

debugging/ltrace

Ltrace

Basic Usage

- Set up build environment

```
module purge
module add compiler/gnu
```

- Build `stream` benchmark

```
gcc -Ofast -march=native -fopenmp stream.c -o stream -lm
```

- Set up OpenMP environment

```
export OMP_NUM_THREADS=4
export OMP_PROC_BIND=TRUE
export OMP_PLACES=cores
```

- Trace all function calls of benchmark stream

```
ltrace ./stream
```

- Ltrace

- Filter for `alloc` and `free` functions calls within the stream binary
(ignoring these calls within libraries).
- Discard standard output

```
OMP_NUM_THREADS=1 \
ltrace \
    --demangle \
    -e *alloc*@stream+free@stream \
    ./stream \
    >/dev/null

stream->aligned_alloc(64, 0x4c4b400, 0x7f68d8, 1)
stream->aligned_alloc(64, 0x4c4b400, 0x147acb1a8040, 0x147acb1a8000)
stream->aligned_alloc(64, 0x4c4b400, 0x147ac655c040, 0x147ac655c000)
stream->free(0x147acb1a8040)
```

```

stream->free(0x147ac655c040)
stream->free(0x147ac1910040)

-> memory allocation and free for vectors a, b and c

```

- Ltrace

- Filter for `alloc` and `free` functions calls within the stream binary (ignoring these calls within libraries).
- Trace child processes to follow OpenMP Threads.
- Only count matching function calls.

```

OMP_NUM_THREADS=2 \
ltrace \
-f \
--demangle \
-e *alloc*@stream+free@stream \
-c \
./stream >/dev/null

% time      seconds   usecs/call    calls      function
----- -----
 55.62    0.006587     1097        6  free
 37.91    0.004490      748        6 aligned_alloc
   6.47    0.000766      766        1 exit_group
----- -----
100.00    0.011843                  13 total

-> Each OpenMP Thread does its own memory allocation and free

```

Usage scenarios with OpenMPI

- Set up build environment

```

module purge
module add \
    compiler/gnu \
    mpi/openmpi
module add devel/strace

```

- Build `rank_league` benchmark

```
mpicc -O2 -march=native rank_league.c -o rank_league
```

- Ltrace all MPI ranks to individual files (e.g. for comparison)

```

mpirun -np 4 bash -c \
'ltrace -o ltrace.out.${OMPI_COMM_WORLD_RANK} \
./rank_league'
ll -h ltrace.out.*

```

```
-rw-r--r-- 1 bq0742 hk-project-scs 191K May  5 11:05 ltrace.out.0
-rw-r--r-- 1 bq0742 hk-project-scs 188K May  5 11:05 ltrace.out.1
-rw-r--r-- 1 bq0742 hk-project-scs 188K May  5 11:05 ltrace.out.2
-rw-r--r-- 1 bq0742 hk-project-scs 188K May  5 11:05 ltrace.out.3
```

- Ltrace

- Only on first MPI rank (e.g. for data reduction)
- Redirect trace to file

```
mpirun -np 4 bash -c \
'if [[ ${OMPI_COMM_WORLD_RANK} -eq 0 ]]; then
    exec ltrace -o ltrace.out \
        ./rank_league
else
    exec ./rank_league
fi'

ll -h ltrace.out
-rw-r--r-- 1 bq0742 hk-project-scs 191K May  5 11:20 ltrace.out
```

- Ltrace

- Only on first MPI rank (e.g. for data reduction)
- Count calls to MPI functions

```
mpirun -np 4 bash -c \
'if [[ ${OMPI_COMM_WORLD_RANK} -eq 0 ]]; then
    exec ltrace -c -e *MPI* \
        ./rank_league
else
    exec ./rank_league
fi'

% time      seconds   usecs/call      calls      function
----- -----
 32.58     1.344215    1344215          1 MPI_Finalize
 28.26     1.165933    1165933          1 MPI_Init
 18.16     0.749022      936        800 MPI_Isend
 17.42     0.718733      898        800 MPI_Irecv
  2.75     0.113337    14167          8 MPI_Waitall
  0.38     0.015681    15681          1 exit_group
  0.15     0.006058      757          8 MPI_Sendrecv
  0.13     0.005490      686          8 MPI_Wtime
  0.09     0.003766      941          4 MPI_Barrier
  0.04     0.001478      492          3 MPI_Recv
  0.02     0.000697      697          1 MPI_Comm_size
  0.02     0.000671      671          1 MPI_Get_processor_name
  0.01     0.000581      581          1 MPI_Comm_rank
```