# Program #4 (Cpt S 223)

Fall 2010

Due: November 15, 2010, 5pm on Angel

## General Guidelines

- All source code must be in C++. You can use Windows or Unix environments at your discretion.
- Each program project should be accompanied by a COVER SHEET (see details below). Assignments without cover page will NOT be graded.
- **Team work is allowed and recommended for this project** (although not mandatory). Each team can comprise of at most two members. Grading will not differentiate between 1-member and 2-member teams. It is an option that is intended to encourage joint discussion and implementation.
- The final material for submission should be entirely written by you/your team. If you decide to consult with others outside your team, or refer materials online, you MUST give due credits to these sources (people, books, webpages, etc.) by listing them on the cover sheet mentioned below. Note that no points will be deducted for referencing these sources. However, your discussion/consultation should be limited to the initial design level. Sharing or even showing your source code/assignment verbiage to anyone else in the class, or direct reproduction of source code/verbiage from online resources, will all be considered plagiarism, and therefore will be awarded ZERO points and subject to the WSU Academic Dishonesty policy. (Reproducing from the Weiss textbook is an exception to this rule, and such reproduction is encouraged wherever possible.)
- Grading will be based on correctness, coding style, implementation efficiency, exception handling, source code documentation, and the written report. For 2-member teams, the same grade will be assigned to the whole team.
- Angel Submission: Make just ONE SUBMISSION PER TEAM. The submission should include a cover sheet as usual and contain the other documents listed in the assignment. The submission should have one zip file (or tar file) and it should have the last names of the individual(s) in the team. Submit as an attachment on Angel's email and sent to: "All course faculty". This will send out an email to both the instructor and the TA. Submissions are due by 5pm.

PS: Submissions through any other means or from general email addresses will be discarded and will NOT be graded. So please follow the above instructions carefully.

Late submission policy: A late penalty of 10% will be assessed for late submissions within the next 24-hour. Note: Submissions will be accepted only through the Angel webmail. Any other submission outside the Angel portal will be discarded and not graded. (In the unlikely event of an Angel server outage on the day of submission, assisgnments can be emailed to the instructor's EECS email account.)

# 1 Problem Introduction

You will be primarily using priority queues for this programming assignment. Among priority queues, you are free to pick and choose the most appropriate/efficient heap implementation. Your design can also contain any other data structure we have discussed so far in class, as long as the use of such a data structure could be justified from a design perspective. In fact, go for a simpler data structure wherever possible. An unnecessarily overcomplicated implementation (example: using a heap where actually a simple array would suffice) will lead to reduction in points. The idea is to let you make the design choice, and identify the most suitable data structure for each system component that will effectively address performance and functional tradeoffs.

Use of any appropriate STL containers is highly encouraged whenever possible (*priority_queue, set, map, list, vector*, etc.). You are also free to re-use/reproduce any source code that is provided in the Weiss textbook. Most of its source code are available from the following website: `http://www.cs.fiu.edu/~weiss/dsaa_c++/Code/`.

# 2 Problem Statement

Your organization has purchased a new parallel computer (or "cluster") which has $p$ processors. Your task is to design and implement a simple *shortest-job-first scheduler* that allows multiple users to access the cluster at the same time, as per the specifications in Section 3. Figure 1 is a conceptual illustration of the different components in the system, how they interact, and the necessary data containers you will have to implement. Refer to this figure while reading the specification in Section 3.

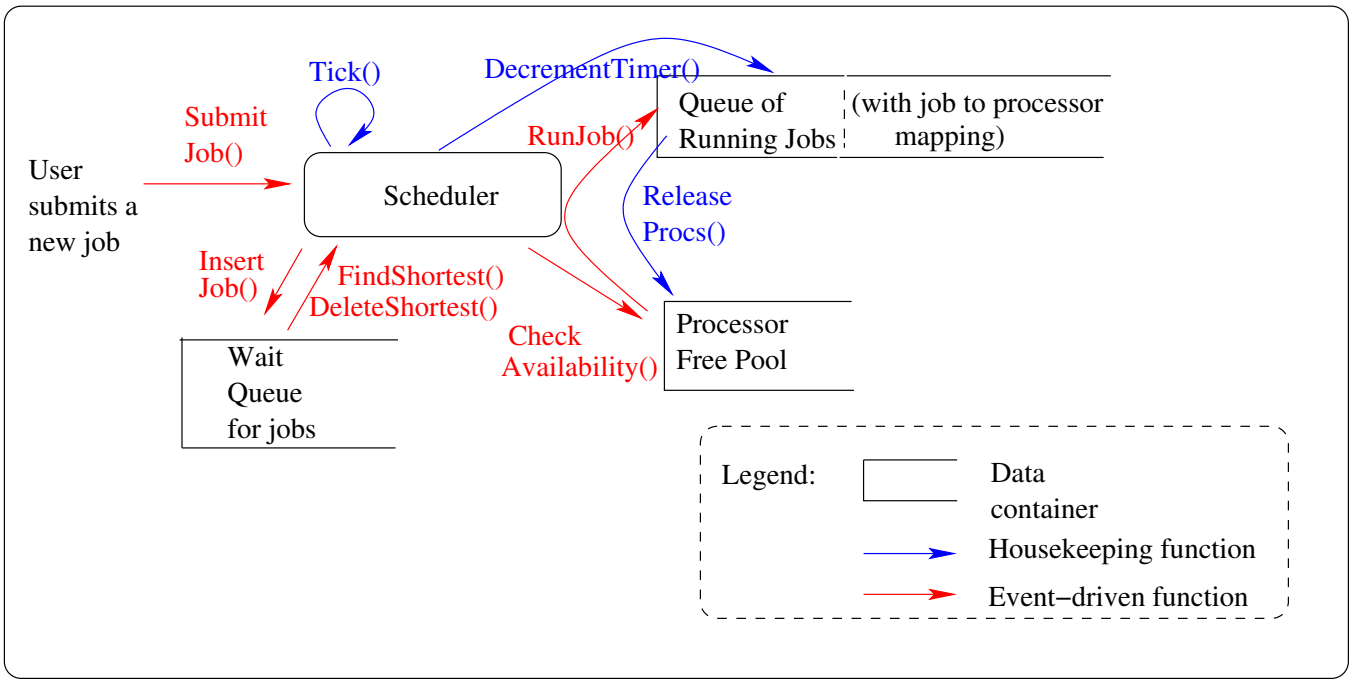


Figure 1: Organization of the scheduler system for a parallel cluster that has a total of $p$ processors.

**Algorithm 1** *Pseudocode for the scheduler' Tick procedure*

**Tick()**

0. Prompt user for submitting new jobs, one at a time.
1. Insert all user-submitted new jobs into the wait queue;
2. Decrement timer for all running jobs in the running queue;
3. Release all the processors corresponding to the completed/expired jobs from the running queue to the free pool. A job is said to have "expired" if its timer becomes zero during this tick. Also, any expired job should be automatically removed from the running job queue;
4. Find the next shortest job ( $< J_i, p_i, t_i >$) from the wait queue and check if sufficient number of processors are available in the free pool to run it. If so:
   - Remove job $J_i$ from the wait queue;
   - Insert job $J_i$ into the running queue and assign $p_i$ processors (from free pool) to that job;

Figure 2: Sequence of steps to be carried out during each Tick() of the scheduler.

## 3 Problem Specification

The scheduler performs all required functions at regular timesteps called "ticks". The tick in other words, is the pulse of the scheduler. During each tick, various events could happen and the scheduler's responsibility is to address all these events before moving on to the next tick. Because of this, each tick could take variable amount of time to complete, and for design simplicity you can assume that there is sufficient real time within each tick to complete all required tasks. The various events and the respective action that is required to be carried out by the scheduler are outlined below. (Refer Figures 1 & 2 as you read this.)

1. ***Event:*** The user has input one or more jobs at the start of any tick. Each job arrives to the scheduler in the following format: <job_description, n_procs, n_ticks>.

   ***Action:*** For each job submission, the system assigns the job a new, unused integer id (referred to as "job_id") and then calls the **InsertJob()**[1] function with the following parameters:

   <job_id, job_description, n_procs, n_ticks>, where:

   (i) "job_description" contains details on what command to execute to run the job (i.e., "<program name> <its arguments>"). It is a string data type;

   (ii) "n_procs" is the number of processors that this job needs to run on;

   (iii) "n_ticks" is the number of ticks that this job is going to need on each of the n_proc processors. Note that this number is the equivalent of the job's estimated running time, from start to finish, as the job will be launched on all n_procs processors at the same time to run in parallel for n_ticks number of ticks.

   The **InsertJob()** function first checks if ($0 <$ n_procs $\leq p$) and (n_ticks$> 0$). If so, it inserts the job into a " wait queue". Otherwise, a job submission error is raised with an appropriate message.

2. ***Event:*** The job wait queue is *not* empty.

   ***Action:*** Among all the jobs in the wait queue, find the job that is expected to take the least time to complete (**FindShortest()**) – i.e., a job with the minimum value for n_ticks. If there

[1] Note: The notation **InsertJob()** does *not* necessarily mean there are no arguments to this function. You are free to decide what arguments each function signature should have. I have purposefully left that choice to you as part of design. This applies to all the functions.

are more than one job to choose from this category, then pick any one of them. Let $J_i$ denote the selected job, and $p_i$ denote its n_procs. Now, check if there are at least $p_i$ processors currently available in the processor "free pool" (**CheckAvailability()**). If not, then leave the wait queue intact and return. Otherwise:

(a) remove $J_i$ from the wait queue (**DeleteShortest()**), and

(b) remove $p_i$ arbitrary processors from the free pool and assign them to that job (**RunJob()**), so that they can start executing the job from the next tick[2]. When a job goes to the running state, along with it a countdown "timer" is started; initialize this to n_ticks.

3. ***Housekeeping Event:*** The running job queue is *not* empty.
   ***Action:*** The job timer for *ALL* running jobs is decremented by one (**DecrementTimer()**).

4. ***Housekeeping Event:*** A job that was in the running state completed during this tick — i.e., its timer has reached a value of zero after decrement.
   ***Action:*** The scheduler should release all the processors attached to this job, back into the free pool so that they become available for other jobs during later ticks (**ReleaseProcs()**).

When the algorithm starts, the system is initialized with $p$ processors in the free pool. Figure 2 shows the sequence of steps that the scheduler should perform during each tick as a pseudocode.

# 4 API and Testing

Follow the below instructions to test your code.

- Write a separate main.cpp program which runs an indefinite number of iterations. Each iteration corresponds to one tick. In each iteration, the current tick number should be displayed. Then the code should call the Tick() function. Define the Tick() function as a member function of your Scheduler class.

- At the start of each tick, the program should prompt the user to input an arbitrary number of jobs through the standard input, one job at a time. Basically prompt the user for three entries: (i) Job descriptor (a string), (ii) n_procs (an integer), and (iii) n_ticks (an integer).

[2]Remember, your code does not have to concern itself about running the actual job. You can assume that will be taken care by the individual processors which were assigned to the job.

After each entry the user will press enter to signal the end of that entry. If the user has no more new jobs to submit, then the user will signal that by entering the string "null" for job description.

- Then the tick procedure should carry out the steps outlined in the Tick() pseudocode in Figure 2.
- At the end of the Tick() function, display back a system-defined Job_id number unique to each job submitted during this tick or an error message upon failure to insert.
- If there was any job that completed during this tick, then *print* its job_id along with information on how many processors got released. Also, at the end of the tick, print separately i) the contents of the wait queue (just the job ids), ii) list of all currently running jobs (job ids, present timer value, number of processors being used), and iii) the number of processors in the free pool.
- The test program should terminate if the user inputs "exit" for the job descriptor in the standard input.

## 5 Design Document - (1 page limit)

Along with your source code, also submit a one-page design document (as Word Doc or PDF - no other formats allowed).

In the first part of this document, provide a figure that best describes your design for the scheduler system. It should basically be a modified version of Figure 1, in which you should state what data structures/STL containers you used to implement the three data containers (wait queue, running job queue, free pool) .

In the second part, make a table to simply state the worst-case run-time complexities in your implementation for each of the event functions (InsertJob,FindShortest, etc.).

In the third part, add a short paragraph about what you think are the main shortcomings/bottlenecks of this shortest-job-first strategy, from both a performance and functionality point of view. Support your claims with brief but compelling rationale.

## 6 FINAL CHECKLIST

Make only one submission per team. In the zip file's name include the last names of the individuals in the team. For 2-member teams, any one of the two individuals can submit and copy (CC) the other. The zip file archive should contain:

__ Cover sheet (just one per team)

__ Source code folder with main.cpp and all project files

__ 1-page design document