

Exercise Setup

- For most of the exercises, we'll assume the following setup:
 - We have two processes to schedule.
 - One process is a **compute job**: it just eats CPU and never does IO or blocks for any reason.
 - One process is **IO bound**: it does very little computation, but frequently requests IO (and therefore blocks).
- In each exercise, you'll be told the amount of computation the IO bound process does and how long its IO takes. (In some cases, you'll also be told a **time slice** or **quantum**.)
- You should compute latency and throughput after you've filled in the tables.

FIFO

- I/O Job does 1 ms of computation then 10 ms of I/O.
- Compute Job is pure computation.

Time	I/O Job	Compute Job
0	Running	Ready

Round Robin w/ Long Time Slices

- I/O Job does 1 ms of computation then 10 ms of I/O.
- Compute Job is pure computation.
- **100 ms timeslice**

Time	I/O Job	Compute Job
0	Running	Ready

RR w/ Short Time Slices

- I/O Job does 1 ms of computation then 10 ms of I/O.
- Compute Job is pure computation.
- **1 ms timeslice**

Time	I/O Job	Compute Job
0	Running	Ready

STCF: Shortest Time to *Completion* First

- I/O Job does 1 ms of computation then 10 ms of I/O.
- Compute Job is pure computation.

Time	I/O Job	Compute Job
0	Running	Ready

STCF: Shortest Time to *Completion* First

- Process 1 does 1 ms of computation then 10 ms of I/O.
- Process 2 is pure computation.

Time	I/O Job	Compute Job
0	Running	Ready
1	Blocked (I/O)	Running
11	Running	Ready
12	Blocked (I/O)	Running
22	Running	Ready
23	Blocked (I/O)	Running
33	Running	Ready

Problems to discuss and solve

1. Prove that STCF will produce the minimum average response time.
2. Relative to RR and FCFS, how will STCF compare on maximum response time?
3. Since STCF cannot be implemented, derive a scheduler that attempts to approximate it.