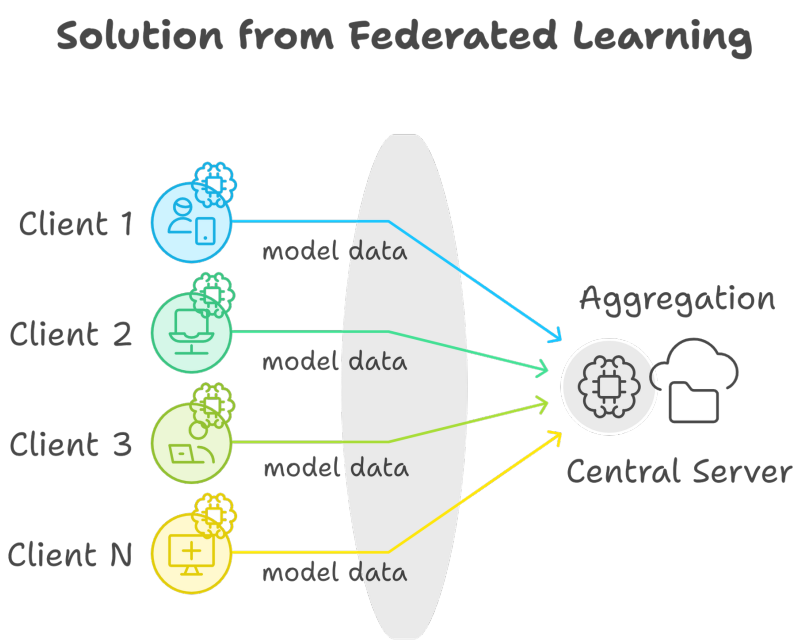

Federated Learning framework

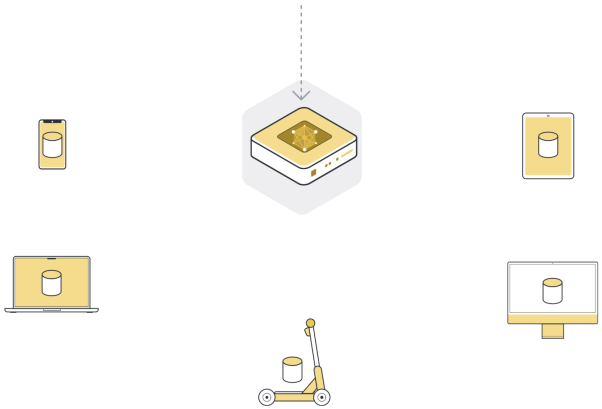

Step 1: Global model init.

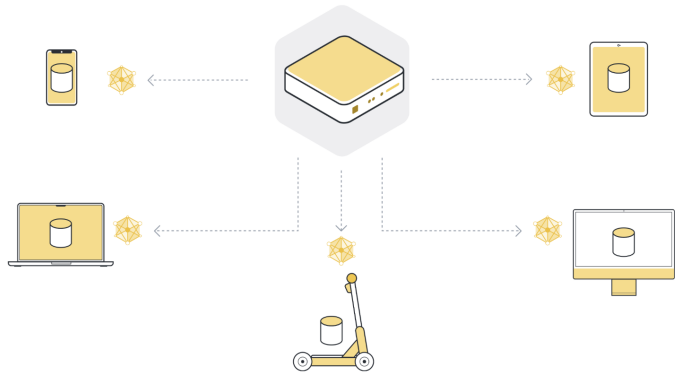

Step 2: Send model to clients

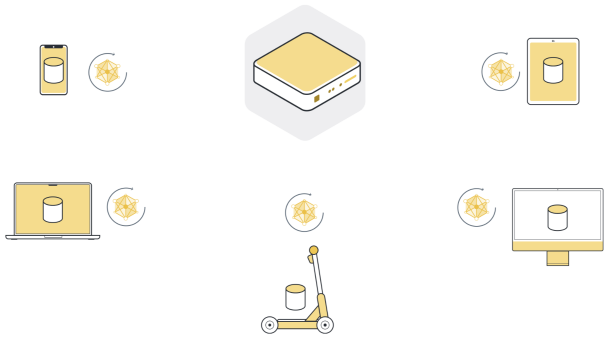

Step 3: Local training

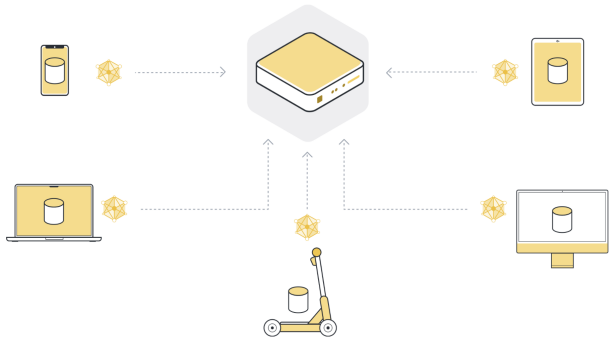

Step 4: Return and aggregate in global

Why? - High energy consumption for AI

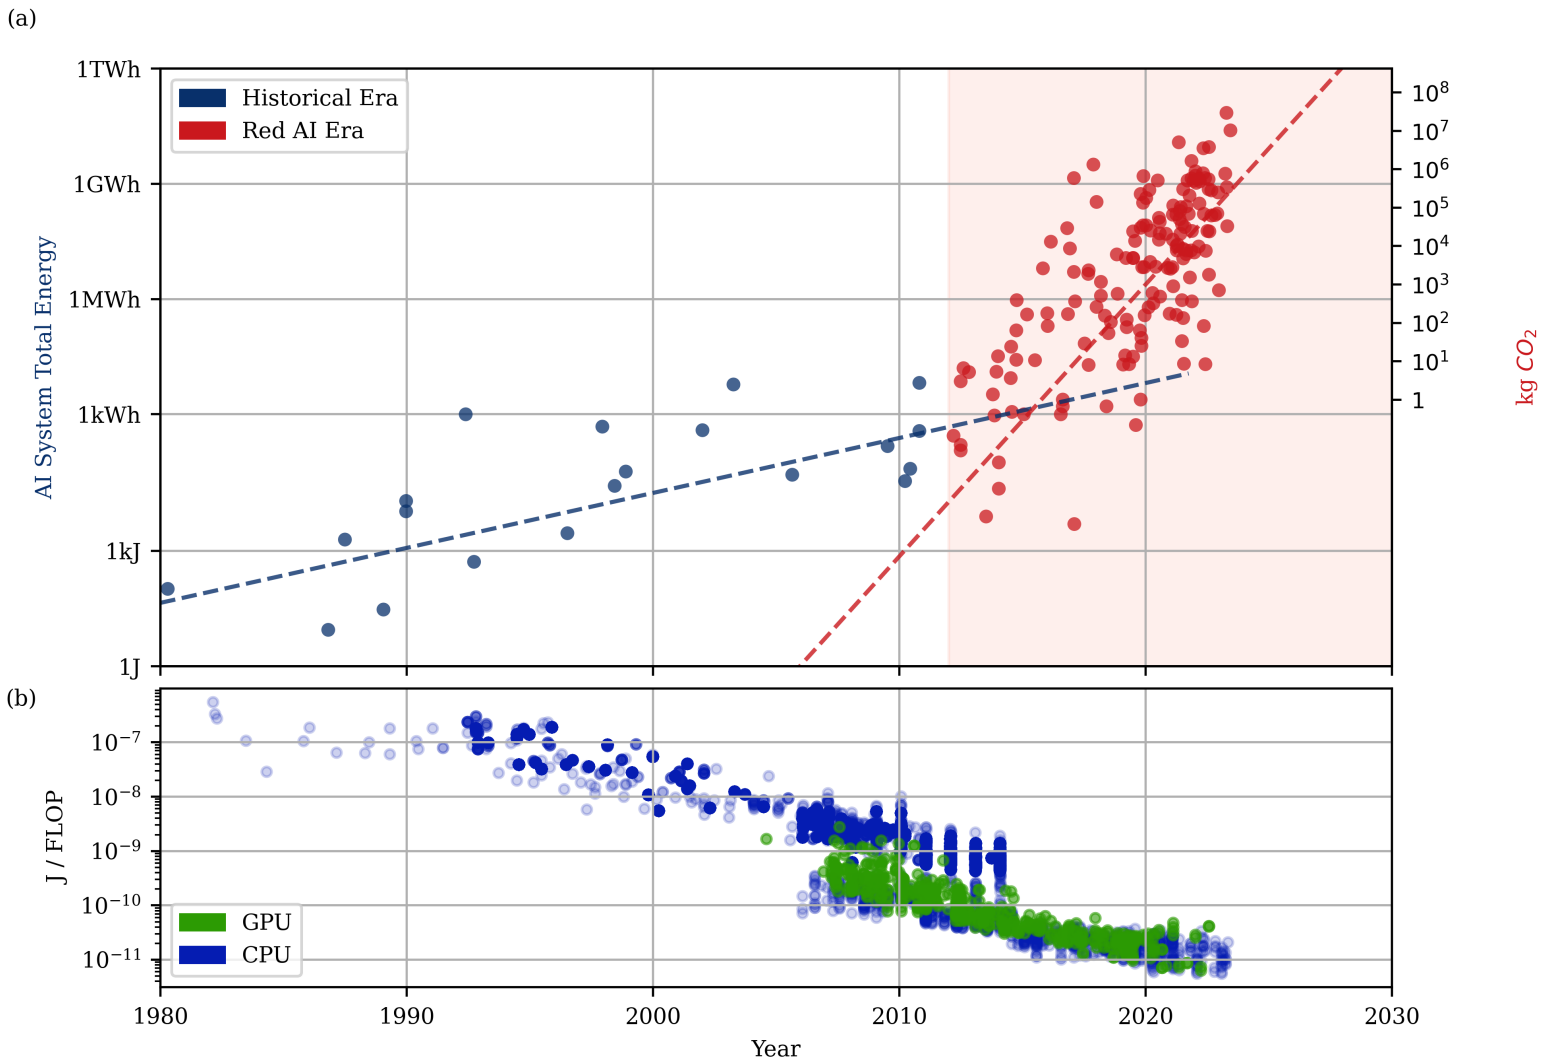

[1]

[1] Measuring the Energy Consumption and Efficiency of Deep Neural Networks: An Empirical Analysis and Design Recommendations (2024), analysis **BUTTER-E** Dataset

From influence to warning

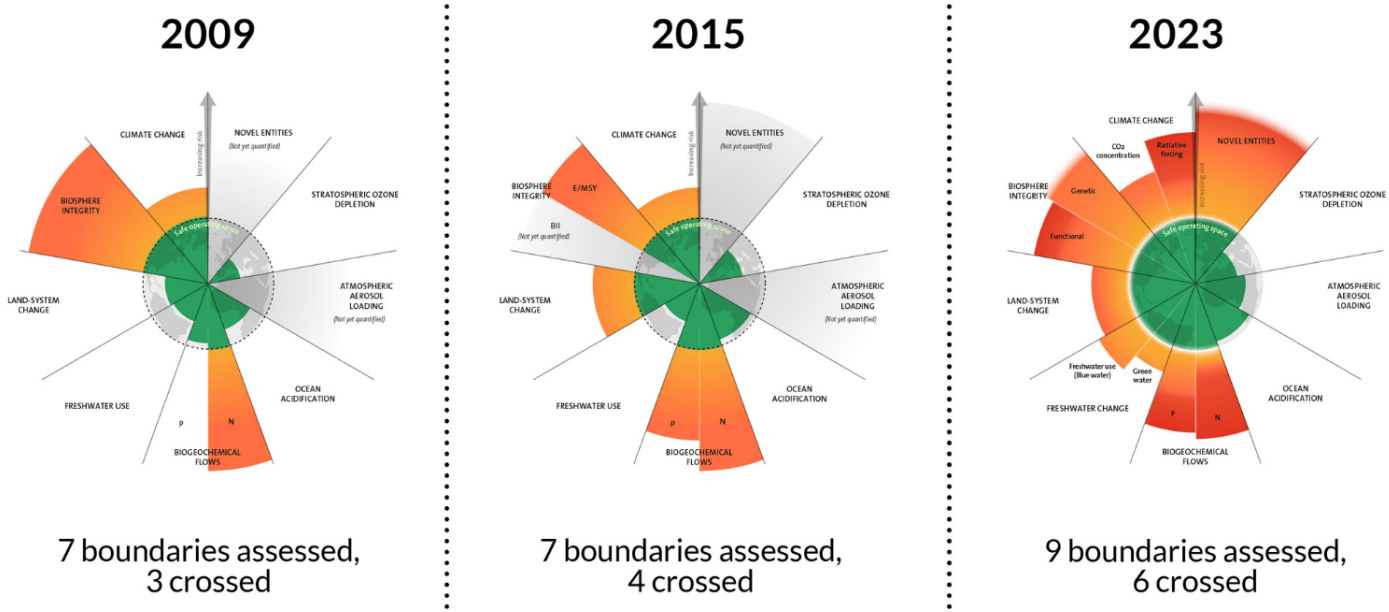

[2] "**Red AI** *refers to AI research that seeks to improve accuracy (or related measures) through the use of* **massive computational power** *while disregarding the cost — essentially* **"buying"** *stronger results."*
"**Green AI** *refers to AI research that yields novel results while taking into account the* **computational cost**, *encouraging a* **reduction in resources spent**."

[2] Schwartz, Roy, et al. "Green ai." Communications of the ACM 63.12 (2020): 54-63.

Why? - We have money - Project Funding

- ANR DELIGHT (a**D**vancing f**E**derated **L**earn**I**ng while reducin**G** t**H**e carbon foo**T**print) project.
- Target: incorporate energy efficiency as one of the metrics of FL to push FL towards sustainability.
- 3 main parts:
  - Reproducible framework of energy - performance tradeoff (my thesis).
  - Improve, develop FL model.

Objective and Planning

Objective: develop method for estimating energy, build a automatic framework to explore the trade-offs between Energy and FL performance.
Phases of the thesis:

- Set up an experimental environment on Grid'5000 (g5k) to gather performance and energy metrics.
  - Create a use-case for the Flower framework.
  - Build a reproducible and automated framework for obtaining metrics for this use case
- Propose, formulate energy model, and implement the different leverages.
- Explore the impact of the leverages on both energy and performance.

General framework

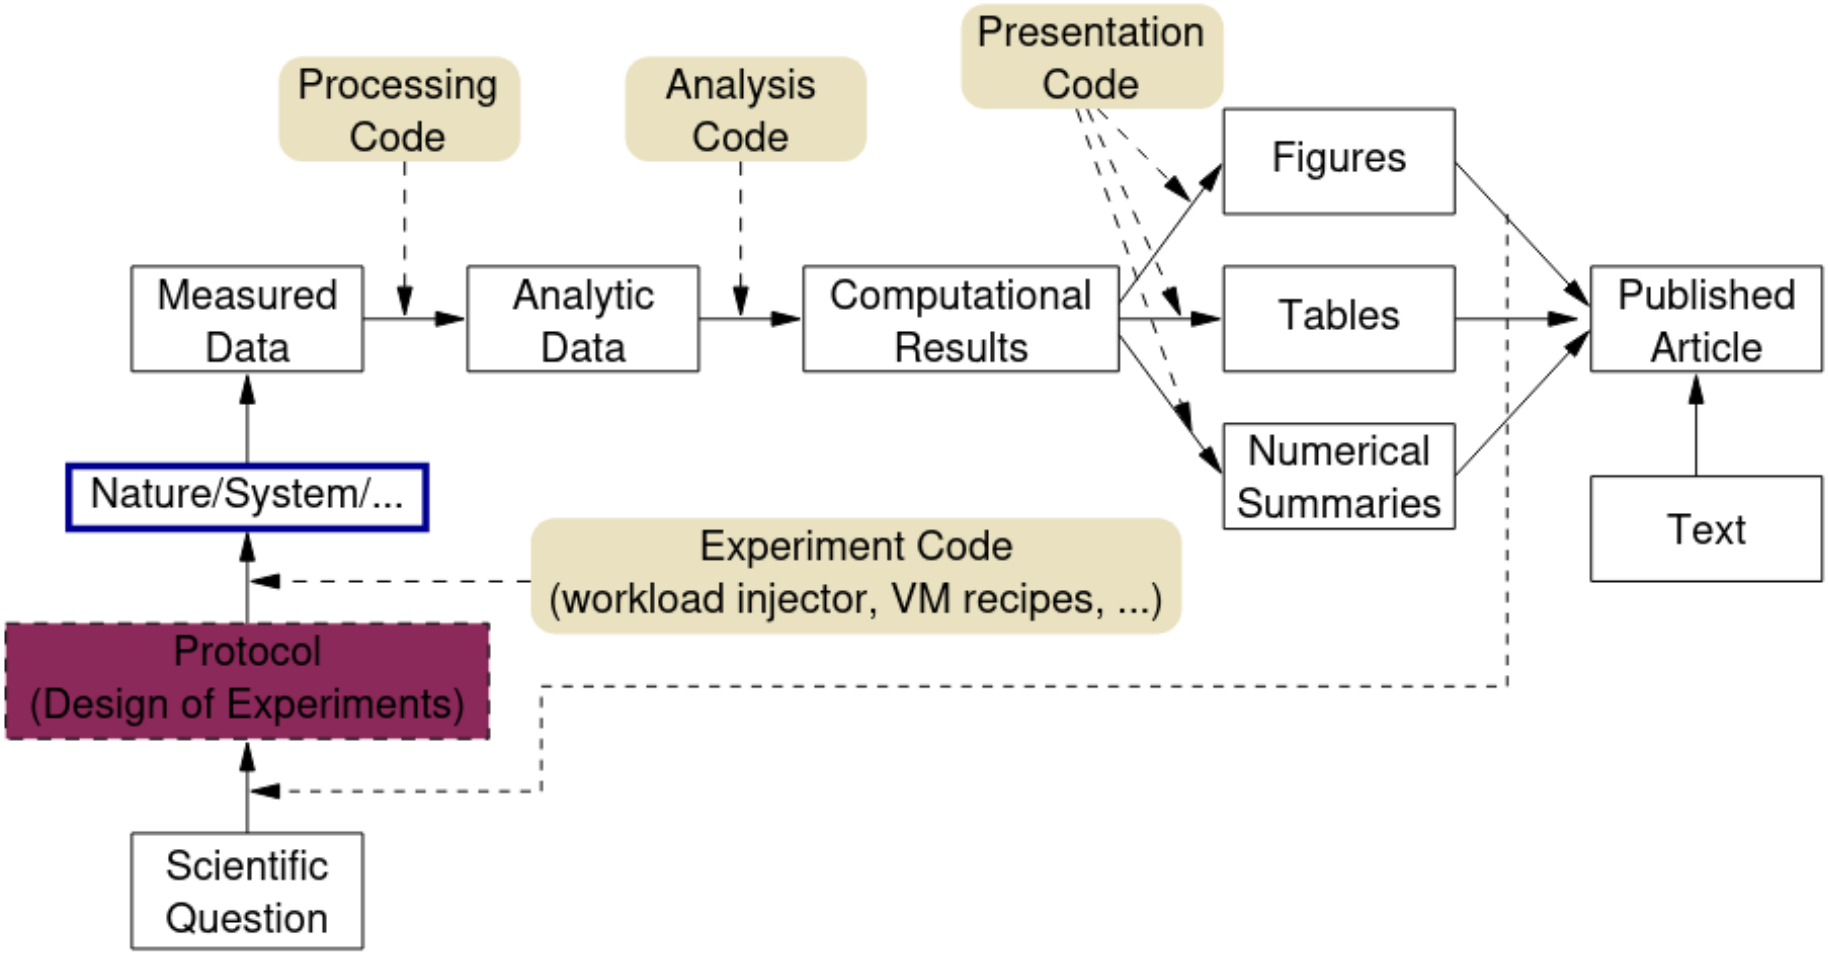

(from Millian slide - or Arnaud Legrand?)

General framework

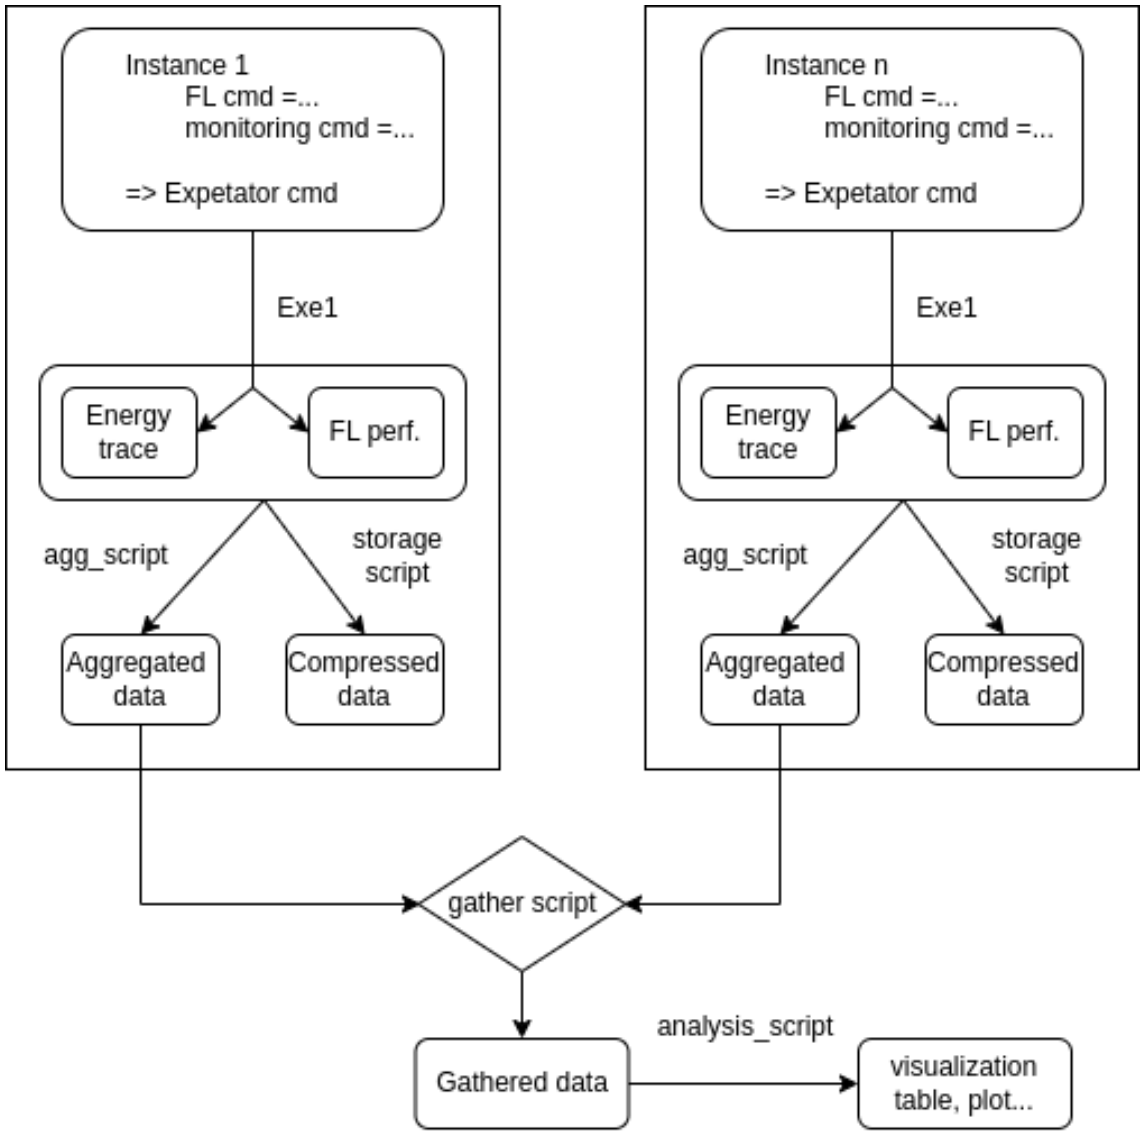

Experiment

Process

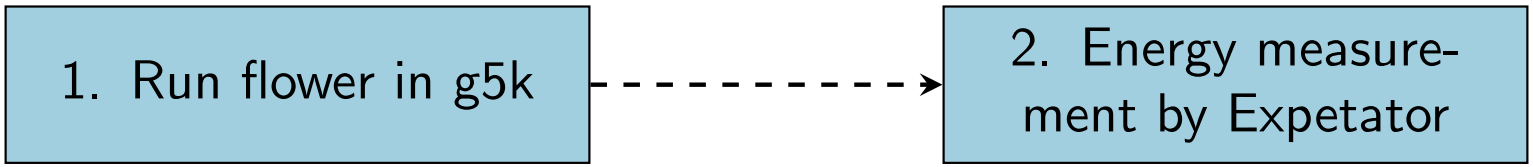

- Flower implement - see more
- Set of sensors - see more

Process

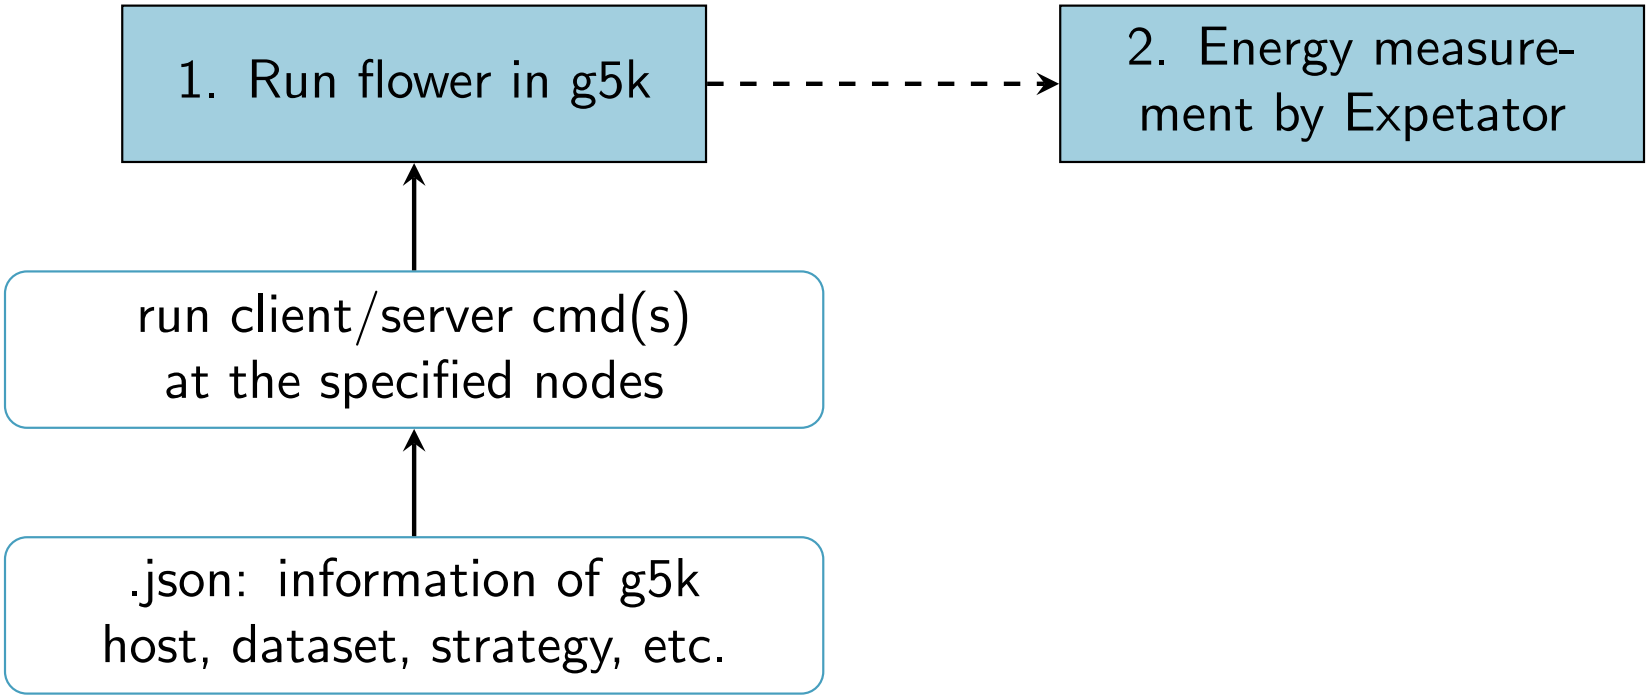

▶ Flower implement - see more ▶ Set of sensors - see more

Process

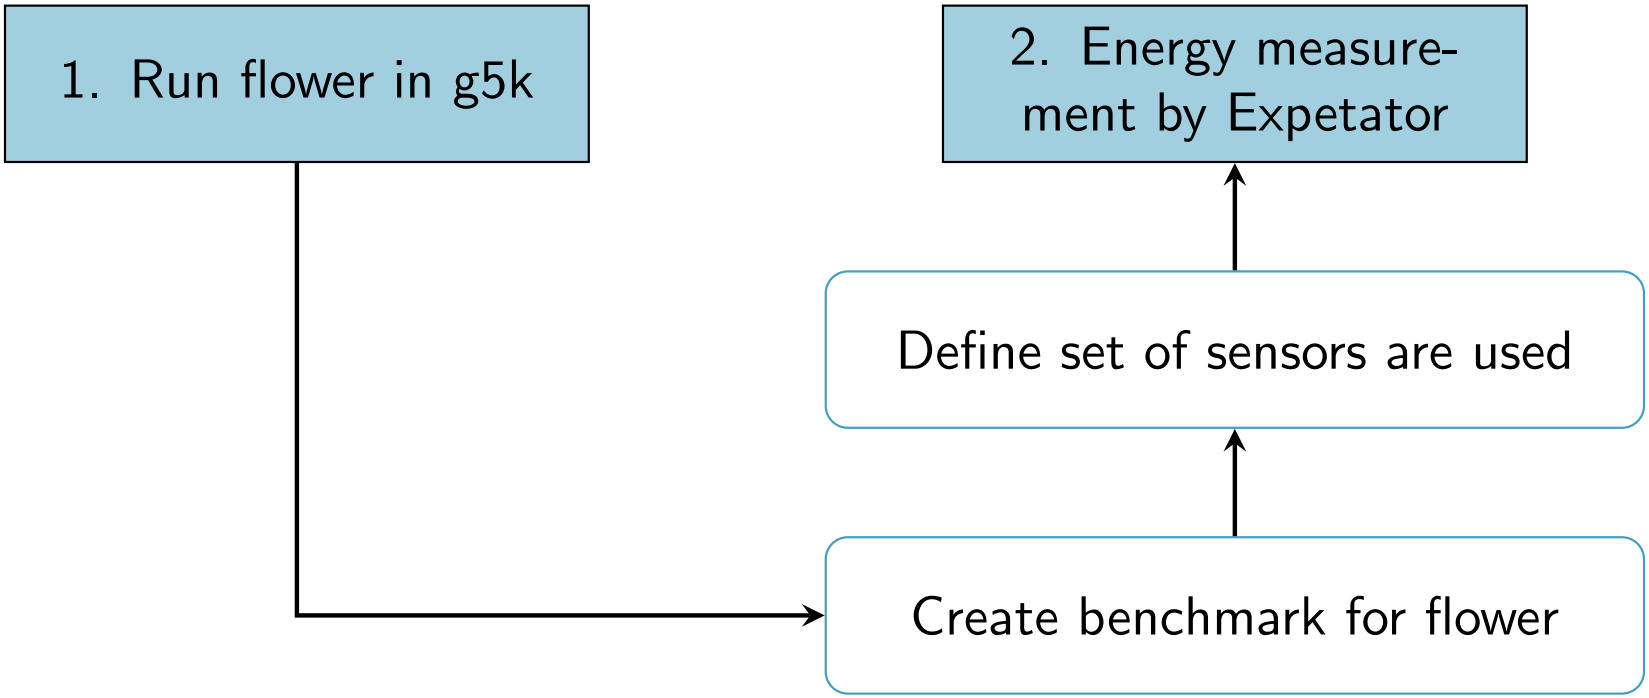

▸ Flower implement - see more ▸ Set of sensors - see more

G5k platform

Choose the nodes:

- Site: Nancy
- Cluster: Gros
- Cpu: Intel Xeon Gold 5220 - 18 Cores - x86_64
- Mem: 96GB
- Storage: 480GB SSD + 960GB SSD*
- Net: 2x25 Gbps (SR-IOV)
- Include Kwollect Powermetter

2 monitors

Keyword	**Kwollect**	**RAPL**
Purpose	SW/HW-based monitors usage monitoring	Energy measurement and thermal management monitoring CPU power usage and system performance
Measurement	Power consumption for various components (CPU, memory, entire server)	focus on CPU cores, within Intel CPUs, package (entire processor), and DRAM (memory)...

Code structure

Repo[3]
Readme[4]

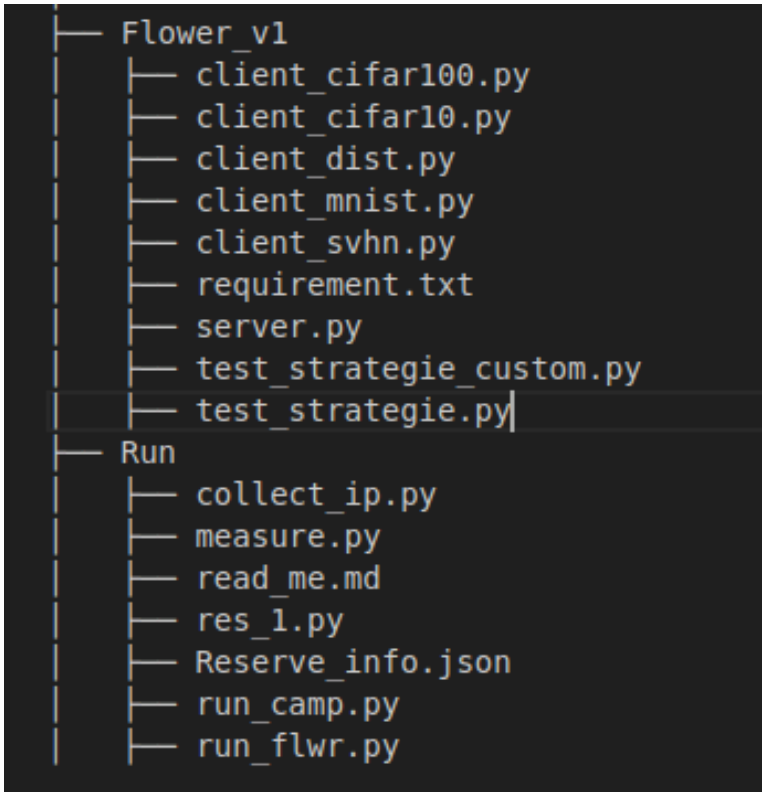

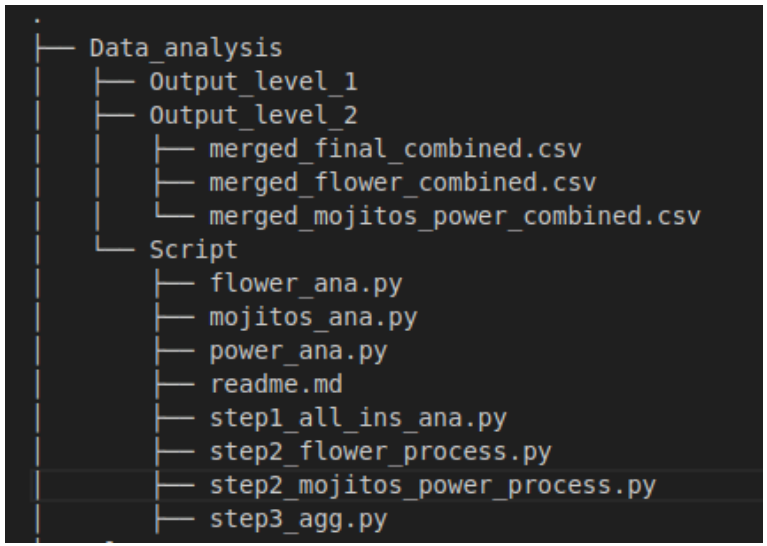

[3] https://gitlab.irit.fr/huongdm/huong_delight/Huong_Journey

[4] https://gitlab.irit.fr/huongdm/huong_delight/Huong_Journey/-/blob/main/Run/read_me.md?ref_type=heads

Data storing /Log/

```
    ├── Flower_test_5
    ├── Flower_cifar10_ori_1
    │   ├── Flwr_20241020_135940
    │   │   ├── Client_172.16.48.10
    │   │   ├── Client_172.16.48.13
    │   │   ├── Client_172.16.50.3
    │   │   └── Server_172.16.51.1
    │   ├── Flwr_20241020_140030
    │   │   ├── Client_172.16.48.10
    │   │   ├── Client_172.16.48.13
    │   │   ├── Client_172.16.50.3
    │   │   └── Server_172.16.51.1
    │   ├── _hercule-1.lyon.grid5000.fr_1729425564
    │   ├── _hercule-1.lyon.grid5000.fr_1729425564_mojitos
    │   │   ├── hercule-1.lyon.grid5000.fr_flower_1729425580
    │   │   ├── hercule-1.lyon.grid5000.fr_flower_1729425630
    │   │   ├── orion-3.lyon.grid5000.fr_flower_1729425580
    │   │   ├── orion-3.lyon.grid5000.fr_flower_1729425630
    │   │   ├── taurus-10.lyon.grid5000.fr_flower_1729425580
    │   │   ├── taurus-10.lyon.grid5000.fr_flower_1729425630
    │   │   ├── taurus-13.lyon.grid5000.fr_flower_1729425580
    │   │   └── taurus-13.lyon.grid5000.fr_flower_1729425630
    │   ├── _hercule-1.lyon.grid5000.fr_1729425564_power
    │   │   ├── hercule-1.lyon.grid5000.fr_flower_1729425580
    │   │   └── hercule-1.lyon.grid5000.fr_flower_1729425630
    │   └── Metadata.json
    └── Flower_cifar10_ori_2
        ├── Flwr_20241020_140137
        │   ├── Client_172.16.48.10
```

Data Processing /Data_analysis/

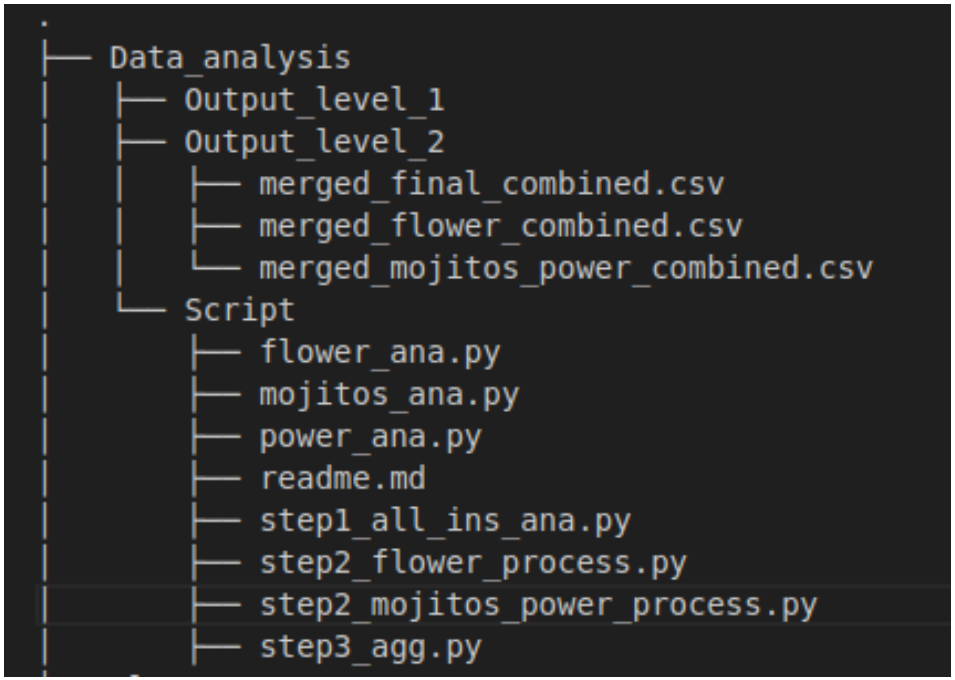

- mojitos_ana, power_ana, flower_ana: pre-process mojitos, kwollect, FL performance data of each instance
- step1_all_ins_ana: run 3 pre-processes above for all files in 1 campaign log
- step2_flower_process: extract needed perf information from flwr
- step2_mojitos_power_process: extract needed energy info from mojitos and kwollect.
- step3_agg: combine, agg, compute to export 1 final csv.

Results

Data_analysis > Output_level_2 > merged_final_combined.csv

```
dataset,strategy,num_round,log_timestamp,time (s),final_loss,file_number,mean_E_RAPL,max_time_s_RAPL,num_nodes,mean_P_Kwollect,max_time_s_Kwollect
cifar10,fedAvg,1,2024-11-05 19:43:47,8.39,1.550370693206787,1730832227,1573996.9177708772,48.70115637000004,4,85.554375,49.03917193412781
cifar10,fedAvg,5,2024-11-05 19:45:31,34.78,0.8216357231140137,1730832331,2951092.421780954,42.66357536700002,4,114.81083380601248,42.04303693771362
cifar10,fedMedian,1,2024-11-05 19:53:27,7.87,1.3695908784866333,1730832807,2745925.980583561,13.138403516000151,4,106.25,13.011778831481934
cifar10,fedMedian,5,2024-11-05 19:54:29,35.88,0.8208371996879578,1730832869,2776180.716609589,43.64776340900016,4,111.39659468438538,43.043659925460815
cifar10,new,1,2024-11-05 19:47:27,8.12,1.3844554424285889,1730832447,2746707.306390977,13.197975008999949,4,108.67994505494504,13.01743197441101
cifar10,new,5,2024-11-05 19:48:26,36.22,0.8037880659103394,1730832506,2740971.961907449,44.18649467599994,4,111.3517316017316,43.03302884101868
cifar10,ori,1,2024-11-05 19:50:25,8.15,1.1542521715164185,1730832625,1924451.6610267283,13.524328816999969,4,91.01785714285715,13.01209807395935
cifar10,ori,5,2024-11-05 19:51:28,36.67,0.9512131810188292,1730832688,1789009.1145875151,42.219894161999946,4,89.97660575858251,42.03613519668579
mnist,fedAvg,1,2024-11-05 19:56:30,8.76,0.0517027899622917,1730832990,2356143.538596492,16.947971714999994,4,106.52941176470588,16.011781930923462
mnist,fedAvg,5,2024-11-05 19:57:36,40.38,0.0221263654530048,1730833056,2643139.7951609925,48.10968462500023,4,107.89583333333331,48.0528609752655
mnist,fedMedian,1,2024-11-05 20:06:10,8.71,0.0423926003277301,1730833570,2676900.4148509763,13.814431886999955,4,104.93333333333332,14.020021915435793
mnist,fedMedian,5,2024-11-05 20:07:12,40.52,0.0238807089626789,1730833632,2632104.9772498296,48.31467085899976,4,107.35323491098568,48.04829788208008
mnist,new,1,2024-11-05 19:59:44,9.09,0.0575130358338356,1730833184,1958596.9998581011,16.90818759800004,4,93.18055555555556,17.059574842453003
mnist,new,5,2024-11-05 20:00:51,39.58,0.0205710977315902,1730833251,2808162.922707889,46.79084990499996,4,111.88297872340426,46.057049036026
mnist,ori,1,2024-11-05 20:03:00,8.93,0.0473004952073097,1730833380,1776703.582672609,14.117340154999964,4,90.73076923076924,14.014400005340576
mnist,ori,5,2024-11-05 20:04:02,40.94,0.0371751897037029,1730833441,1651069.0905411898,45.79304104100038,4,88.06241134751774,46.0476610660553
```

Results

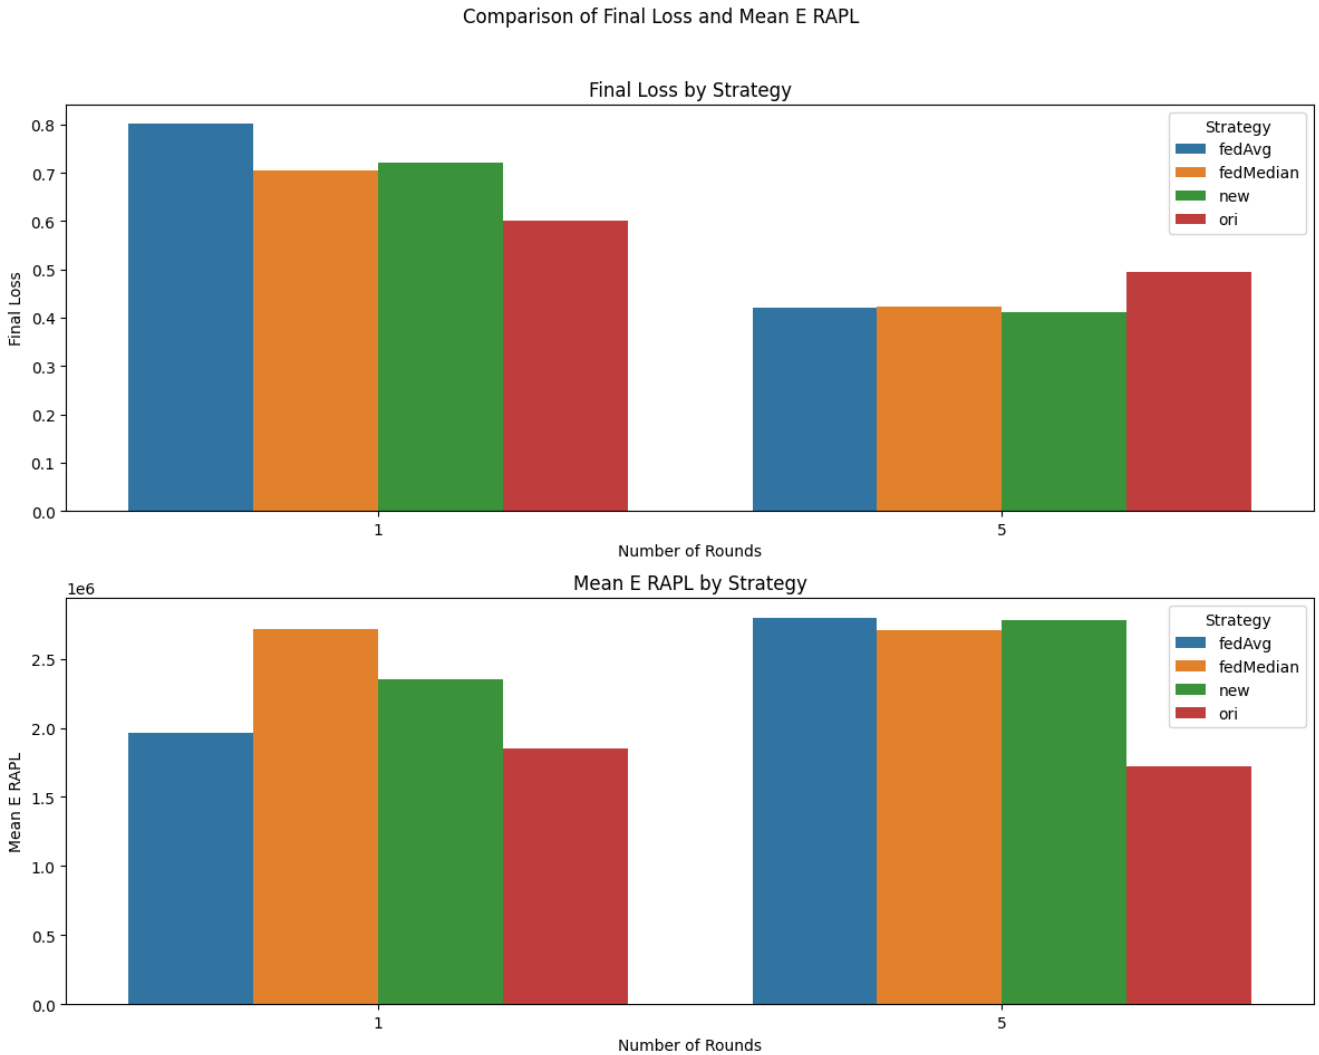
Comparison of Final Loss and Mean E RAPL
Final Loss by Strategy
Final Loss
Number of Rounds
Strategy
fedAvg
fedMedian
new
ori
Mean E RAPL by Strategy
1e6
Mean E RAPL
Number of Rounds
Strategy
fedAvg
fedMedian
new
ori

Thank you!

Flower implement

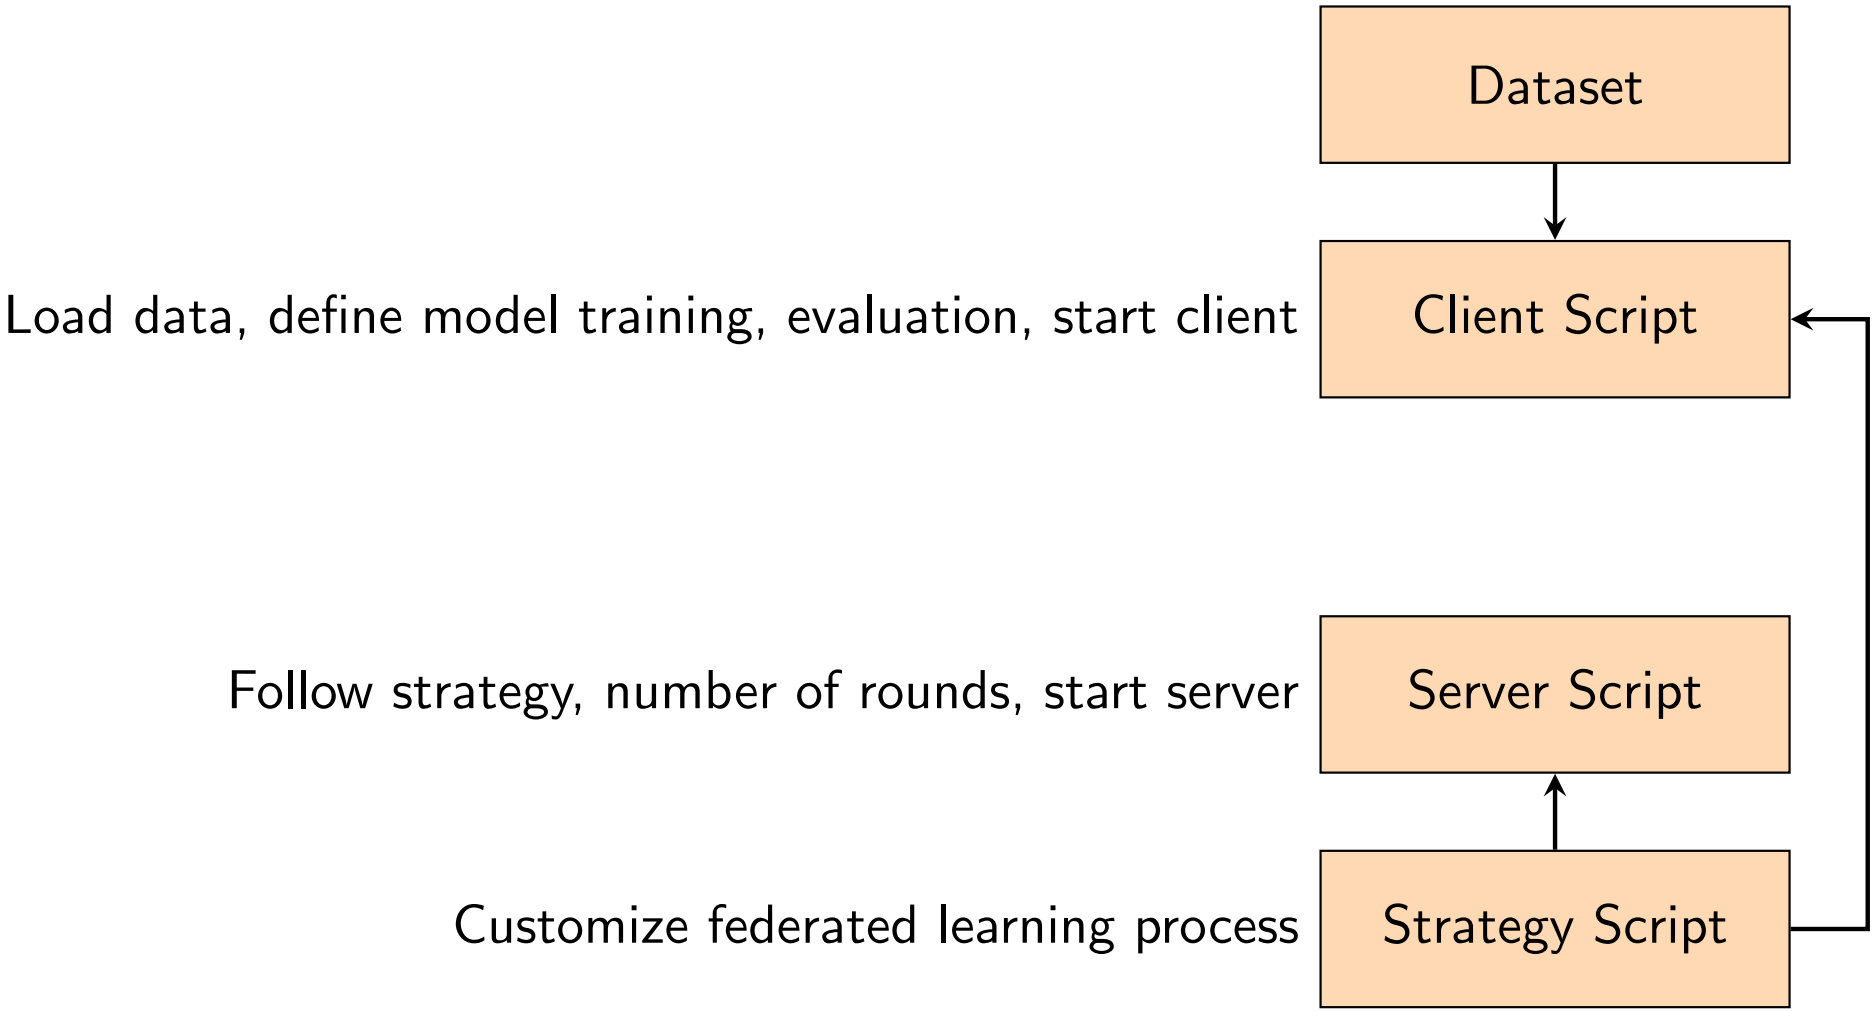

Go back

Energy measurement - Expetator

Run the Flower framework

Define set of sensors are used for measurement

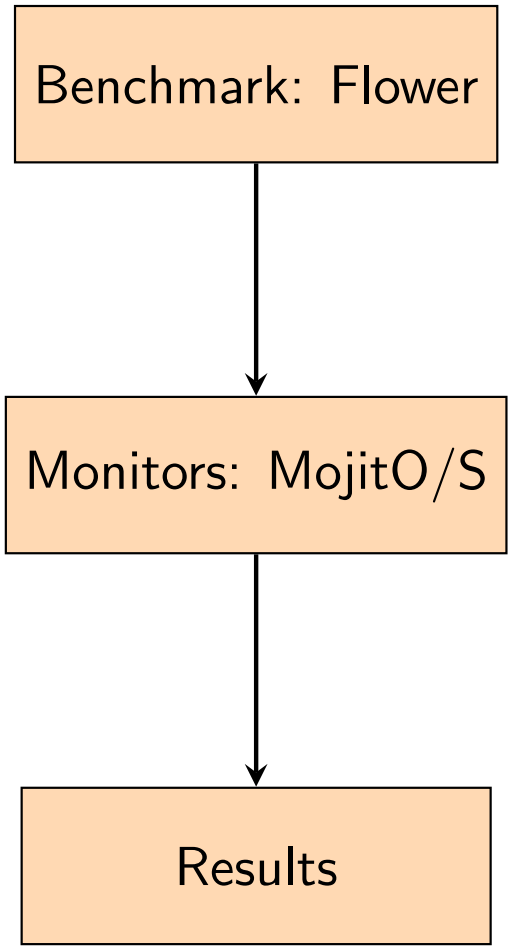

▸ Energy parameters - see more

Energy measurement - Expetator

rxp	number of received packets
rxb	number of received bytes
txp	number of sent packets
txb	number of sent bytes
package	entire sockets
core0	or Power Plane 0, all processor cores on the socket
dram	RAM
idle	no activate status
user	CPU

Go back