

AutoBench: Finding Workloads That You Need Using Pluggable Hybrid Analyses

Researchers often rely on benchmarks to demonstrate feasibility or efficiency of their contributions. However, finding the right benchmark suite can be a daunting task — existing benchmark suites may be outdated, known to be flawed, or simply irrelevant for the proposed approach. Creating a proper benchmark suite is challenging, extremely time consuming, and also — unless it becomes widely popular — a thankless endeavor. In this paper, we introduce a novel approach to help researchers find relevant workloads for their experimental evaluation needs. Our approach relies on the huge number of open-source projects available in public repositories, and on unit testing having become best practice

in software development. Using a repository crawler employing pluggable static and dynamic analyses for filtering and workload characterization, we allow users to automatically find projects with relevant workloads. Preliminary results presented here show that unit tests can provide a viable source of workloads, and that the combination of static and dynamic analyses improves the ability to identify relevant workloads that can serve as the basis for custom benchmark suites.

Author: Walter Binder (walter.binder@usi.ch)