

# Thomas Campistron

## Education

2014–2019 **MSc in Computer Science**, University of Lille, France.



*Relevant coursework:* Virtualization, cloud computing, distributed systems, OS design, system programming, computer architecture, compilation, cryptography

## Work Experience

Mar 2021 – Now **Core team developer**, Meilisearch, Remote.



Meilisearch is an open-source search engine written in rust. I work as much on meilisearch as on milli (the core library of meilisearch). I wrote all the benchmarks, set up a CI to execute the benchmarks on each commit, and push the result to see all the unexpected performance changes. I also fuzzed multiple parts of Meilisearch or milli to reproduce contributors' bugs without any dataset. Here is a little list of fun things I did; I implemented the indexing and searching part of the geosearch, I rewrote the parser entirely for our filters and drastically improved the error messages; We found two tricky bugs and one memory leak in our utilization of LMDB where I used RR to understand what was going wrong with this endless, hideous and goto polluted C code.

On a day-to-day basis, I work with contributors, I answer questions on github issues, discussions or our community slack channel. I fix bugs, refactorize the code, create tooling, improve the CI and do everything you would expect from a developer. Sometimes I read scientific articles and publications on new DBs or distributed systems.

Mar–Aug 2020 **Improve the performance of Mindspore**, Huawei, Remote.



MindSpore is an open-source AI computing framework that implements a lot of state-of-the-art technologies. I was hired to work on an OCaml PoC for simplifying mathematical operations on the control flow graph of MindSpore. The aim of the project is also to maximize the parallelization of the programs and hence get the maximum perfs out of the Huawei Ascend processor or any gpu.

Mar–Aug 2019 **Evaluation of the P4 Language for DDoS mitigation**, Worldline, Lille.



Worldline is one of the European leaders of digital payments. The first aim of this internship was to test a new language called P4. This language stands for “Programming Protocol-Independent Packet Processors”. Our test case was to implement a SynCookie Proxy and to test it against handmade XDP or DPDK solutions.

Apr–Aug 2017 **Optimizing and storing the Firewall rules**, StormShield, Lille.



Stormshield is a European leader in digital infrastructure security, they sell hardware Firewalls. My internship was about optimizing the access to the Firewall rules. Rules are now accessible in read-only by multiple daemons at the same time without any mutex, memory or loading time in  $O(\log_2(n))$  instead of  $O(n^2)$ .

## Technologies

**C** < 100 000 LoC *My professional experience at StormShield was almost entirely in C. I also used C for some personal projects and during my education. I worked with OpenMPI, Valgrind, gdb, SDL2, AFL. I did a bit of IoT with it and it was my main go-to language until I discovered Rust.*

**Ruby, Python** < 10 000 LoC *Mostly used as scripting languages. I know how to extend ruby with C and I did a little bit of C ffi in python. I developed a JIT compiler for a subset of a language in ruby.*

**Haskell, VHDL, Prolog** < 2000 LoC *I enjoyed these languages a long time ago but do not work with it. I would love to though.*

**OCaml** < 10 000 LoC *I learned the OCaml language specifically for my last job at Huawei and apart from some little personal project I used it almost exclusively for this project.*

**Rust** < 150 000 LoC *That's the only language I used for the last 3 years and I plan to continue working with it for a few more years. Here is some old project I did with friends or alone; A bare-metal Rust crate to easily program a teensy3.2. We've also worked on a working async subset of a Minecraft server. I also developed the Mandelbrot and Julia set using parallelization and SIMD. I developed a chip-8, brainfuck, Argh! interpreters. I also often participate in some problem-solving challenges like project Euler and advent of code or the Google Hashcode.*

## Personal

**On my free time** I enjoy thinking of high performance and parallelized code, profiling code to find bottlenecks and then trying to apply as many techniques as possible going from optimizing the algorithm/data structure to things like data-oriented programming, branchless programming, multi-thread or SIMD. I also like bare-metal embedded system (even if I'm not that good at reading the constructor specs). Language and networking are also two domains that draw my attention a lot.

**Travel** I love to travel, so far I visited the following countries:

- USA – Chicago: 2 months
- Japan – Tokyo-Kyoto: 1 month
- Ireland – Dublin: 2 weeks
- Norway – Bergen: 1 week
- Finally, I went back to Japan for 2 more months (September–November 2019) right after my graduation and visited a lot of cities!

But also when it's possible I love to go hiking in the south of France or do bike trips around my city, my longest trip was 160Km in one day but I usually stay below 100Km per day.

**Other** I listen to **a lot** of music. I'm almost always listening to something, usually from the 70'. I also love good food, and especially good cheeses, cooking or playing video games. And finally, I have a strange addiction for keyboards that runs on QMK, currently, I own an ergodox-ez, a preonic and I should get a Kyria in the following days.

Here is my preonic

