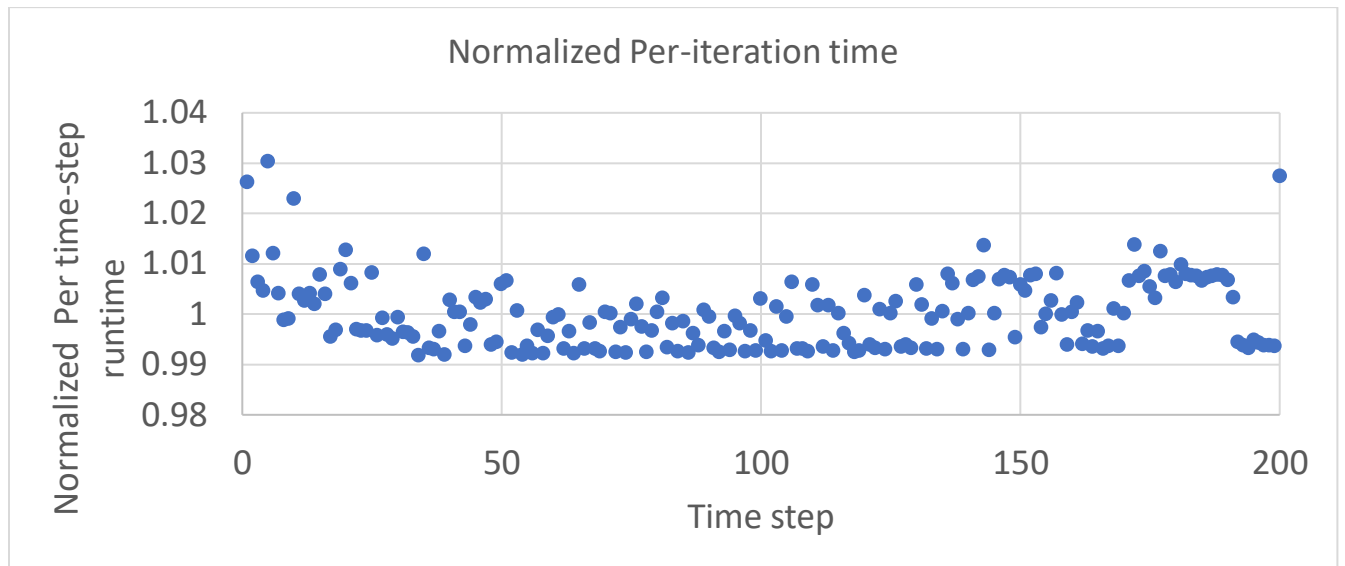



1. Per-program, is there time step-dependent performance behavior?



There is trivial time step-dependent performance behavior.

2. Per-loop per-time step runtime is also with similar trend of variations.



Per time-step normalized runtime of Cleverleaf kernel accelerate1, across 200 time steps.

Each iteration:

76/113 kernel invocations of the 33 kernels lasts less than 20 ms (grid size $\leq 25 \times 25$)

37/113 last 7500~40000ms (grid size are always: 4000×4000)

This means:

1. 76/113 kernel executions should use serial code path
2. 37/113 kernel executions should use OMP code path
3. The separation boundary for this two kernels is grid size.

3. DT (time step) is used rarely (5 out 33 kernels), as a multiplication factor within loops

e.g,

$$a = dt * b$$

which does not have observable impacts on performance.

4. What is next?

Cleverleaf w/ more inputs

Kripke

LULESH

strong scaling: will change the amount of work for different configurations.

Lulesh

#MPI only first, maybe hybrid w/ OMP later.

vary #MPI ranks/nodes, #nodes

same inputs but different code variants

baseline always O3