

RM-Replay for Cluster Scheduling Project Report 3/15-3/27

Zhen Huang, Blake Ehrenbeck

These two weeks we are trying to examine the correctness of the RM-Replay that we installed, since we were uncertain about the output of metric.log in regard to utility and average wait time as well as work on running some Theta jobs in RM-Replay. We studied the trace_metric.c to understand the output and compared to the binary trace file, and we concluded that our RM-Replay should be working correctly. We are also nearly completion of importing the data from Theta to our new database.

Progress:

1. The output of metric.log was wrong because there is a bug in the trace_metrics.c, we found that in the *compute_metrics* method starts from line 362 each categories (mc, GPU and all) is calling the function *computer_metrics*.

```
361 }
362 printf("all=%llu preset=%llu (otherp=%llu)\n", njobs, njobs preset, njobs otherp);
363 if (c_mc && c_gpu) {
364     compute_metrics(job_arr_all, njobs_all, Nnodes, &makespan, &util, &avg_wait, &std_wait, &min_wait, &max_wait, &
365     printf("[ALL=%llu]\t Makespan=%ld\t Util=%.8f\t Avg Wait=(%.8f,%.8f,%ld,%ld,%ld,%.4f)\t Dispersion=%.8f Slowdown
min_wait, max_wait, coefvar_wait, dispersion, slowdown, nthrougput_mc+nthrougput_gpu);
366 }
367
368 if (c_mc) {
369     compute_metrics(job_arr_mc, njobs_mc, Nnodes_mc, &makespan, &util, &avg_wait, &std_wait, &min_wait, &max_wait,
370     printf("[MC=%llu]\t Makespan=%ld\t Util=%.8f\t Avg Wait=(%.8f,%.8f,%ld,%ld,%ld,%.4f)\t Dispersion=%.8f Slowdown
n_wait, max_wait, coefvar_wait, dispersion, slowdown, nthrougput_mc);
371 }
372
373 if (c_gpu) {
374     compute_metrics(job_arr_gpu, njobs_gpu, Nnodes_gpu, &makespan, &util, &avg_wait, &std_wait, &min_wait, &max_wai
375     printf("[GPU=%llu]\t Makespan=%ld\t Util=%.8f\t Avg Wait=(%.8f,%.8f,%ld,%ld,%ld,%.4f)\t Dispersion=%.8f Slowdown
min_wait, max_wait, coefvar_wait, dispersion, slowdown, nthrougput_gpu);
376 }
377
378 free(job_arr);
379 free(job_arr_all);
380 free(job_arr_mc);
```

However, the function *compute_metrics* doesn't reset the variable *nwait* to 0, so the output in metrics was off by one when the metrics output get printed out, and we get the wrong number of waits (should be 1 but 2):

```
all=1 preset=0 (otherp=0)
4
[ALL=1] Makespan=198 Util=0.37500000 Avg Wait=(4.00000000,0.00000000,1.4,4,0.0000) Dispersion=1.00000000 Slowdown=1.02020202 Throughtput=1
[MC=0] Makespan=9223372036854775807 Util=0.00000000 Avg Wait=(0.00000000,0.00000000,1.9223372036854775807,0,-nan) Dispersion=-nan Slowdown=-nan Throughtput=0
4
[GPU=1] Makespan=198 Util=0.42857143 Avg Wait=(2.00000000,1.41421356,2.4,4,0.7071) Dispersion=0.58578644 Slowdown=1.02020202 Throughtput=1
[slurm@1838cc774d52 submitter]$ ls
```

The format of Avg_Wait(...) is as follows:

(Average Wait, Standard Deviation Wait, Number of Waits, Min Wait, Max Wait, Coef Var Wait)

After we change the code by resetting `nwait` to 0 every time it gets called, we get the correct output:

```
void compute_metrics(job_trace_t *job_arr, unsigned long long njobs, unsigned long long Nnodes, long
t, double *dispersion, double *slowdown, double *coefvar_wait)
{
    *nwait=0;
    unsigned long long j;
    long *time_submit_arr;
    long *time_exec_arr; // end-start
    long *time_wait_arr; // start-eligible
```

```
[slurm@1838cc774d52 submitter]$ ./trace_metrics -w /slurm/data/replay.03.23.trace.03.23.testing..0001.27/replay.03.23.trace -r 4
Range: min_start=1553366933 [0,4] start_range=1553366933 end_range=1553381333
all 1 preset 0 (otherp 0)
Start: 1553366933Submit: 1553366929waited for 4[ALL=1] Makespan=198 Util=0.37500000 Avg_Wait=(4.00000000,0.00000000,1,4,4,0.0000)
[MC=0] makespan=9223372036854775807 Util=0.00000000 Avg_Wait=(nan,nan,0,9223372036854775807,0,-nan) Dispersion=-nan Slowdo
Start: 1553366933Submit: 1553366929waited for 4[GPU=1] Makespan=198 Util=0.42857143 Avg_Wait=(4.00000000,0.00000000,1,4,4,0.0000)
[slurm@1838cc774d52 submitter]$
```

(MC output here is nonsensical since the job we submitted was using GPU nodes)

The wait time seen above we found was actually in *seconds* not minutes as we previously thought. We figured this out by printing the start and submit timestamps and converting them from a UNIX timestamp to a more readable date-time format. This few second long wait checks out with logs from Slurm.

2. By studying the configuration of Jarvis provided by Yao, we were able to calculate the theoretical utility, and we check with the output of RM-Replay which shown to be correct:

Jarvis has 7 node with GPU and 1 without. So total number of nodes is 8. Our testing job use 3 nodes to run a MPI job, so the overall utilization is $3/8 = 0.375$. GPU utilization is also correct: We used 3/7 GPU enabled nodes, which is about 42% utilization.

```
[slurm@1838cc774d52 submitter]$ ./trace_metrics -w /slurm/data/replay.03.23.trace.03.23.testing..0001.27/replay.03.23.trace -r 4
Range: min_start=1553366933 [0,4] start_range=1553366933 end_range=1553381333
all 1 preset 0 (otherp 0)
Start: 1553366933Submit: 1553366929waited for 4[ALL=1] Makespan=198 Util=0.37500000 Avg_Wait=(4.00000000,0.00000000,1,4,4,0.0000)
[MC=0] makespan=9223372036854775807 Util=0.00000000 Avg_Wait=(nan,nan,0,9223372036854775807,0,-nan) Dispersion=-nan Slowdo
Start: 1553366933Submit: 1553366929waited for 4[GPU=1] Makespan=198 Util=0.42857143 Avg_Wait=(4.00000000,0.00000000,1,4,4,0.0000)
[slurm@1838cc774d52 submitter]$
```

We are now continuing to work on feeding Theta logs into RM-Replay. We have nearly completed recreating the slurm_acct_db. Both the jarvis_job_table and resv_table schemas are created.

Also we have created the CSV file mapping the Cobalt logs from Theta over to something we can import into the slurm_acct_db we built.

We are now working on importing the CSV file into the database we built. We've just run into a minor issue with the datatypes not lining up completely, but it just requires a minor fix.