

SABS: An Adaptive Benchmarking Solution

Maxime Chevalier-Boisvert

In the world of academic research, and particularly in the areas of computer science and engineering, there is often a need to perform benchmarks and to gather various metrics in order to compare the characteristics or measure the performance of various software systems. This process is usually automated using simple scripts written by members of research groups themselves. These scripts are most often custom tailored for the needs of each specific research applications and make it difficult to obtain data in a consistent and comparable manner, particularly over different computing platforms. Furthermore, the fact that such scripts are often being rewritten by many different individuals brings in the potential for errors.

Our goal was to design and implement an adaptive benchmarking system that would be sufficiently adaptive to suit the very diverse needs of the academic research world, flexible enough to work with a wide variety of different programming languages, platforms and configurations and easily extensible by anyone who wishes to add new features. This system also needed to be complete enough so that it would not require any modifications for the majority of typical uses. We present you our current design and implementation of this system, which we have called SABS (Sable Adaptive Benchmarking System).