

Constructing Backpressure Resistant Closed-Loop Stream Benchmarks

Iain Dixon, Matthew Forshaw, Joe Matthews

Stream processing system (SPS) benchmarks allow practitioners to load a system with realistic and representative workloads and assess its performance. These benchmarks can be constructed as open-loop, where the benchmark generator and system under test (SUT) are formalised and independent, and closed-loop, where this separation is not explicitly guaranteed. Due to their simplicity, closed-loop designs are common both in bespoke benchmarks (YCSB, NEXMark) as well as the benchmarks provided by SPS providers like Apache Flink (playground suite).

In this presentation we will demonstrate how the closed-loop paradigm can result in generators becoming backpressured, thereby failing to achieve the load from desired arrival processes. We present a conceptual model of SPS rate-throttling mechanisms such as load shedding and backpressure, and where these mechanisms operate differently under closed-loop when compared with deployment. We implement a generator cutoff to mitigate this issue and walk the audience through the complexity of this space by implementing various representative arrival processes, starting with a simple uniform arrivals case to provide intuition of the problem and moving to jump and Poisson arrival processes to demonstrate realistic workloads. Our work with static, windowed, and dynamic cutoffs presents an exploration into how careful benchmark implementation can produce accurate and non-confounding benchmark loads in a closed-loop stream benchmark. We hope this presentation will help guide future benchmark development and help future benchmarking practitioners basing their analysis on confounded results.