# AdaptivePerf: a portable, low-overhead, and comprehensive code profiler for single- and multi-threaded applications

Maksymilian Graczyk (CERN), Stefan Roiser (CERN)

## What is AdaptivePerf?

- **Open-source code profiler** for Linux, based on "perf" with custom patches
- **Traces every spawned thread and process**
- **Fixes "perf"'s broken stacks** (compiling a tested program with frame pointers required)
- Samples **both on-CPU and off-CPU** activity
- Produces **interactive flame graphs and charts** viewable in a web browser
- Main functionality designed with **hardware portability** in mind
- Supports custom sampling-based "perf" events for **profiling interactions with hardware**
- **Detects automatically inappropriate kernel and CPU configurations**
- **Allows TCP streaming** of profiling data to a separate machine for real-time processing

## How does it compare to similar maintained profilers?

| | Hardware-vendor-portable | Profiles software-hardware interaction* | Low profiling overhead | Open-source | Off-CPU profiling | Heterogeneous architecture support |
|---|---|---|---|---|---|---|
| **AdaptivePerf** | **Yes** | **Yes** | **Yes** | **Yes** | **Yes** | **Planned!** |
| Original "perf" | Yes | Yes | Yes | Yes | Limited | No |
| Intel VTune Profiler | No | Yes | Yes | No | Yes | Intel GPUs/FPGAs only |
| AMD μProf | No | Yes | Yes | No | Yes | AMD GPUs only |
| valgrind | Yes | No | No | Yes | No | No |
| gprof | Yes | No | Needs CI** | Yes | No | No |
| gperftools | Yes | No | Needs CI** | Yes | No | No |
| NVIDIA profilers | No | Yes | No | No | Yes | NVIDIA GPUs only |

*If supported by a user's hardware architecture.
**Code instrumentation other than not omitting frame pointers.

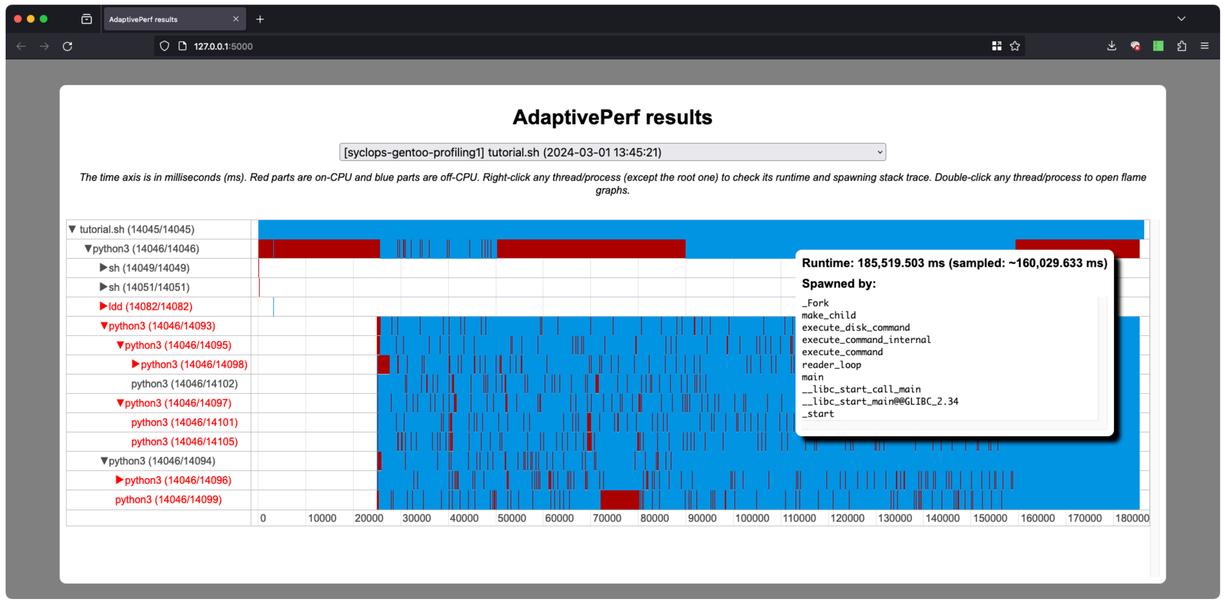


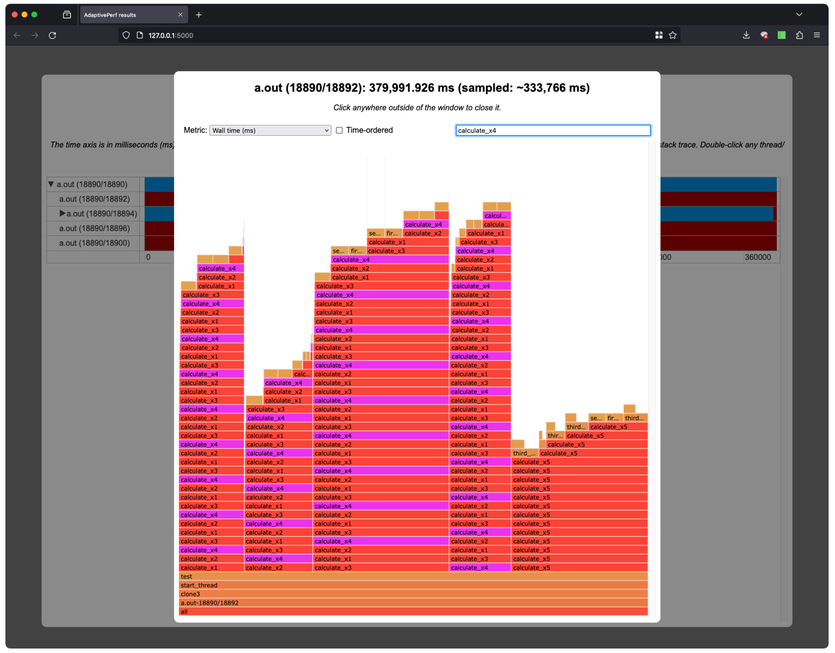


Screenshots as of 4 March 2024. AdaptivePerf is in beta and evolving, so it may differ now.

## Envisaged applications

- Profiling physics data analysers and simulators (e.g. ROOT, Madgraph5, Geant4)
- Profiling software used for online and offline computing at physics experiments
- Software-hardware co-design (e.g. RISC-V core customisation done in the SYCLOPS EU project, triggering and DAQ system development at the LHC experiments)
- And more!

## Planned improvements

- Profiling wide-ranged heterogeneous architectures (RISC-V, GPUs, FPGAs, TPUs, and other accelerators) in a maximally open-source way
- Matching non-sampling-based metrics (such as power consumption) to code segments
- Removing or weakening the frame pointer compilation requirement
- Other feedback and suggestions for improvements are welcome!

Developed in the context of the SYCLOPS project, which is funded by the European Union HE research and innovation programme under grant agreement No 101092877.

**Scan me for the information slides + the download instructions or ask for a demo!**

You can also send a message: maksymilian.graczyk@cern.ch