

## CSC 343 Operating Systems, Spring 2021

**Dr. Dale E. Parson, Assignment 4, modeling swapping using a FIFO queue and a min queue ordered on process logical memory size.**

This assignment is due via **make turnitin** from the swap4prj2021 directory by **11:59 PM on Friday April 30**.

I am handing out the STM code for a simulation that uses swapping; it places processes into a FIFO queue when they must wait to map their logical memory to physical memory via a relocation register. You must write an alternative implementation that uses a min-priority queue, with ordering based on logical memory size, to store the waiting processes. There is very little code to change for this project. The emphasis is more on analysis. You must read and understand the code to understand your assignment, and you must update a **README.txt** file that I have started.

Perform the following steps to get my handout. You will need to test on machine mcgonagall as previously explained (**ssh mcgonagall** from acad).

```
cd $HOME          # or start out in your login directory
cd ./OpSys
cp ~parson/OpSys/swap4prj2021.problem.zip swap4prj2021.problem.zip
unzip swap4prj2021.problem.zip
cd ./swap4prj2021
make testfifo     # This tests the handout code. It should work.
make testmin      # This tests your code. It will not work until you make your changes.
```

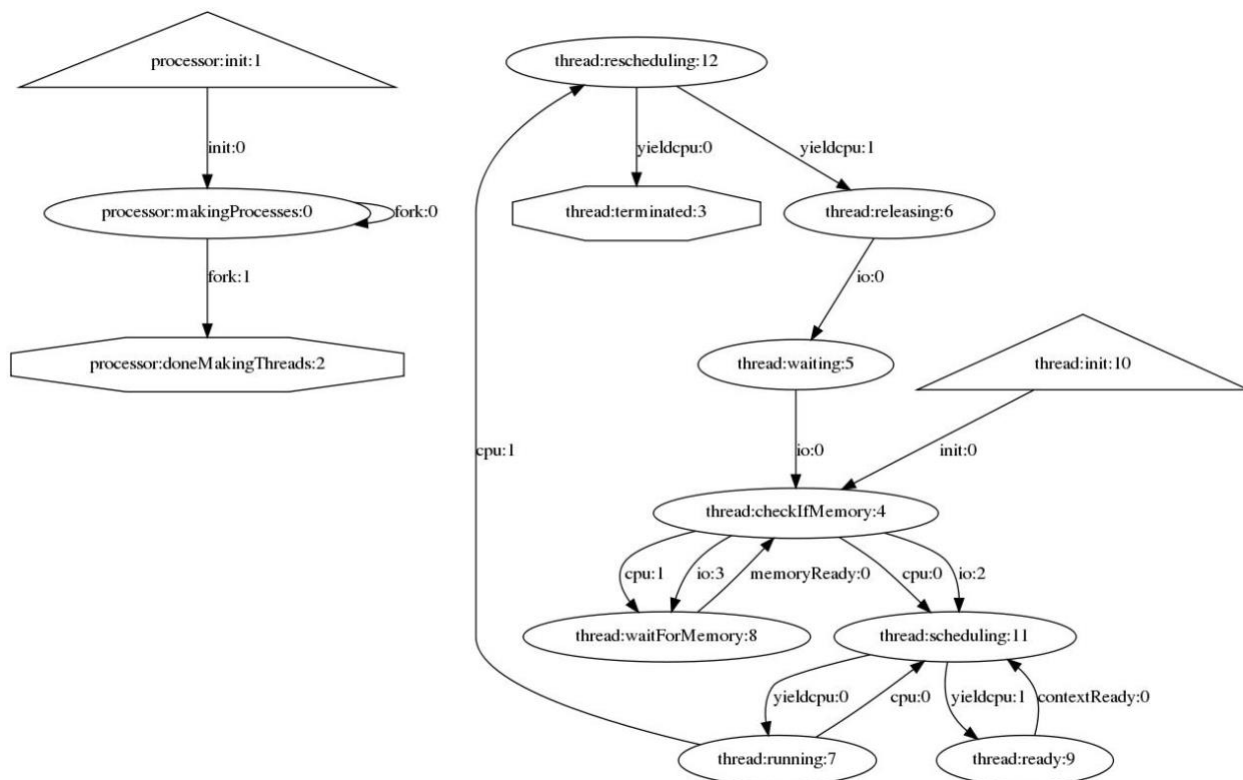
### 1. Understand the workings of the handout code (See STM.doc.txt for library documentation.)

Page 2 shows the JPEG graph for all of the STMs in this assignment. Your code changes must not change this graph topology. Here is the list of STM files handed out with the project, tested by **make testfifo**. All STMS in this assignment used round robin CPU scheduling.

```
rr_ff_swap.stm # Locate memory for a process using first fit.
                # Activate (release from memory queue) only one process waiting for memory at a time.
rr_bf_swap.stm # Locate memory for a process using best fit.
allrr_ff_swap.stm # Locate memory for a process using first fit.
                  # Activate all processes waiting for memory at a time.
allrr_bf_swap.stm # Locate memory for a process using best fit.
                  # Activate all processes waiting for memory at a time.
```

You will perform the following steps, and then edit the new “min” files with your code changes. **DO NOT MODIFY THE HANDOUT .stm FILES LISTED ABOVE.**

```
cp rr_ff_swap.stm rr_ff_min.stm
cp rr_bf_swap.stm rr_bf_min.stm
cp allrr_ff_swap.stm allrr_ff_min.stm
cp allrr_bf_swap.stm allrr_bf_min.stm
```



**State machine graph for all STM files in assignment 4.**

Here is a summary of the main transitions in the diagram. We will go over the code in class.

1. The **processor state machine** fork()s 24 single-threaded processes after initializing the following data fields and some unrelated fields.

```
# Two gig of available memory, initially all free, starts at phyaddr 0.
processor.physicalMemory = 2147483648 ;
# Each entry in this table is a (LIMIT, BASE) pair, with values giving
# (SIZE, START_ADDRESS) of a region in physical memory. As memory gets
# allocated and freed by processes, this initial region gets chopped
# up into multiple smaller regions and the merged where possible
# when freeing. See macros and lambda functions in the thread state
# machine for the code the manages the processor.freeregions and
# pcb.relocationRegister.
processor.freeregions = [(processor.physicalMemory, 0)];
# The waitMemQ is the Queue where processes wait until they get memory.
processor.waitMemQ = Queue(ispriority=False);
```

2. The **init -> checkIfMemory** transition initializes some memory management variables and performs the first memory allocation request for physical memory.

```
swapdevice = len(processor.fastio)-1 ;
# The range of process sizes drives contention for memory.
# Each process averages  $(1/16 + 1/2) / 2 = .28125$  of physical memory,
```

```

# so we'd expect 3 to 4 processes to fit into memory at one time,
# and never fewer than 2 will fit at one time. See processesToDo
# in the processor STM above.
pcb.logicalMemorySize = sample(int(processor.physicalMemory/16),
    int(processor.physicalMemory/2), 'uniform');
# pcb.relocationRegister will get a piece of processor.freeregions.
# Its value when memory is allocated is (SIZE, BASE), else None.
pcb.relocationRegister = None ;

```

3. All transitions into state **checkIfMemory** invoke macro **getFirstFit** (or **getBestFit** in best fit STMs) to attempt to allocate memory. When this macro succeeds, it sets variable **findcost** to the cost in ticks for locating the memory or trying to find memory, and it sets **pcb.relocationRegister** to the pair (SIZE, LOC), where SIZE is the number of locations in the process' logical memory, and LOC is the location in physical memory, i.e., the BASE ADDRESS. The **pcb.relocationRegister** is None (Python's null value) when insufficient memory is available; **findcost** is still valid in that case.
4. The transitions between **checkIfMemory** and **waitForMemory** illustrate the fact that, when there is insufficient memory to satisfy the memory request, transitions into **waitForMemory** enqueue the thread into the **processor.waitMemQ**, and transitions back to **checkIfMemory** make another attempt to allocate physical memory after being dequeued and awoken by a **memoryReady** event sent by another thread releasing its physical memory.
5. After acquiring memory, a thread advancing to state **scheduling** performs the round robin CPU scheduling algorithm discussed a few weeks ago.
6. The transitions into states **terminated** and **releasing** release their physical memory while terminating or preparing to perform application I/O, waking up one or more other processes waiting for memory. The *all* versions wake up all process threads waiting in **processor.waitMemQ** by sending a **memoryReady** event, while the non-all versions wake up only the thread at the front of **processor.waitMemQ**.

## YOUR JOB

Code changes are worth 50% of the assignment. README.txt is the other 50%.

Your job after performing the cp copying steps at the bottom of page, thereby creating these four files, is as follows.

```

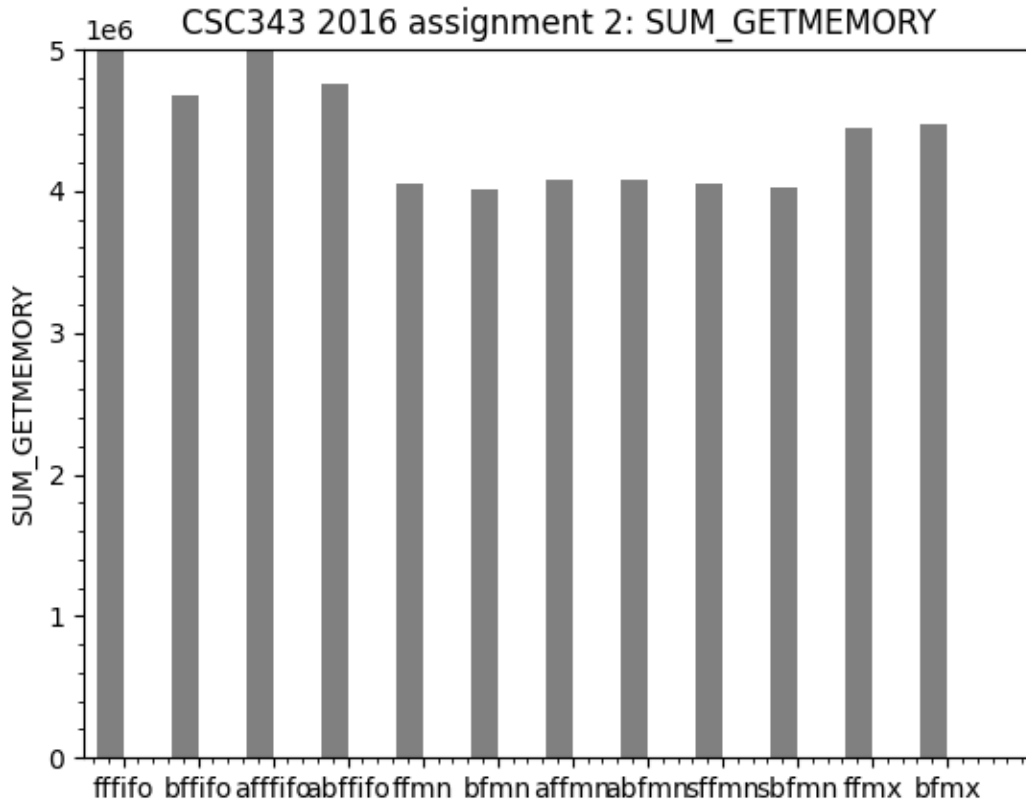
rr_ff_min.stm
bf_min.stm
allrr_ff_min.stm
allrr_bf_min.stm

```

- A. Change **processor.waitMemQ** from a FIFO queue to a priority queue. Python priority queues are min-queues, meaning that the minimum priority value goes to the front of the queue upon enqueueing. To see a min-queue, see field **processor.readyq** in **your last srtf CPU scheduler**. (Don't change **processor.readyq** to a priority queue.)
- B. Change all **processor.waitMemQ.enqueue** function calls to use the process logical memory size as the priority value.
- C. At this point, **make testmin** should work, as should **make clean test**. make clean runs make testmin testfifo.
- D. Update all of the files you edit correctly, adding your name, and describing the algorithm properly in

the comments at the top of the file, including the file name.

Here is a bar graph that shows SUM\_GETMEMORY, which is the combined sum of time spent in states xxx and xxx for all 24 processes, with the “a” prefix signifying the “all” STMs, “fifo” meaning the FIFO xxx and “mn” the min-queue. I also coded a max queue on logical memory size; you are not required to do so. The “ff” prefix is first fit, and “bf” is best fit.



[http://acad.kutztown.edu/~parson/SUM\\_GETMEMORY.png](http://acad.kutztown.edu/~parson/SUM_GETMEMORY.png)

See the actual, sorted values in distributed file **README.txt**.

**Answer the 5 questions found at the bottom of README.txt below each question as before.**

The following instructions for tracing memory allocation & freeing should be helpful in answering the 5 questions.

### **VIEWING A TRACE OF FIRST FIT AND BEST FIT MEMORY ALLOCATIONS AND DEALLOCATIONS**

Invocations of the macro **takeMemory** for transferring memory from **processor.freeregions** to **pcb.relocationRegister**, and **putBackRegion** for transferring memory from **pcb.relocationRegister** to **processor.freeregions**, are already working in the handout code. Each prints a trace of **pcb.relocationRegister** and then **processor.freeregions** at both the start and end of each of these macros. Here are the **msg** invocations within these macros.

```
msg("DEBUG_TKM1|" + str(pcb.relocationRegister) + "|" + str(processor.freeregions) + "|");
# Code for takeMemory appears between the two msg statements.
msg("DEBUG_TKM2|" + str(pcb.relocationRegister) + "|" + str(processor.freeregions) + "|");

msg("DEBUG_PBR1|" + str(pcb.relocationRegister) + "|" + str(processor.freeregions) + "|");
# Code for putBackRegion appears between the two msg statements.
msg("DEBUG_PBR2|" + str(pcb.relocationRegister) + "|" + str(processor.freeregions) + "|");
```

After a successful test run, you can inspect the full trace of memory allocation and deallocation / compaction in the appropriate simulation log file. For example, for handout rr\_ff\_swap.stm you would inspect rr\_ff\_swap.log in the following ways.

```
grep 'DEBUG_[TP]' rr_ff_swap.log # 1424 lines come to the screen
OR
grep 'DEBUG_[TP]' rr_ff_swap.log > junk.txt # lines to temporary file junk.txt
```

Here are a few of the 1424 lines from the console, or from junk.txt when run with file redirection.

```
000000000001,MSG,thread 0 process 0,DEBUG_TKM1|None|[(2147483648, 0)]
000000000001,MSG,thread 0 process 0,DEBUG_TKM2|(290113888, 0)|[(1857369760, 290113888)]
000000000002,MSG,thread 0 process 1,DEBUG_TKM1|None|[(1857369760, 290113888)]
000000000002,MSG,thread 0 process 1,DEBUG_TKM2|(525642137, 290113888)|[(1331727623, 815756025)]
000000000003,MSG,thread 0 process 2,DEBUG_TKM1|None|[(1331727623, 815756025)]
000000000003,MSG,thread 0 process 2,DEBUG_TKM2|(569390073, 815756025)|[(762337550, 1385146098)]
000000000004,MSG,thread 0 process 3,DEBUG_TKM1|None|[(762337550, 1385146098)]
000000000004,MSG,thread 0 process 3,DEBUG_TKM2|(466581130, 1385146098)|[(295756420, 1851727228)]
000000000005,MSG,thread 0 process 4,DEBUG_TKM1|None|[(295756420, 1851727228)]
000000000005,MSG,thread 0 process 4,DEBUG_TKM2|(180058597, 1851727228)|[(115697823, 2031785825)]
000000000006,MSG,thread 0 process 5,DEBUG_TKM1|None|[(115697823, 2031785825)]
000000000006,MSG,thread 0 process 5,DEBUG_TKM2|None|[(115697823, 2031785825)]
... Lines 13 through 44 are not shown here.
000000000023,MSG,thread 0 process 22,DEBUG_TKM1|None|[(115697823, 2031785825)]
000000000023,MSG,thread 0 process 22,DEBUG_TKM2|None|[(115697823, 2031785825)]
000000000024,MSG,thread 0 process 23,DEBUG_TKM1|None|[(115697823, 2031785825)]
000000000024,MSG,thread 0 process 23,DEBUG_TKM2|None|[(115697823, 2031785825)]
000000000047,MSG,thread 0 process 1,DEBUG_PBR1|(525642137, 290113888)|[(115697823, 2031785825)]
000000000047,MSG,thread 0 process 1,DEBUG_PBR2|None|[(525642137, 290113888), (115697823, 2031785825)]
000000000047,MSG,thread 0 process 5,DEBUG_TKM1|None|[(525642137, 290113888), (115697823, 2031785825)]
000000000047,MSG,thread 0 process 5,DEBUG_TKM2|(192721691, 290113888)|[(332920446, 482835579), (115697823, 2031785825)]
000000000757,MSG,thread 0 process 0,DEBUG_PBR1|(290113888, 0)|[(332920446, 482835579), (115697823, 2031785825)]
000000000757,MSG,thread 0 process 0,DEBUG_PBR2|None|[(290113888, 0), (332920446, 482835579), (115697823, 2031785825)]
000000000757,MSG,thread 0 process 6,DEBUG_TKM1|None|[(290113888, 0), (332920446, 482835579), (115697823, 2031785825)]
000000000757,MSG,thread 0 process 6,DEBUG_TKM2|(314142014, 482835579)|[(290113888, 0), (18778432, 796977593), (115697823, 2031785825)]
000000000978,MSG,thread 0 process 2,DEBUG_PBR1|(569390073, 815756025)|[(290113888, 0), (18778432, 796977593), (115697823, 2031785825)]
000000000978,MSG,thread 0 process 2,DEBUG_PBR2|None|[(290113888, 0), (588168505, 796977593), (115697823, 2031785825)]
000000000978,MSG,thread 0 process 7,DEBUG_TKM1|None|[(290113888, 0), (588168505, 796977593), (115697823, 2031785825)]
000000000978,MSG,thread 0 process 7,DEBUG_TKM2|(484957072, 796977593)|[(290113888, 0), (103211433, 1281934665), (115697823, 2031785825)]
000000001022,MSG,thread 0 process 3,DEBUG_PBR1|(466581130, 1385146098)|[(290113888, 0), (103211433, 1281934665), (115697823, 2031785825)]
000000001022,MSG,thread 0 process 3,DEBUG_PBR2|None|[(290113888, 0), (569792563, 1281934665), (115697823, 2031785825)]
000000001022,MSG,thread 0 process 8,DEBUG_TKM1|None|[(290113888, 0), (569792563, 1281934665), (115697823, 2031785825)]
000000001022,MSG,thread 0 process 8,DEBUG_TKM2|(307537605, 1281934665)|[(290113888, 0), (262254958, 1589472270), (115697823, 2031785825)]
000000001286,MSG,thread 0 process 1,DEBUG_TKM1|None|[(290113888, 0), (262254958, 1589472270), (115697823, 2031785825)]
000000001286,MSG,thread 0 process 1,DEBUG_TKM2|None|[(290113888, 0), (262254958, 1589472270), (115697823, 2031785825)]
```

Simulation time 1 shows process 0 with an unmapped **pcb.relocationRegister** at the start of **takeMemory** (DEBUG\_TKM1), and with 290,113,888 elements in **pcb.relocationRegister** at the end of **takeMemory** (DEBUG\_TKM2). Both **pcb.relocationRegister** and **processor.freeregions** show memory regions as (SIZE, BASEADDRESS) pairs.

Process 5 through 23 are not able to satisfy their needs for contiguous memory from the 115,697,823 locations still available in **processor.freeregions** (the “L” in the trace stands for a Python long integer), so their **DEBUG\_TKM2** lines show **None** for a failed allocation in **pcb.relocationRegister**. Those processes must then wait in the **processor.waitMemQ** in state **waitForMemory** until released by a process releasing

memory going into state releasing or terminated. Inspecting the full .log files shows the entry into state **waitForMemory** for threads blocking within the **processor.waitMemQ**.

At time 47 in the trace, process 1 releases sufficient memory for process 5 to proceed. Recall that process 5 was the first process to fail to acquire memory at time 6, so it is at the front of the FIFO **processor.waitMemQ**.

You can create slightly shorter lines in your DEBUG trace with this Unix command.

```
grep 'DEBUG_[TP]' rr_ff_swap.log | sed -e 's/MSG.*thread 0 //'
OR
grep 'DEBUG_[TP]' rr_ff_swap.log | sed -e 's/MSG.*thread 0 //' > junk.txt
```

The **sed** command is a stream editor, and the command in single quotes removes all characters from MSG through thread 0 in the output from **grep**. Make sure to put a space on each side of the 0, and type your single quotes; do not copy & paste quotes from Word documents. This output appears as follows.

```
000000000001,process 0,DEBUG_TKM1|None|[(2147483648, 0)]
000000000001,process 0,DEBUG_TKM2|(290113888, 0)|[(1857369760, 290113888)]
000000000002,process 1,DEBUG_TKM1|None|[(1857369760, 290113888)]
000000000002,process 1,DEBUG_TKM2|(525642137, 290113888)|[(1331727623, 815756025)]
```

Do not try to wade through over 1000 lines of this trace output. When considering answers for the **5** questions in **README.txt**, first think about the effects of queuing (FIFO versus min-queue on process logical memory size) and memory allocation algorithm (first fit versus best fit), and then look at the trace to confirm your suspicions. For example, in comparing **rr\_ff\_swap.log** (“swap” is the FIFO log file) with **rr\_ff\_min.log** (the min-queue), take note of the difference at time 591.

```
grep 'DEBUG_[TP]' rr_ff_swap.log | sed -e 's/MSG.*thread 0 //' | less
grep 'DEBUG_[TP]' rr_ff_min.log | sed -e 's/MSG.*thread 0 //' | less
```

```
FIFO memory scheduling
000000000757,process 0,DEBUG_PBR1|(290113888, 0)|[(332920446, 482835579), (115697823, 2031785825)]
000000000757,process 0,DEBUG_PBR2|None|[(290113888, 0), (332920446, 482835579), (115697823, 2031785825)]
000000000757,process 6,DEBUG_TKM1|None|[(290113888, 0), (332920446, 482835579), (115697823, 2031785825)]
000000000757,process 6,DEBUG_TKM2|(314142014, 482835579)|[(290113888, 0), (18778432, 796977593), (115697823, 2031785825)]MIN-
QUEUE MEMORY SCHEDULING
000000000757,process 0,DEBUG_PBR1|(290113888, 0)|[(332920446, 482835579), (115697823, 2031785825)]
000000000757,process 0,DEBUG_PBR2|None|[(290113888, 0), (332920446, 482835579), (115697823, 2031785825)]
000000000757,process 17,DEBUG_TKM1|None|[(290113888, 0), (332920446, 482835579), (115697823, 2031785825)]
000000000757,process 17,DEBUG_TKM2|(202185513, 0)|[(87928375, 202185513), (332920446, 482835579), (115697823, 2031785825)]
```

The FIFO queue allocates 314,142,014 memory locations to process 6 because of its FIFO arrival order, while the min-queue allocates 202,185,513 locations to process 17 because of its minimal size.

First think about the combined effects of FIFO versus min-queuing, and first fit versus best fit, and formulate each answer for **README.txt**. Use the memory trace as needed, but don't get bogged down in it. If I give hints, it will be in class. Plan to be there.

After **make clean test** passes and you have reread all instructions and written your **5** answers in **README.txt**, run **make turnitin** on mcgonagall or **make turnitin** on acad before the project deadline.