

Evaluate All the Things with `benchkit`


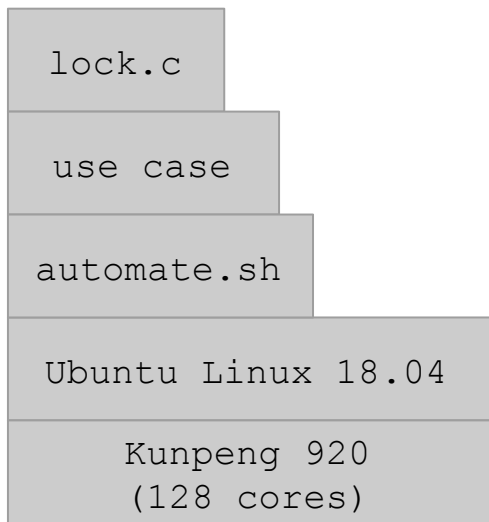
Antonio Paolillo

29th of January, 2025 – FOSDEM BOF Track C

Multicore & Concurrency:
Algorithms, Performance, Correctness



Example use case: evaluating locks



Locks implementation are plenty

e.g. [libvsync](#)


lock.c

use case

automate.sh

Ubuntu Linux 18.04

Kunpeng 920
(128 cores)





libvsync / include / vsync / spinlock /

lilith218 Release v3.6.0 (#12)

| Name |
|------------------|
| .. |
| arraylock.h |
| caslock.h |
| clhlock.h |
| cnalock.h |
| hclhlock.h |
| hemlock.h |
| hmcslock.h |
| mcslock.h |
| rec_mcslock.h |
| rec_seqlock.h |
| rec_spinlock.h |
| rec_ticketlock.h |
| rwlock.h |
| semaphore.h |
| seqcount.h |
| seqlock.h |
| ticketlock.h |
| ttaslock.h |
| twalock.h |

Benchmarking locks in many use cases



| |
|----------------------------|
| lock.c |
| use case |
| automate.sh |
| Ubuntu Linux 18.04 |
| Kunpeng 920 (128 cores) |



Writing shell script to evaluate locks?

lock.c

use case

automate.sh

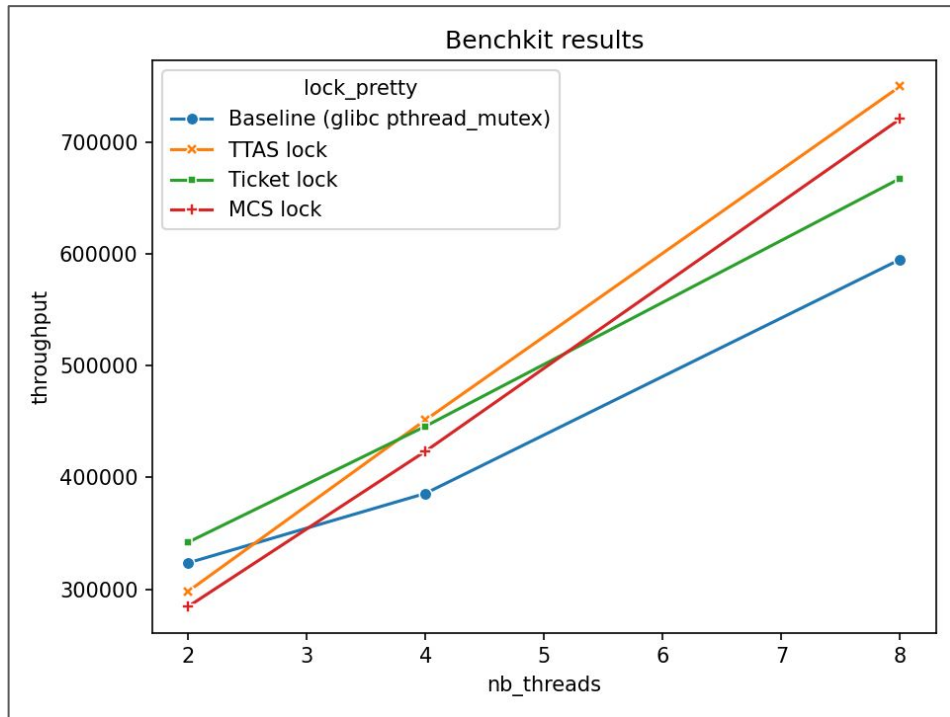
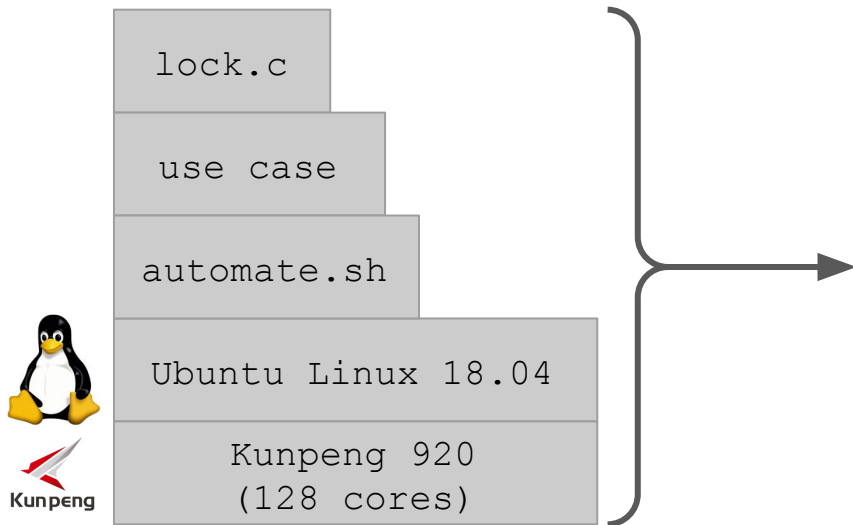
Ubuntu Linux 18.04

Kunpeng 920
(128 cores)



```
nb_threads="1 2 4 8"
bench_name=readrandom
locks="spin mcs ticket clh"
lses="on off"
echo "bench_name;lock;lse;nb_thread;microsop" | tee /tmp/results.csv
for lse in ${lses}
do
  for lock in ${locks}
  do
    make -C ../tilt ${lock} LSE=${lse}
    for nb_thread in ${nb_threads}
    do
      taskset --cpu-list 0 env LD_PRELOAD="../tilt/${lock}.so" ./db_bench
      --benchmarks=${bench_name} --threads=${nb_thread} \
        > /tmp/leveldb_results.txt
      results=$(cat /tmp/leveldb_results.txt | tail -n 1 | grep -o "[0-9.]\+ micros/op" |
cut -d ' ' -f 1)
      echo "${bench_name};${nb_thread};${results}" | tee -a /tmp/results.csv
    done
  done
done
```

Getting some results...



Shell scripts are getting complex quickly

lock.c

use case

~~auto test sh~~

Ubuntu Linux 18.04

Kunpeng 920
(128 cores)




```
nb_threads="1 2 4 8"
bench_name=readrandom
locks="spin mcs"
lses="on off"
echo "bench_name=${bench_name};microsop" | tee /dev/null
for lse in $lses
do
  for lock in $locks
  do
    make -C ../tilt ${lock}
    for nb_thread in $nb_threads
    do
      taskset --cpu=0-127 ./bench
      --benchmarks=${bench_name}
      results=$(./bench ${lock} ${nb_thread} | tail -n 1 | sed 's/micros/op' |
      cut -d ' ' -f 2)
      echo "${lock} ${nb_thread};${results}" | tee -a results.csv
    done
  done
done
```

Benchmark flow

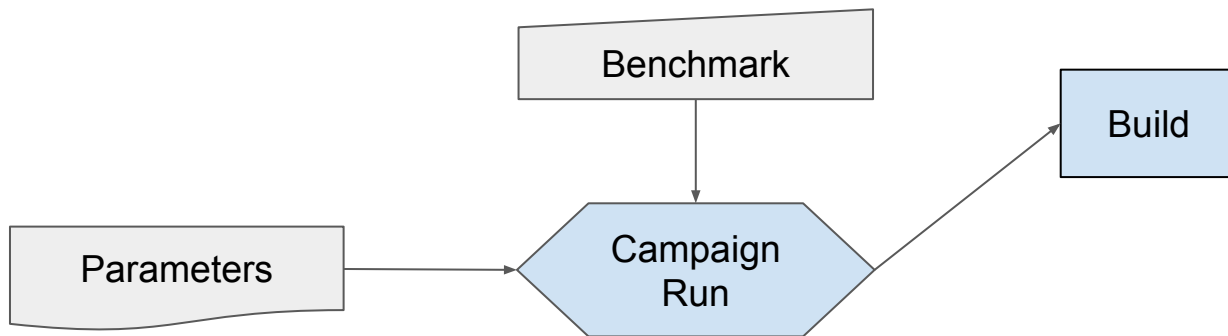
Parameters

Benchmark

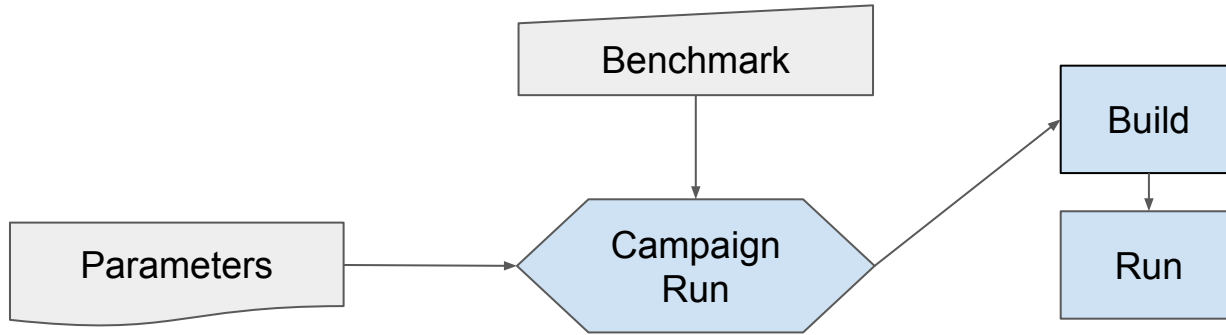
Benchkit flow



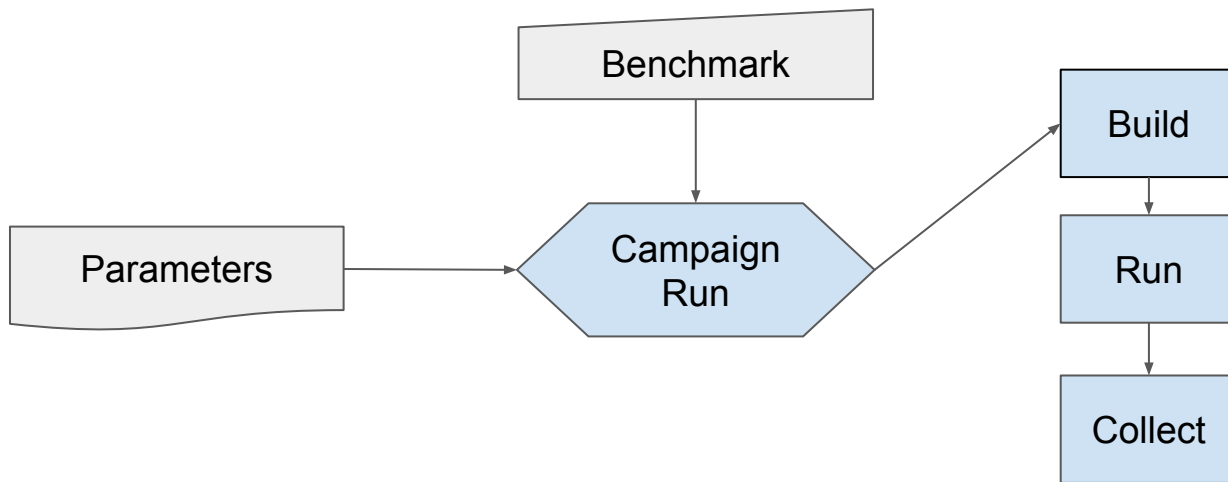
Benchmark flow



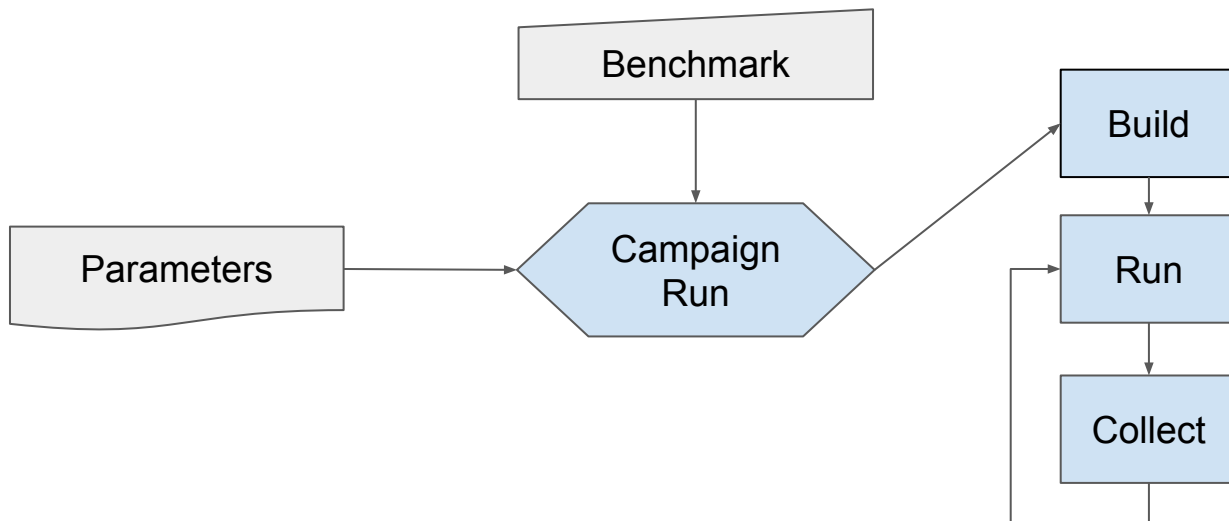
Benchmark flow



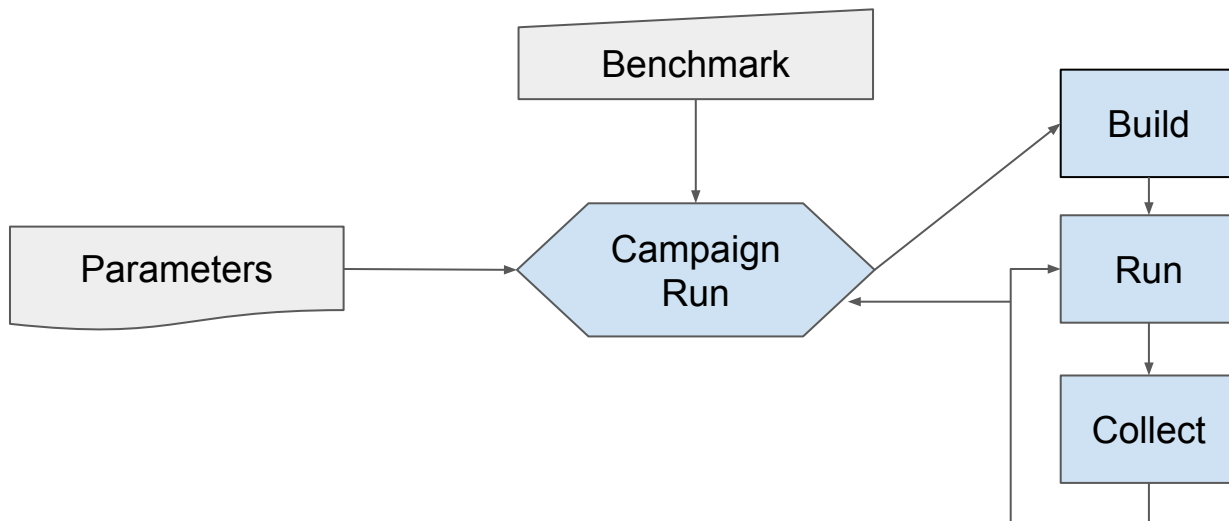
Benchmark flow



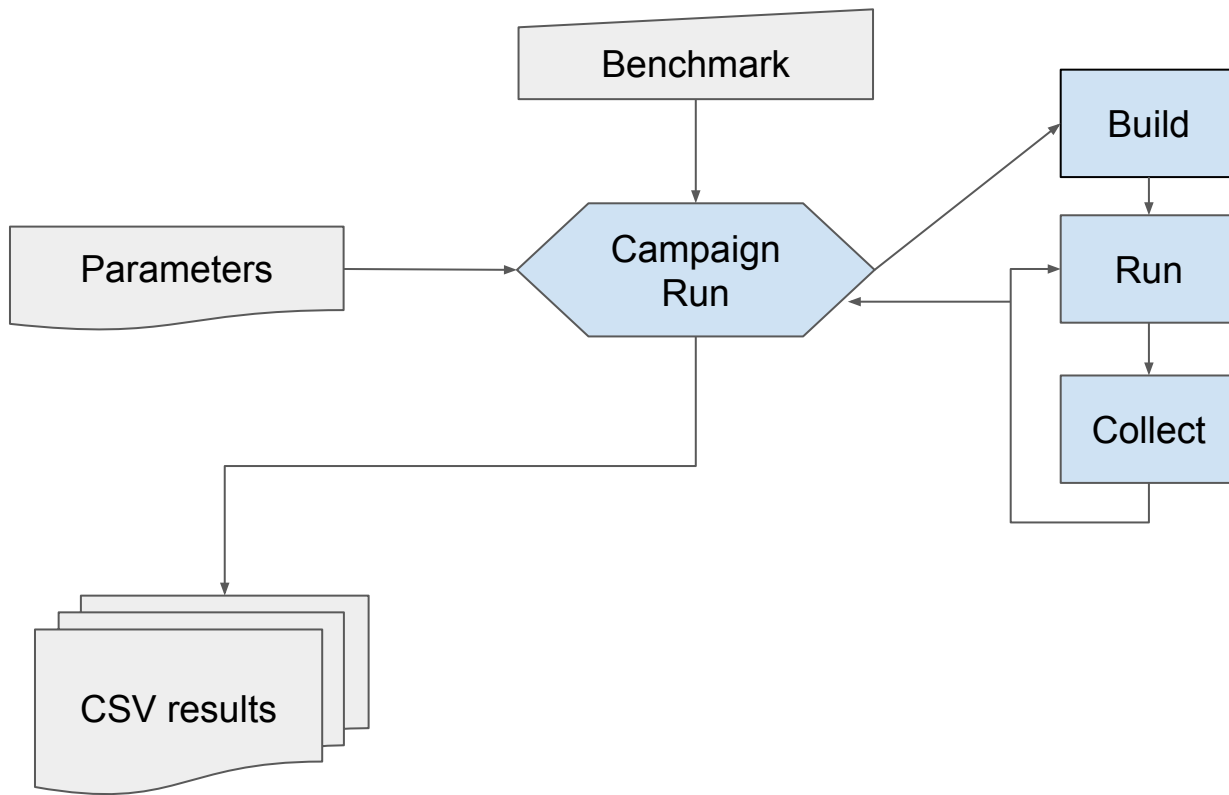
Benchmark flow



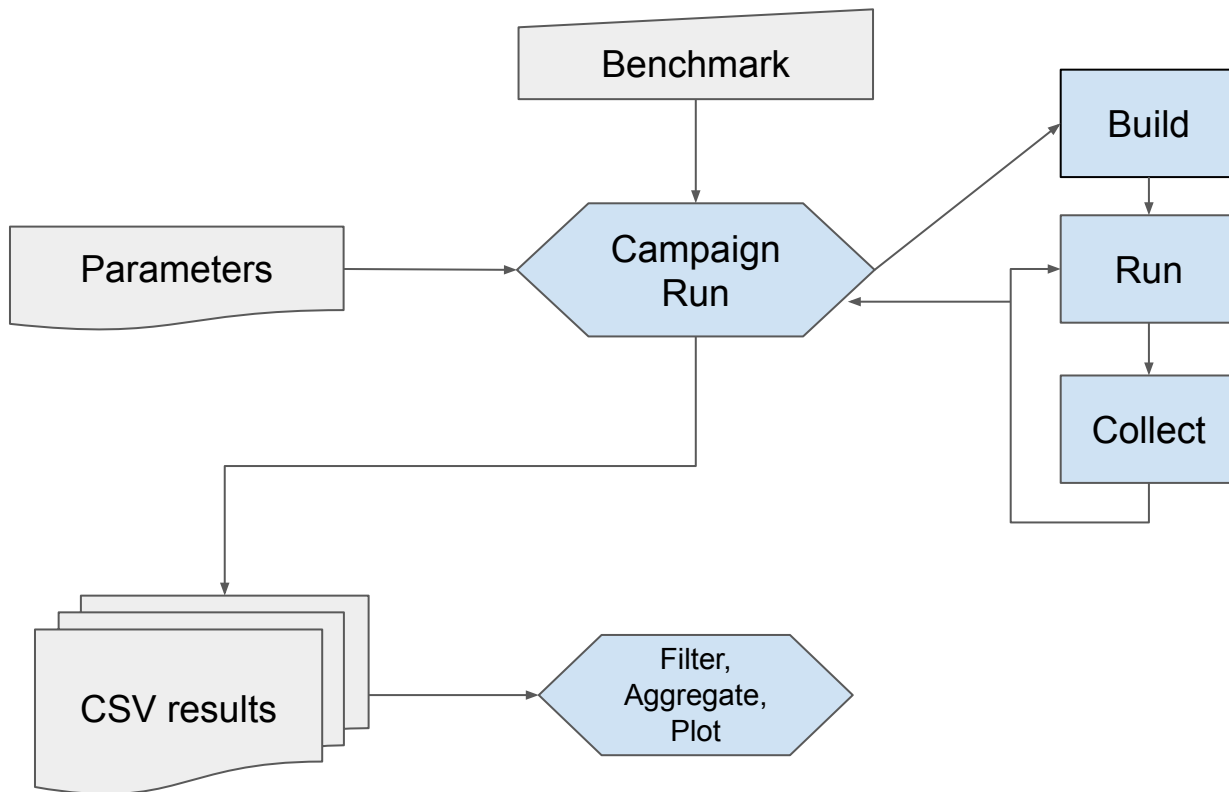
Benchmark flow



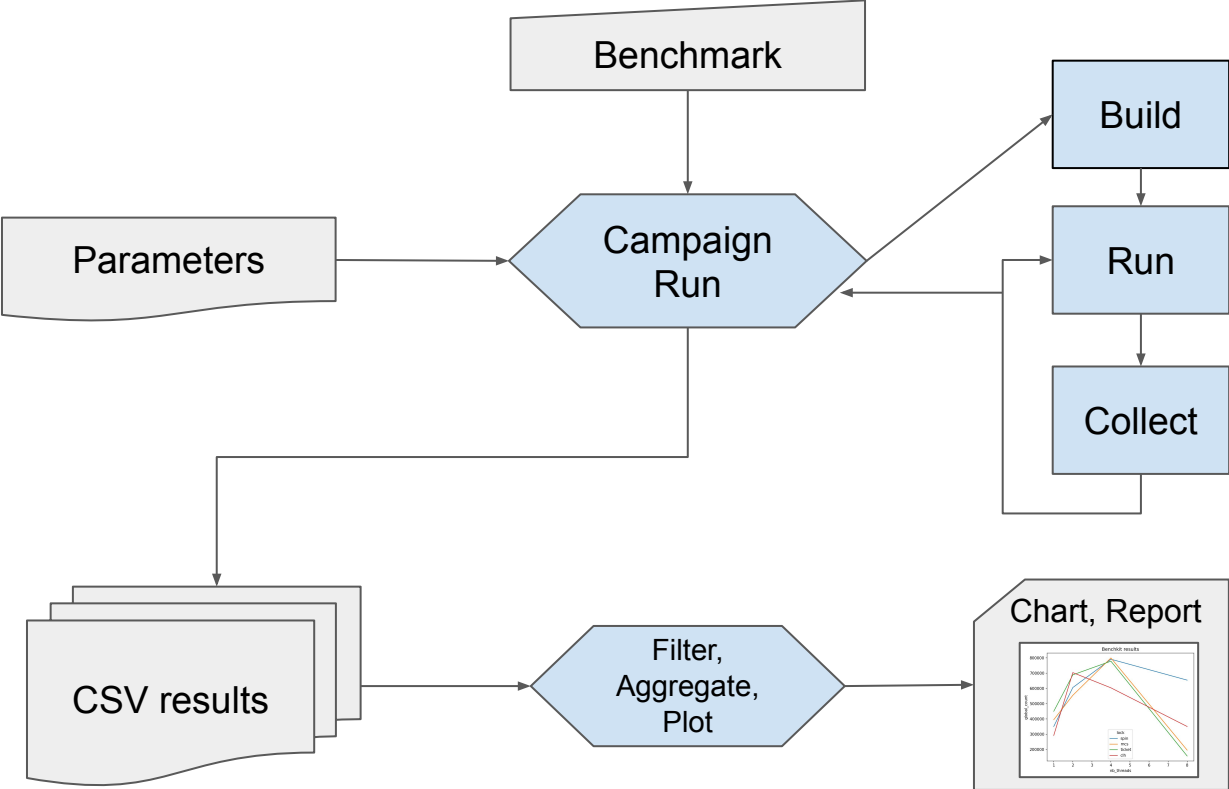
Benchmark flow



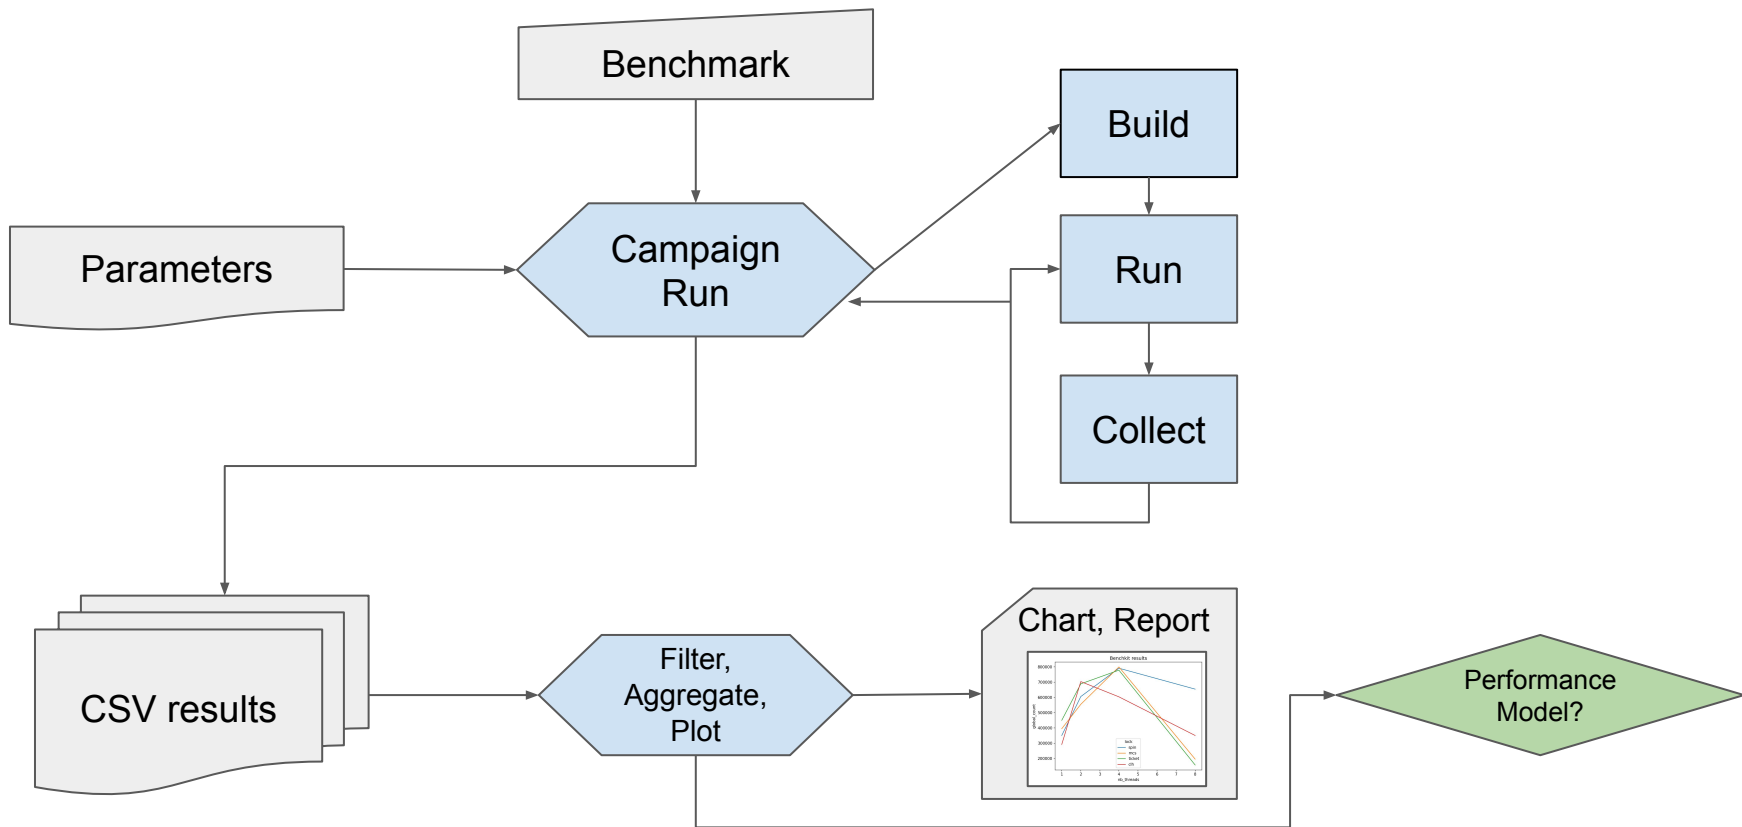
Benchkit flow



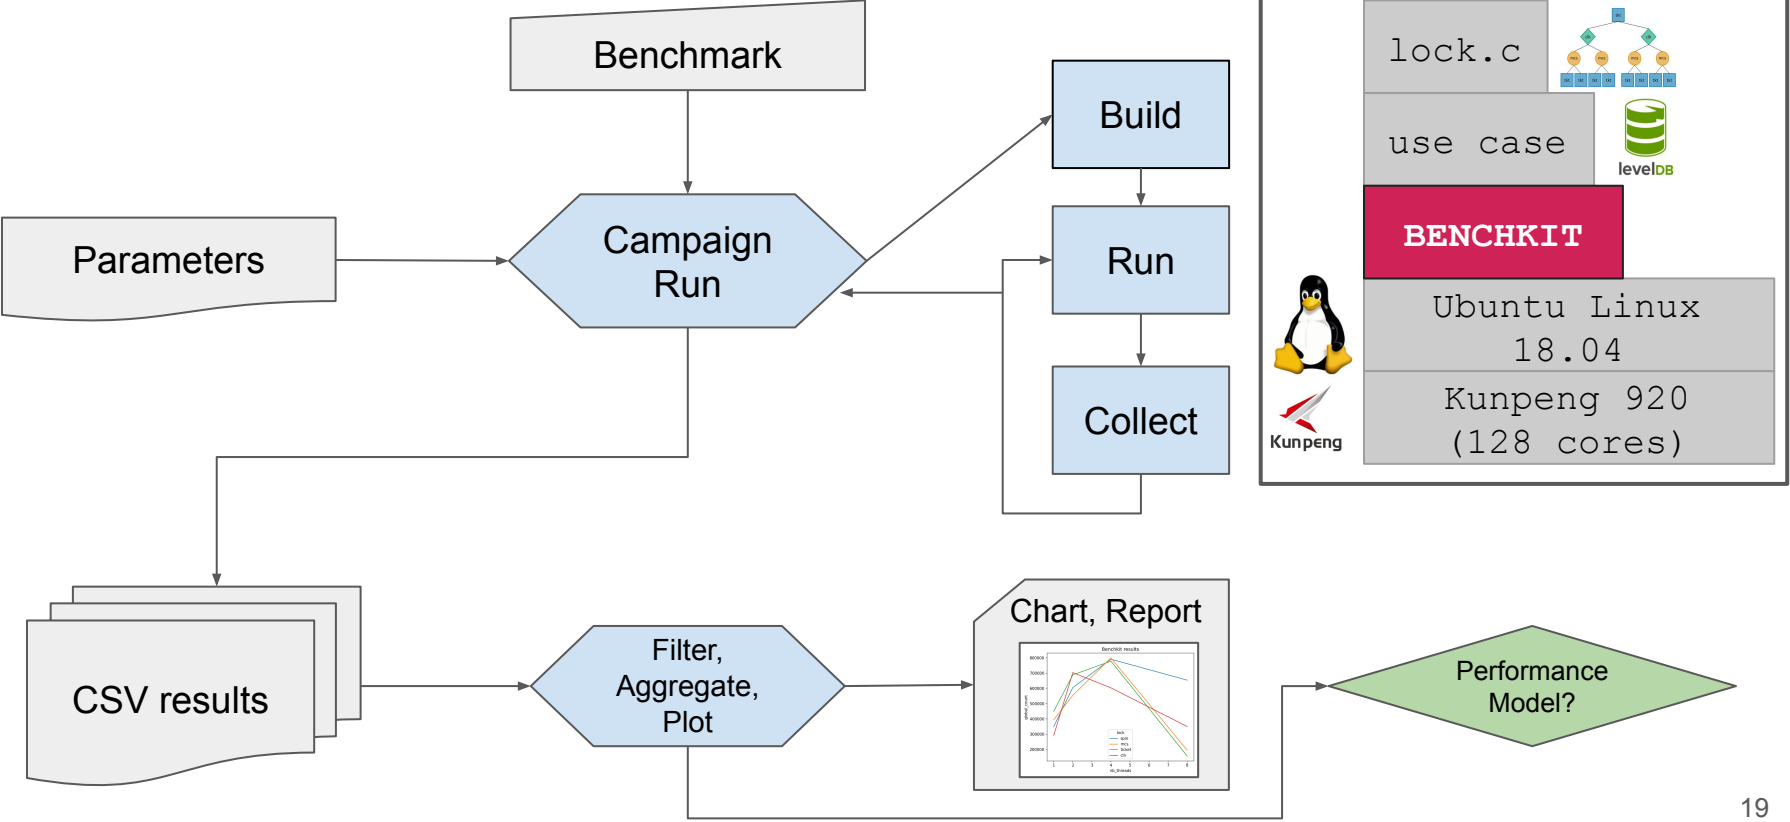
Benchmark flow



Benchmark flow



Benchmark flow



benchkit provides an alternative

lock.c

use case

BENCHKIT

Ubuntu Linux 18.04

Kunpeng 920
(128 cores)



```
nb_threads="1 2 4 8"
bench_name=readrandom
locks="spin mcs ticket clh"
lses="on off"
echo "bench_name;lock;lse;r
for lse in ${lses}
do
  for lock in ${locks}
  do
    make -C ../tilt ${lock}
    for nb_thread in ${nb_thr
    do
      taskset --cpu-list 0 e
--threads=${nb_thread} \
  > /tmp/leveldb_res
  results=$(cat /tmp/lev
  echo "${bench_name};${
  done
done
done

ipython3 < EOF
df = pd.read_csv('/tmp/resu
ax = sns.lineplot(data=df,
                  x='nb_thr
                  y='micros
                  hue='benc


plt.show()
EOF
```

```
if __name__ == '__main__':
    campaign = leveldb_campaign(
        nb_runs=5,
        bench_name=['readrandom', 'fillseq', 'fillrandom'],
        locks=['spin', 'mcs', 'ticket', 'clh'],
        cpu_order=('desc',),
        use_lse=[False, True],
        atomics=('a64', 'blt'),
        nb_threads=[1, 2, 4, 8],
        shared_libs=[get_assignlib(),
                    get_fledgedtilt()],
        command_wrappers=[TasksetWrap()],
    )

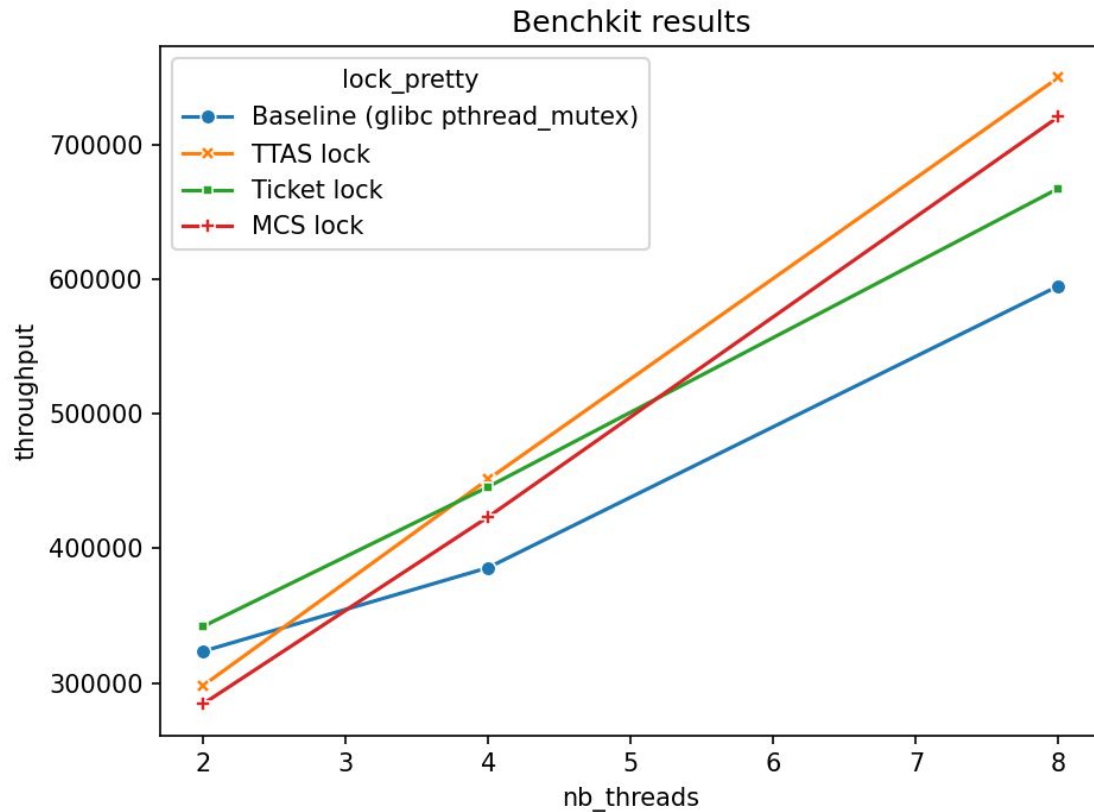
    campaign.run()
    campaign.generate_graph(plot_name='lineplot',
                            x='nb_threads',
                            y='global_count',
                            hue='lock')
```

demo

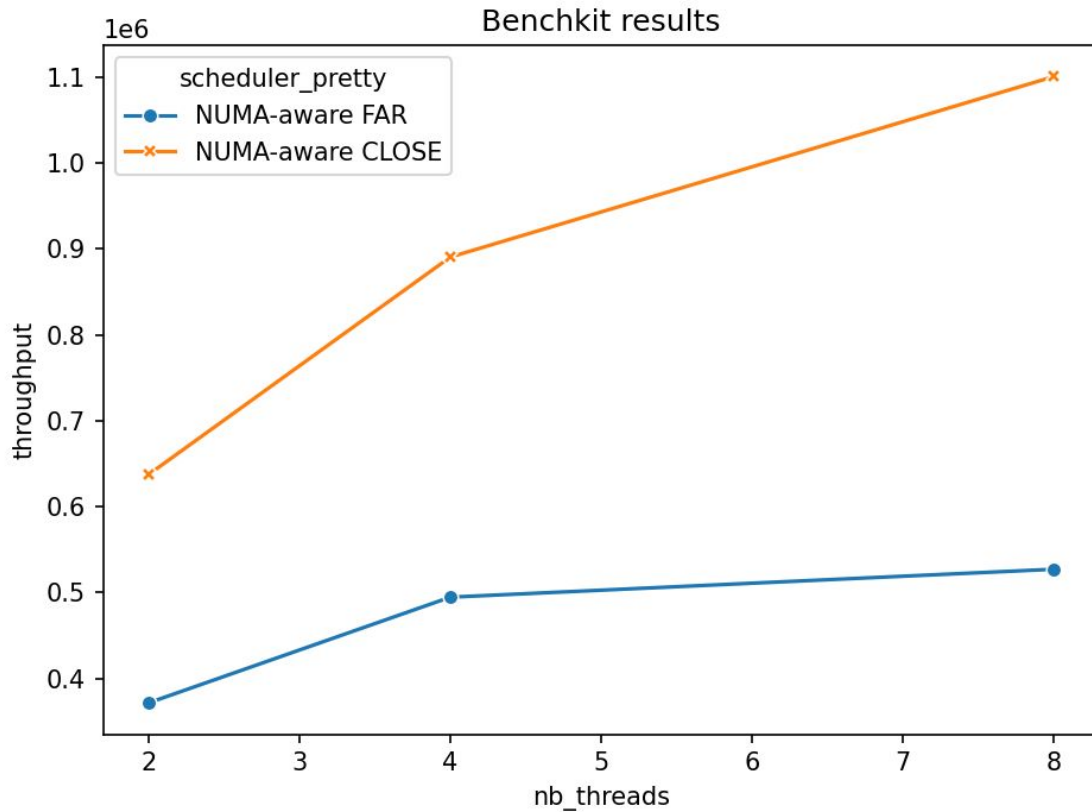
00 - LevelDB




01 - LevelDB - various locks




02 - LevelDB - various schedulers



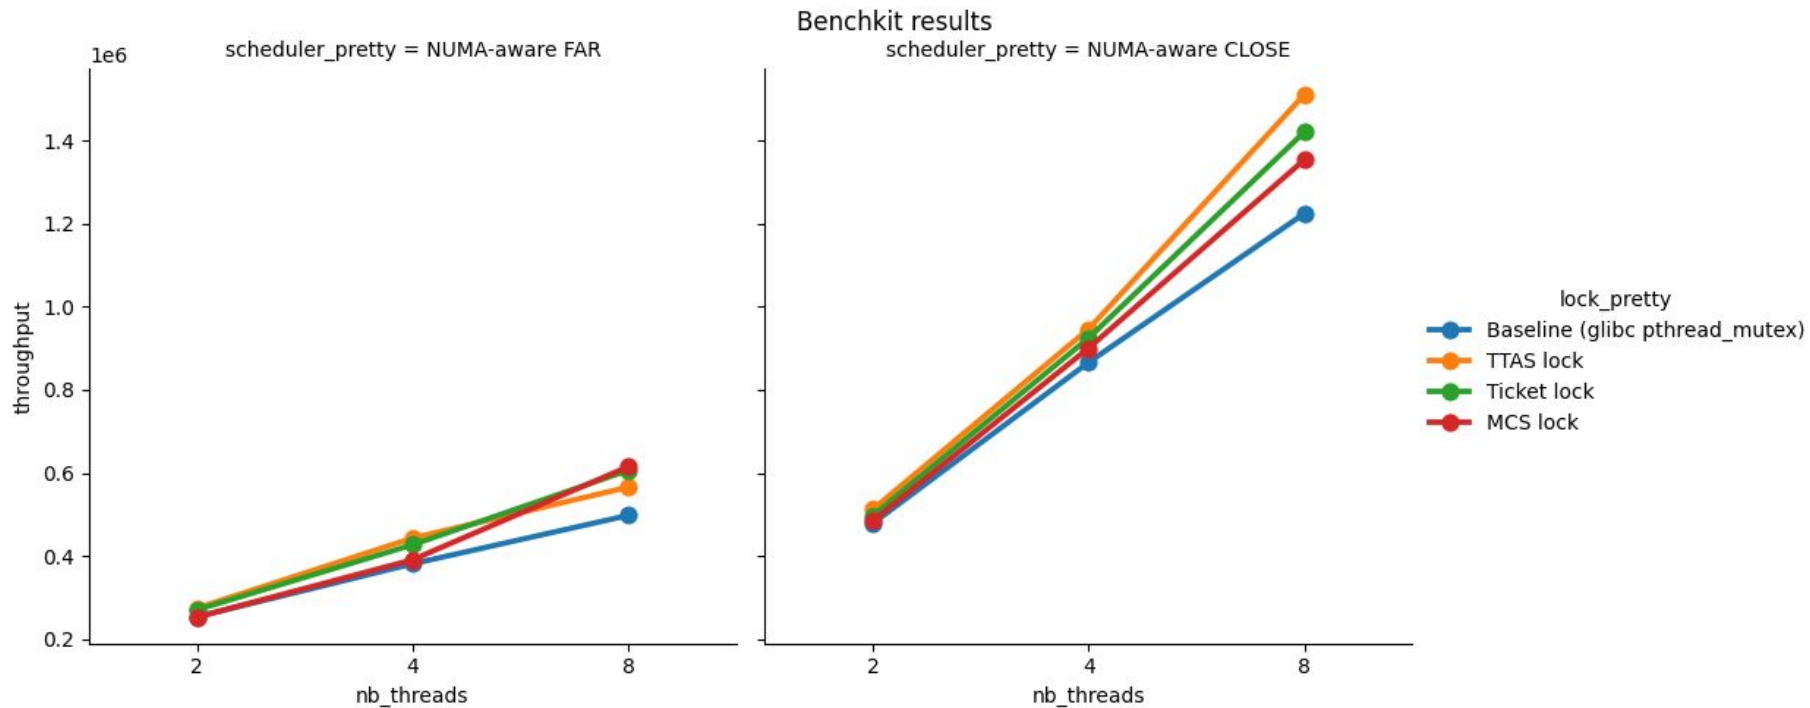
03a - LevelDB - various schedulers & locks



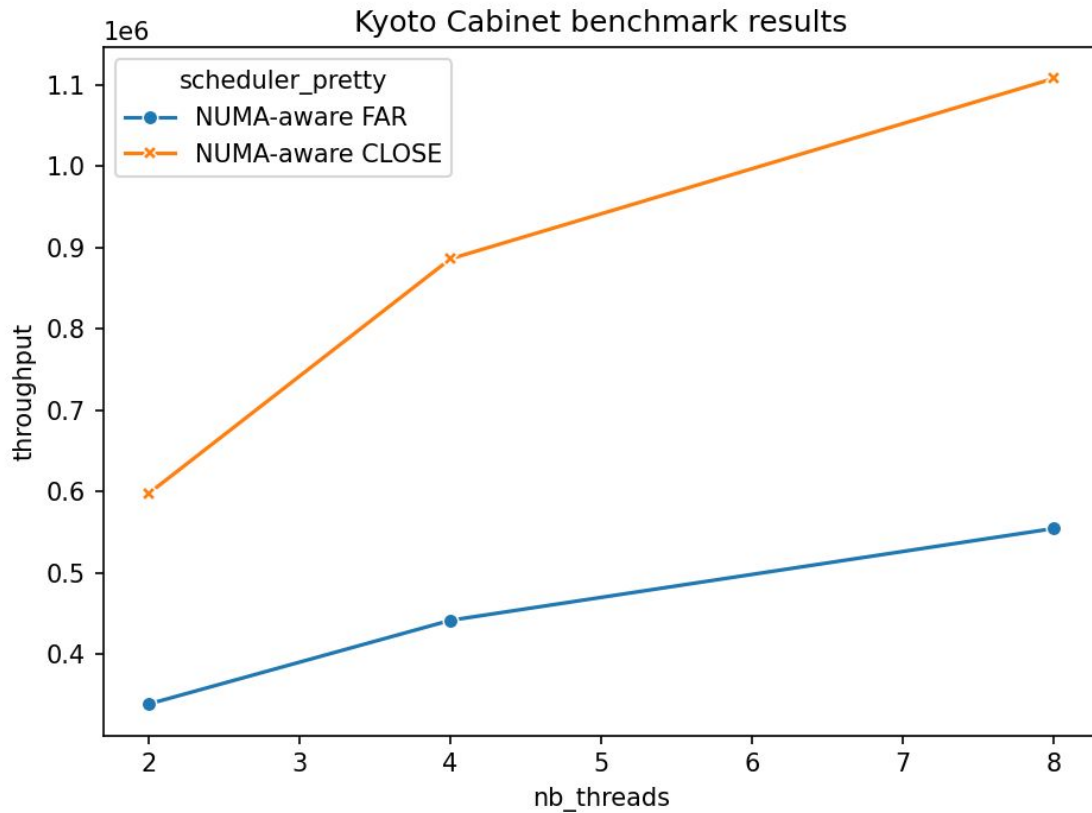
03b - LevelDB - various schedulers & locks



03c - LevelDB - various schedulers & locks



04 - Kyoto Cabinet - various schedulers



Try it! → <https://github.com/open-s4c/benchmark>



On GitHub



MIT license

The screenshot shows the GitHub repository page for `benchmark`. At the top, there is a list of files: `LICENSE` (updated 4 months ago), `MAINTAINERS` (initial public release last year), `README.md` (updated 2 months ago), `ROADMAP.md` (updated last month), and `pyproject.toml` (updated 4 months ago). Below this is the `README` file, which is licensed under MIT. The main heading is **benchmark: Performance Evaluation Framework**. A metadata bar shows `pypi v0.0.1`, `license MIT`, and `downloads 164/month`. The text describes `benchmark` as a push-button end-to-end performance evaluation pipeline. It mentions that users can define a set of experiments called a **campaign**. The **Main principles** section states that the project was born from the need to apply a systematic method to evaluate computer systems.

| File | Description | Last Update |
|----------------|---|--------------|
| LICENSE | Update license to reflect VUB & personal con... | 4 months ago |
| MAINTAINERS | Initial public release | last year |
| README.md | README: Add shields (#135) | 2 months ago |
| ROADMAP.md | Roadmap: add viz tool of firefox (#168) | last month |
| pyproject.toml | Package information for v0.0.1 (#121) | 4 months ago |

Contributors 17

+ 3 contributors

Languages

- Python 98.2%
- Other 1.8%

Interested in *Multicore & Concurrency*?

Contact us!

db7@sdf.org

hernanl.leon@huawei.com

antonio.paolillo@vub.be

