`chunkedseq` Benchmarking User's Guide

Umut Acar Arthur Charguéraud Mike Rainey

5 August 2014

Synopsis

chunkedseq.byte [*ACTION*] [*PARAMETERS*]...

Description

chunkedseq.byte is a script whose purpose is to benchmark the implementation of our chunked-sequence data structure. The script automates all aspects of the benchmarking: building binaries, generation of input data, running of experiments, and output of experimental data, including plots, tables, and raw data.

Options

Actions

The action selects the overall behavior of the script. *ACTION* can be one of the following:

`configure` Generate configuration files that are required by PASL.
`generate` Generate graph files that are to be used by the graph-traversal experiments.
`fifo` Run the "fifo" benchmark.
`lifo` Run the "lifo" benchmark.
`split_merge` Run the "split-merge" benchmark.
`bfs` Run the serial BFS benchmark.
`dfs` Run the serial DFS benchmark.
`pbfs` Run the parallel BFS benchmark.
`filter` Run the parallel filter benchmark.
`map` Run the dynamic-dictionary benchmark.

`report` Generate a table reporting on execution times of benchmarks runs.
`all` Build binaries, generate graphs, and then run all benchmarks.

Parameters

Parameters select finer details of the behavior of the script. *PARAMETERS* can be zero or more of the following:

`-runs` ***n*** Specifies the number of times n to execute for each combination of benchmark parameters. Default is 1.
`-timeout` ***n*** Force a specific timeout for runs. The timeout value n is measured in seconds.
`-mode` ***m*** Where m is `normal` (discard all previous results) or `append`. (append to previous results) or `replace` (discard results that are ran again) or `complete` (add results that are missing).
`--virtual_run` Only show the list of commands that would be called to run the benchmarks.
`--virtual_generate` Only show the list of commands that would be called to generate the graphs used by the graph-search experiments.
`-skip` ***a1*,*a2*,...** Skip selected actions. Note: `-skip run` automatically activates `-skip make`.
`-only` ***a1*,*a2*,...** Perform only selected actions.
`-path_to_graph_data` ***PATH*** Force a specific path in which to store graph data. Default is `_data`.
`-path_to_pasl` ***PATH*** Force a specific path to the root of the PASL source tree. Default is `..`

Compilation parameters

`-allocator` ***a*** Select a drop-in replacement for malloc/free by specifiying a custom library *a*.
`--use_hwloc` Compile PASL binaries with support for hwloc.
`-path_to_`***PACKAGE PATH*** Configure PASL to look for package named *PACKAGE* in the path *PATH*.

Sample applications

- Configure all PASL binaries to use hwloc.

  ```
  ./chunkedseq.byte configure --use_hwloc -path_to_hwloc /pathto/hwloc
  ```

- Configure all PASL binaries to use tcmalloc.

  ```
  ./chunkedseq.byte configure -allocator tcmalloc -path_to_tcmalloc /pathto/tcmalloc
  ```

- Run all experiments but do not plot.

  ```
  ./chunkedseq.byte -skip plot all
  ```

- Run just graph experiments and neither plot nor build dependencies.

  ```
  ./chunkedseq.byte -skip plot,make dfs
  ./chunkedseq.byte -skip plot,make bfs
  ./chunkedseq.byte -skip plot,make pbfs
  ```

- Generate a table reporting on the three experiments from above.

  ```
  ./chunkedseq.byte report
  ```

See also

The chunkedseq source code and all documentation can be downloaded from http://deepsea.inria.fr/chunkedseq/