

Practical Performance Analysis with Score-P/Vampir

Perform all the following tasks in your study group.

Contents

Task 1: Build and trace a small MPI C program (15 min)	1
Task 2: Analyze the output traces of Task 1 in Vampir (20 min)	3
Task 3: Repeat the steps for an optimized code (10 min)	4
Task 4: Build and trace a small Python MPI program (20 min)	5
Task 5: Analyze the output traces of Task 4 in Vampir (15 min)	6

Task 1: Build and trace a small MPI C program (15 min)

Build, trace, and inspect a small MPI program with Score-P.

Steps

1. Connect with X11 forwarding if you want to use graphical tools remotely:

```
1 ssh -X user@glogin-p2.hpc.gwdg.de
```

-X enables X11 forwarding, so graphical applications running on the remote system can be displayed locally.

2. Load the required modules:

```
1 module spider scorep/8.4  
2 module load gcc/14.2.0 openmpi/4.1.7 scorep/8.4
```

3. Save the following MPI C code as `code.c` (You can find the code in the website page):

```
1 #include <mpi.h>  
2 #include <stdio.h>  
3 #include <stdlib.h>  
4 #include <unistd.h>  
5  
6 static void heavy_work(long long n)  
7 {  
8     volatile double x = 0.0;  
9     for (long long i = 0; i < n; i++) {  
10         x += (i % 7) * 0.000001;  
11     }
```

```

12 }
13
14 int main(int argc, char** argv)
15 {
16     MPI_Init(&argc, &argv);
17
18     int rank, size;
19     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
20     MPI_Comm_size(MPI_COMM_WORLD, &size);
21
22     long long scale = 1;
23     if (argc > 1) {
24         scale = atoll(argv[1]);
25         if (scale < 1) {
26             if (rank == 0) {
27                 fprintf(stderr, Scale factor must be >= 1\n);
28             }
29             MPI_Finalize();
30             return 1;
31         }
32     }
33
34     long long base_work_slow = 400000000LL;
35     long long base_work_fast = 100000000LL;
36     long long work = (rank == 0) ? base_work_slow * scale : base_work_fast * scale;
37
38     heavy_work(work);
39
40     /* Everybody waits here for the slow rank */
41     MPI_Barrier(MPI_COMM_WORLD);
42
43     int value = rank;
44
45     if (rank == 0) {
46         int* recvbuf = malloc(size * sizeof(int));
47         if (!recvbuf) {
48             perror(malloc);
49             MPI_Abort(MPI_COMM_WORLD, 1);
50         }
51
52         recvbuf[0] = value;
53
54         for (int src = 1; src < size; src++) {
55             MPI_Recv(&recvbuf[src], 1, MPI_INT, src, 0, MPI_COMM_WORLD,
56                     MPI_STATUS_IGNORE);
57         }
58
59         FILE* f = fopen(mpi_bad_output.txt, w);
60         if (!f) {
61             perror(fopen);
62             free(recvbuf);
63             MPI_Finalize();
64             return 1;
65         }
66
67         long long repeats = 2000LL * scale;
68         for (long long repeat = 0; repeat < repeats; repeat++) {
69             for (int i = 0; i < size; i++) {
70                 fprintf(f, repeat=%lld rank=%d value=%d\n, repeat, i, recvbuf[i]);
71                 fflush(f);
72             }
73         }
74
75         fclose(f);
76         free(recvbuf);
77     } else {

```

```
78     MPI_Send(&value, 1, MPI_INT, 0, 0, MPI_COMM_WORLD);
79 }
80
81 MPI_Barrier(MPI_COMM_WORLD);
82
83 MPI_Finalize();
84 return 0;
85 }
```

4. Instrument and compile the code with Score-P:

```
1 scorep --mpp=mpi --thread=none --nomemory mpicc -O2 -o scorep_mpi code.c
```

5. Run the program (you can change the scale factor 50 to run the code for longer):

```
1 export SCOREP_ENABLE_TRACING=true
2 export SCOREP_TOTAL_MEMORY=4GB
3 srun -p <partition> -N 1 -n 4 ./scorep_mpi 50
```

6. You will see a Score-P output directory in the execution location. Explore its contents:

```
1 scorep-20260409_XXXXXXXXXXXXXXXXXXXXXXXX
```

7. Optionally inspect the generated files, for example:

```
1 ls scorep-*/
2 scorep-score scorep-*/profile.cubex
```

Task 2: Analyze the output traces of Task 1 in Vampir (20 min)

Open the generated trace in Vampir and identify the main performance issues (It may take a few seconds, please be patient).

Steps

1. Load Vampir (You must have connected with -X option):

```
1 module load vampir
```

2. Open the trace:

```
1 vampir /path/to/output/traces.otf2
```

3. In Vampir, inspect:

- load imbalance across ranks
- waiting time at the barrier
- communication concentration at rank 0

4. As discussed during the session, you can also install a limited Vampir version locally, copy the full output directory to your local machine, and analyze it there.

Task 3: Repeat the steps for an optimized code (10 min)

Save the following code as `fixed_code.c` (You can find the code in the website page), do the exact same steps in task 1, for the following optimized code, and compare the results

```
1 #include <mpi.h>
2 #include <stdio.h>
3 #include <stdlib.h>
4
5 static void heavy_work(long long n)
6 {
7     volatile double x = 0.0;
8     for (long long i = 0; i < n; i++) {
9         x += (i % 7) * 0.000001;
10    }
11 }
12
13 int main(int argc, char** argv)
14 {
15     MPI_Init(&argc, &argv);
16
17     int rank, size;
18     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
19     MPI_Comm_size(MPI_COMM_WORLD, &size);
20
21     /* Optional scale factor */
22     long long scale = 1;
23     if (argc > 1) {
24         scale = atoll(argv[1]);
25         if (scale < 1) {
26             if (rank == 0) {
27                 fprintf(stderr, Scale factor must be >= 1\n);
28             }
29             MPI_Finalize();
30             return 1;
31         }
32     }
33
34     /* Fixed version:
35        all ranks do the same amount of work */
36     long long base_work = 100000000LL;
37     long long work = base_work * scale;
38     heavy_work(work);
39
40     /* Improved communication:
41        use collective communication instead of rank 0 receiving one by one */
42     int value = rank;
43     int* all_values = NULL;
44
45     if (rank == 0) {
46         all_values = malloc(size * sizeof(int));
47         if (!all_values) {
48             perror(malloc);
49             MPI_Abort(MPI_COMM_WORLD, 1);
50         }
51     }
52
53     MPI_Gather(&value, 1, MPI_INT,
54              all_values, 1, MPI_INT,
55              0, MPI_COMM_WORLD);
56
57     /* Improved I/O:
58        rank 0 writes once in larger buffered form */
59     if (rank == 0) {
60         FILE* f = fopen(mpi_fixed_output.txt, w);
61         if (!f) {
```

```

62     perror(fopen);
63     free(all_values);
64     MPI_Finalize();
65     return 1;
66 }
67
68 long long repeats = 2000LL * scale;
69 for (long long repeat = 0; repeat < repeats; repeat++) {
70     for (int i = 0; i < size; i++) {
71         fprintf(f, repeat=%lld rank=%d value=%d\n,
72             repeat, i, all_values[i]);
73     }
74 }
75
76 fclose(f);
77 free(all_values);
78 }
79
80 MPI_Finalize();
81 return 0;
82 }

```

Task 4: Build and trace a small Python MPI program (20 min)

Create a Python environment, install the required packages, and trace a small MPI Python program with Score-P.

Steps

1. Load Miniforge and create a local Python environment:

```

1 module load miniforge3
2 conda create --prefix ./myenv python=3.12
3 source activate ./myenv

```

2. Load the needed modules:

```

1 module load gcc/14.2.0 openmpi/4.1.7 scorep/8.4
2 module load openmpi

```

3. Install mpi4py:

```

1 pip install mpi4py
2 pip install numpy

```

4. Check that mpi4py is installed correctly:

```

1 python -c import mpi4py

```

5. Clone and install the Score-P Python bindings:

```

1 git clone https://github.com/score-p/scorep_binding_python
2 cd scorep_binding_python
3 pip install .
4 cd ..

```

6. Save the following code as `script.py`

```
1 from mpi4py import MPI
2 import numpy as np
3 amode = MPI.MODE_WRONLY|MPI.MODE_CREATE
4 comm = MPI.COMM_WORLD
5 fh = MPI.File.Open(comm, './datafile.contig', amode)
6 buffer = np.empty(10, dtype=np.int32)
7 buffer[:] = comm.Get_rank()
8 offset = comm.Get_rank()*buffer.nbytes
9 fh.Write_at_all(offset, buffer)
10 fh.Close()
```

7. Enable tracing and set Score-P memory:

```
1 export SCOREP_ENABLE_TRACING=true
2 export SCOREP_TOTAL_MEMORY=1837MB
```

8. Run the Python program with Score-P:

```
1 srun -p medium -n 3 python -m scorep --io=runtime:posix --mpp=mpi script.py
```

9. Inspect the generated Score-P output directory:

```
1 scorep-20260409_xxxx_XXXXXXXXXXXXXXXXXXXX
```

Task 5: Analyze the output traces of Task 4 in Vampir (15 min)

Steps

1. Open the generated trace in Vampir, and analyze it (follow the steps in task 2).