Advanced Operating Systems:
Lab 1 – Getting Started with Kernel Tracing / I/O
Part II Assignment

Prof. Robert N. M. Watson

2021-2022

This lab assignment compares several configurations of the benchmark, exploring (and explaining) performance differences between them, as well as the impact of the probe effect arising from DTrace use.

Submitting your completed assignment

Please submit your solution in the form of a single PDF interleaving written answers, plots, tabular data, and source-code excerpts, generated from your JupyterLab notebook. The easiest way to do this is via the Print menu option on JupyterLab's File menu. All submissions are via the course's Moodle page.

Experimental questions

1. Performance

Create a plot illustrating I/O bandwidth across the full range of buffer sizes (powers of 2 from 32 bytes to 16MiB) and a fixed total I/O size (16MiB).

2. Kernel statistics

The benchmark is able to query various kernel statistics about the running process using the `getrusage(2)` API. Considering the `inblock`, `time`, and `stime` metrics, what conclusions can we draw about this workload?

3. Kernel profiling

Using DTrace's profile provider and the `stack()` function, determine for each buffer size what the dominant consumers of system (kernel) time are. Create an annotated plot and explanation as follows:

- Partition the performance graph from (1) by key inflection point, shading each region of the graph.
- For each partition, enumerate the top 5 consumers of CPU time, sorted from highest to lowest.
- Where the kernel behaviour of one partition behaves substantially differentially from its adjacent partition(s), explain how and why.
- Label the graph clearly indicate the potential sources of inflection-point changes, and the behaviours present in each region.

Please include your D scripts as part of your submission.

4. Probe effect

Gather performance benchmark results across the buffer-size parameter space using: (a) no DTrace script running; and (b) a DTrace script that counts the number of `read(2)` and `write(2)` system calls.

- Plot the two performance curves on the same graph, labelling them clearly.
- Briefly summarise any performance differences between the two sets of measurements.
- Despite these differences, do we have reason to think that the insights from our DTrace analysis have useful explanatory power?

Please ensure that your D scripts are part of your submission.