

CS455 Help Sheet

OpenMP Help

OpenMP is an API for parallel programming in shared memory systems.

Common Directives

- `#pragma omp parallel`: Defines a parallel region.
- `#pragma omp for`: Divides loop iterations among threads.
- `#pragma omp sections`: Divides tasks into independent sections.
- `#pragma omp single`: Ensures a block is executed by one thread.
- `#pragma omp critical`: Ensures mutual exclusion.
- `#pragma omp barrier`: Synchronizes threads.
- `#pragma omp reduction(op: list)`: Combines results from threads.

Environment Variables

- `OMP_NUM_THREADS`: Set the number of threads.
- `OMP_SCHEDULE`: Configure loop scheduling (static, dynamic).

Performance Tips

- Use `OMP_SCHEDULE` for load balancing.
- Avoid false sharing by aligning data structures.
- Minimize critical sections to reduce contention.
- Use `omp_get_wtime` for timing in OpenMP.

Examples

Parallel Loop:

```
1 #pragma omp parallel for
2 for (int i = 0; i < n; i++) {
3     a[i] = b[i] + c[i];
4 }
```

Reduction Example:

```
1 int sum = 0;
2 #pragma omp parallel for reduction(+:sum)
3 for (int i = 0; i < n; i++) {
4     sum += a[i];
5 }
```

Critical Section:

```
1 #pragma omp critical
2 {
3     shared_var += local_val;
4 }
```

MPI Help

MPI (Message Passing Interface) is used for distributed-memory parallel programming.

Common Functions

- `MPI_Init`: Initialize the MPI environment.
- `MPI_Comm_rank`: Get the rank of the process.
- `MPI_Comm_size`: Get the total number of processes.
- `MPI_Send / MPI_Recv`: Send/receive messages.
- `MPI_Barrier`: Synchronize processes.
- `MPI_Reduce`: Perform a reduction operation across processes.
- `MPI_Bcast`: Broadcast a message from one process to all other processes.
- `MPI_Finalize`: Clean up and exit the MPI environment.

Environment Variables

- `mpirun -np <num_processes> <program>`: Run a program with the specified number of processes.
- `MPICH_DEVICE`: Define the network device for communication.

Performance Tips

- Minimize communication between processes.
- Use non-blocking communication (`MPI_Isend`, `MPI_Irecv`) for overlap.

- Decompose work to minimize idle time among processes.
- Use `MPI_Wtime` for measuring execution time in MPI programs.

Examples

Basic MPI Program:

```
1 #include <mpi.h>
2 int main(int argc, char** argv) {
3     MPI_Init(&argc, &argv);
4     int rank, size;
5     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
6     MPI_Comm_size(MPI_COMM_WORLD, &size);
7     printf("Hello from rank %d of %d\n", rank, size);
8     MPI_Finalize();
9 }
10 }
```

Send, Receive, and Barrier Example:

```
1 #include <mpi.h>
2 #include <stdio.h>
3
4 int main(int argc, char** argv) {
5     MPI_Init(&argc, &argv);
6     int rank, size, data;
7     MPI_Comm_rank(MPI_COMM_WORLD, &rank);
8     MPI_Comm_size(MPI_COMM_WORLD, &size);
9
10     if (rank == 0) { // Sender
11         data = 42; // Example data
12         MPI_Send(&data, 1, MPI_INT, 1, 0,
13                  MPI_COMM_WORLD);
14         printf("Process 0 sent data %d to Process
15               1\n", data);
16     }
17     else if (rank == 1) { // Receiver
18         MPI_Recv(&data, 1, MPI_INT, 0, 0,
19                  MPI_COMM_WORLD, MPI_STATUS_IGNORE);
20         printf("Process 1 received data %d from
21               Process 0\n", data);
22     }
23     MPI_Barrier(MPI_COMM_WORLD); // Synchronize all
24     processes
25     printf("Process %d reached the barrier\n", rank);
26
27     MPI_Finalize();
28 }
29 }
```

C++ Help

C++ is used in HPC for its efficiency and control over system resources.

Common Data Structures and Functions

- std::vector: Dynamic arrays.
- std::unordered_map: Fast key-value lookups.
- std::unordered_set: Unique elements with fast lookups.
- std::sort: Sorts elements.
- std::move: Transfers data without copying.
- std::unique_ptr: Exclusive ownership.
- std::shared_ptr: Shared ownership.
- std::chrono::high_resolution_clock: Timing function.

Common Flags

- -O0: No optimization (default flag).
- -O1: Basic optimizations.
- -O2: Balanced optimizations.
- -O3: Aggressive optimization (may increase compile time and cause unexpected behavior).
- -Ofast: More aggressive optimization (may cause unexpected behavior).
- -g: Generates debug information.
- -Wall -Wextra: Enable extra warnings.

Examples

Chrono Timing Example:

```
1 #include <iostream>
2 #include <chrono>
3
4 int main() {
5     auto start =
6         std::chrono::high_resolution_clock::now();
7     // Code to time
8     auto end =
9         std::chrono::high_resolution_clock::now();
10    std::chrono::duration<double> elapsed = end -
11        start;
12    std::cout << "Elapsed time: " << elapsed.count()
13        << " seconds\n";
14 }
```

Lambda Example:

```
1 auto add = [](int a, int b) -> int {
2     return a + b;
3 };
```

Linux Help

Linux is widely used in HPC.

Basic Commands

- ls: List directory contents.
- pwd: Present working directory.
- cd <dir>: Change directory.
- cp <src> <dest>: Copy files.
- mv <src> <dest>: Move/rename files.
- rm <file>: Delete files.
- ssh <user>@<host>: Connect to a remote machine.
- scp <src> <dest>: Copy files between local and remote machine.

Git Help

Git is a version control system.

Basic Commands

- git init: Initialize a repository.
- git clone <url>: Clone a repository.
- git status: Display files status.
- git checkout <branch>: Switch branch.
- git add <file>: Stage changes.
- git commit -m "message": Commit changes.
- git push: Push changes to a remote repository.
- git pull: Pull changes from a remote repository.

PBS Pro Help

PBS Pro is a job scheduler commonly used on HPC systems.

Common Commands

- qsub <script>: Submit a job script.
- qstat: Check the status of submitted jobs.
- qdel <job_id>: Cancel a running or queued job.
- pbsnodes -a: Display information about compute nodes.
- tracejob <job_id>: View job history and diagnostic information.

Typical PBS Pro Job Script

```
1 #!/bin/bash
2 #PBS -N MyJob
3 #PBS -l nodes=1:ppn=4
4 #PBS -l walltime=01:00:00
5 #PBS -q short
6 #PBS -j oe
7
8 # Change to working directory
9 cd ${PBS_O_WORKDIR}
10 ./my_program
```

Slurm Help

Slurm (Simple Linux Utility for Resource Management) is another popular job scheduler.

Common Commands

- sbatch <script>: Submit a Slurm job script.
- squeue: View queued and running jobs.
- scancel <job_id>: Cancel a running or pending job.
- sinfo: Display partition and node information.
- srun <executable>: Launch a parallel job step (often used inside scripts).

Typical Slurm Job Script

```
1 #!/bin/bash
2 #SBATCH --job-name=MySlurmJob
3 #SBATCH --nodes=1
4 #SBATCH --ntasks-per-node=4
5 #SBATCH --time=01:00:00
6 #SBATCH --partition=short
7
8 #SBATCH --output=%x.%j.out
9 #SBATCH --error=%x.%j.err
10
11 # Run the program
12 ./my_program
```