Chapter 12 Exercises

1. Describe how you would use a negative feedback loop to manage your bank balance. Describe how you would do so with a positive feedback loop. Explain which you would prefer and why.

ANS:

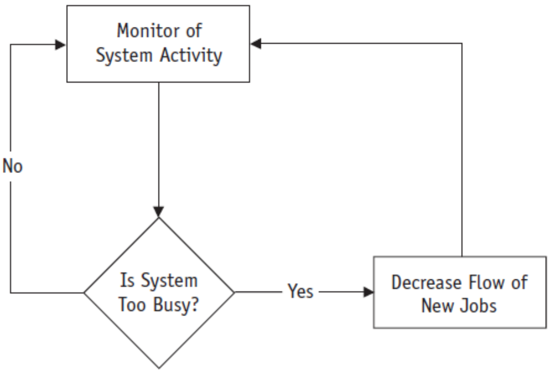

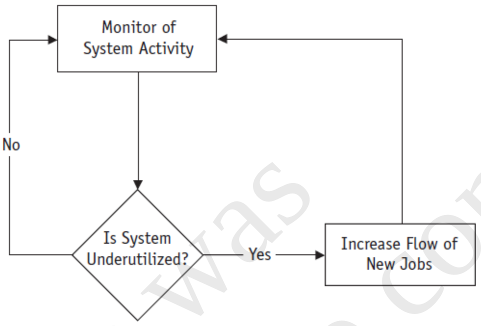

(a) negative feedback loop (b) positive feedback loop

There can be several answers to this question but here is a sample.

Negative feedback loop: Student monitors the checking account balance. Is it too low? If not, continue monitoring. If it is too low, stop writing checks until the next deposit is made. Check the balance and repeat this until the balance is no longer too low.

Positive feedback loop: Student monitors the checking account balance. Is it too high? If not, continue monitoring. If it is too high, start paying down credit card balances and take some time off work. Check the balance and repeat this until the balance is no longer too high.

2. Describe how a supermarket manager could use a positive feedback loop to direct waiting customers to four open cash register lanes, being careful to minimize waiting time for customers and maximize the speed of customer processing. Include a description of how you would monitor the system and measure its success. ANS: The manager could tell all employees to always wait for customers even if it means standing at the register even when no one is in line. This means these employees will have less time to restock the shelves. One way to measure the system is to measure the number of customers processed by each cashier. Another would be to measure the

amount of overtime required to restock shelves… something that could have been done when no one was in line and the cashiers were not needed at the registers.

3. Revisiting the example of the gas station lines and positive feedback loops, why do you think the length of gas lines will change? ANS: If you're looking for a gas station and the first one you find has few cars waiting in line, you collect the data and you react positively. Your processor suggests that you join the line and fill up while the line is short. As a result of using this positive feedback loop, long lines at some gas stations will grow when drivers join the line because it seems to be worth the wait. On the other hand, if you were using a negative feedback loop, you would see if the gas station had "too many cars waiting in line." If so, you would be tempted to drive on to another station.

   Are there are other factors that could impact the length of gas lines? Explain in your own words why or why not. ANS: One factor could include the speed with which customers get their gas and move out of the way for other waiting drivers; Another factor could be the ease of access to the station. Another factor could be the presence of another station with shorter lines. There can be several factors cited here.

4. Imagine that you are managing a university system that becomes CPU-bound at the conclusion of each school semester. If you were allowed to double the number of processors, what effect on throughput would you expect? If you could make one additional change to the system, what would it be? Explain why you'd expect your changes to improve performance? ANS:If the system is CPU-bound, then the addition of more processors would likely increase throughput allowing more jobs to be processed at the same time (assuming the operating systems software allows dynamic multiprocessing). This result is not guaranteed for all systems, but it should give students a starting point for an in depth discussion of system slowdowns and potential fixes.

5. Imagine that you are managing the system for a consulting company that becomes I/O-bound at the end of each fiscal year. What effect on throughput would you expect if you were allowed to double the number of processors? If you could make one additional change to the system, what would it be? Explain in your own words why you'd expect your changes to improve overall system performance? ANS: If a system is I/O bound, which could be caused, for example, by printing lengthy, numerous

reports for year-end analysis, then the addition of more processors will not noticeably speed up this situation. However, if you had the opportunity to add more I/O resources (notably, in this case, printers), these may well improve throughput remarkably.

6. Using the information below,
   - CPU W = 4 ms (the average time to execute 1,000 instructions)
   - Drive A = 45 ms (average data access speed)
   - Drive B = 3 ms (average data access speed)
   - Drive C = 17 ms (average data access speed)
   - Drive D = 10 ms (average data access speed)

   Calculate I/O access speed using this CPU and each of the four disk drives as they evaluate the following track requests in this order: 0, 31, 20, 15, 20, 31, 15. Then, in each case, calculate the access speeds after the track requests are reordered and rank the four disk drives before and after reordering. ANSWER: Using the same assumptions as those for Figure 12.2 (that the arm is already at the first location, track 0), without reordering, the arm will make 6 stops. After reordering, the arm will make 3 stops at tracks 15, 20, and 31. In these four configurations, reordering is always faster, but that is not always the case.

   CPU W + Disk A without reordering: 6 accesses @ 45ms each = 270 ms

   CPU W + Disk A with reordering: 4ms for reordering + 3 accesses @ 45ms each = 139 ms

   CPU W + Disk B without reordering: 6 accesses @ 3 ms each = 18 ms

   CPU W + Disk B with reordering: 4ms for reordering + 3 accesses @ 3 ms each = 13 ms

   CPU W + Disk C without reordering: 6 accesses @ 17ms each = 102 ms

   CPU W + Disk C with reordering: 4ms for reordering + 3 accesses @ 17ms each = 55 ms

   CPU W + Disk D without reordering: 6 accesses @ 10 ms each = 60 ms

   CPU W + Disk D with reordering: 4ms for reordering + 3 accesses @ 10 each = 34 ms

7. Using the information below,
   - CPU X = 0.3 ms (the average time to execute 1,000 instructions)

- Drive A = 0.7 ms (average data access speed)
- Drive B = 4 ms (average data access speed)
- Drive C = .03 ms (average data access speed)

Calculate I/O access speed using this CPU and each of the disk drives as they evaluate the following track requests in this order: 16, 4, 9, 16, 29, 31, 5. Then, in each case, calculate the access speeds after the track requests are reordered and rank the three disk drives before and after reordering. ANS: Using the same assumptions as those for Figure 12.2 (that the arm is already at the first location, track 0), without reordering, the arm will make 7 stops. After reordering, the arm will make 3 stops at tracks 15, 20, and 31. In these cases, reordering is the same or faster, but that is not always the case.

CPU X + Disk A without reordering: 7 accesses @ 0.7ms each = 4.9 ms

CPU X + Disk A with reordering: 0.3ms for reordering + 6 accesses @ 0.7ms each = 4.5 ms

CPU X + Disk B without reordering: 7 accesses @ 4 ms each = 28 ms

CPU X + Disk B with reordering: 0.3ms for reordering + 6 accesses @ 4 ms each = 24.3 ms

CPU X + Disk C without reordering: 7 accesses @ 0.3ms each = 2.1 ms

CPU X + Disk C with reordering: 0.3ms for reordering + 6 accesses @ 0.3ms each = 2.1 ms

8. Remembering that there's a trade-off between memory use and CPU overhead, give an example where increasing the size of virtual memory will improve job throughput. Then give an example where doing so will cause throughput to suffer and explain why this is so. ANS: One way to improve system performance, sometimes remarkably better performance, is to increase the size of virtual memory to supplement the RAM available on a system if it is often CPU-bound system.

   However, if the system has a limited amount of free secondary storage, and increased virtual memory reduces the amount of free storage to critical levels, then overall system performance will suffer until secondary storage is increased to at least a minimum level.

9. Looking back over the past 12 months, let's say your own computer had failed unexpectedly and catastrophically twice in that time. Identify the worst possible time for failure and the best possible time. Then compare the time and cost it would have

required for you to recover from those two catastrophic failures. Describe in your own words the factors that differentiated the worst experience from the best. ANS: The goal of this question is to open a dialog about the need for system availability and reliability (and backing up important data on a regular basis). These answers will vary but might suggest that the worst time would be when the computer had a crucial role during a critical deadline situation. Alternatively, the best time could be when the critical deadline has passed, one has just backed up the system, and there was sufficient time and money available to recover the system. One might also suggest that any computer that fails catastrophically twice in a year is due to be replaced, thus requiring purchasing power and time to customize a new computer out of the box.

10. Describe how you would convince a coworker to better manage a personal computer by performing regular backups and keep the system patches current. ANS: This question offers an opportunity to discuss ways to communicate with coworkers on technical topics. There can be many answers to this question. Here are several:

- Teach them how frequent patches can reduce downtime and improve security.
- Stress the time required to recover the system (compared to the time to back up the computer) if it should crash.
- Stress the vital nature of the data to that individual and the benefits of having the system remain accurate and available.
- Explore the possible legal liability to the organization, and the personal consequences to the coworker, if the data protection is deemed negligent.

11. Calculate the reliability of a hard disk drive with an MTBF of 2,499 hours during the last 40 hours of this month. Assume e = 2.71828 and use the formula.

$$\text{Reliability}(t) = e^{-(1/\text{MTBF})(t)}$$

ANS: (Remember that *t* is the time period and notice that all units for this question are expressed in hours.)

$\text{Reliability}(t) = e^{-(1/\text{MTBF})(t)}$

$\text{Reliability}(t) = 2.71828^{-(1/2499)(40)}$

$\text{Reliability}(t) = 2.71828^{-(0.0004)(40)}$

$\text{Reliability}(t) = 2.71828^{-(0.016)}$

$\text{Reliability}(t) = 0.98412$

Hint: to perform the math, students may find it helpful to use an Exponent Calculator, and remembering that this exponent is a negative number.

12. Calculate the reliability of a hard disk drive with an MTBF of 4,622 hours during the crucial last 16 hours of the last fiscal quarter (the three-month period beginning October 1 and ending December 31). Assume e = 2.71828 and use the reliability formula from the previous exercise. ANS: (Remember that *t* is the time period and all the units for this question are expressed in hours.)

    $Reliability(t) = e^{-(1/MTBF)(t)}$

    $Reliability(t) = 2.71828^{-(1/4622)(16)}$

    $Reliability(t) = 2.71828^{-(0.0002)(16)}$

    $Reliability(t) = 2.71828^{-(0.0032)}$

    $Reliability(t) = 0.99680$

13. Calculate the reliability of a server with an MTBF of 10,500 hours during the summer busy selling season from May 1 through September 15. Assume that the server must remain operational 24 hours/day during that entire time. Hint: Begin by calculating the number of hours of operation during the busy season. Assume e = 2.71828 and use the reliability formula from the previous exercises. ANS: (Begin by multiplying 24 times the number of days: May = 31, June = 30, July = 31, August = 31, September = 15. Therefore, busy season is 138 days * 24 = 3,312 hours.)

    $Reliability(t) = e^{-(1/MTBF)(t)}$

    $Reliability(t) = 2.71828^{-(1/10500)(3312)}$

    $Reliability(t) = 2.71828^{-(.000095)(3312)}$

    $Reliability(t) = 2.71828^{-(0.3146)}$

    $Reliability(t) = 0.73008$

14. Calculate the availability of a server with an MTBF of 75 hours and an MTTR of 3 days (72 hours). ANS: 0.510

    Availability = MTBF/(MTBF+MTTR)

    Availability = 75/(75+72)

    Availability = 75/(147) = 0.510

15. Calculate the availability of a hard disk drive with an MTBF of 2,040 hours and an MTTR of 8.5 hours. ANS: 0.995

    Availability = MTBF/(MTBF+MTTR)

    Availability = 2040/(2040+8.5)

    Availability = 2040/(2048.5) = 0.995

16. Calculate the availability of a server with an MTBF of 50 weeks and an MTTR of 798 hours. ANS: 0.913

    Assume continuous around-the-clock operation. First calculate the number of MTBF hours so the units are all expressed in hours:

    50 weeks in hours = 24 (hours) * 7 (days) * 50 (weeks) = 8,400 hours

    Then calculate the availability:

    Availability = MTBF/(MTBF+MTTR)

    Availability = 8400/(8400+798)

    Availability = 8400/(9198) = 0.913

Advanced Exercises

17. Compare and contrast availability and reliability. In your opinion, which is more important to a system manager? Substantiate your answer in your own words. ANS: These answers will vary. For different systems, different measures of success will be more important. Use this exercise to explore the importance of availability and reliability to different systems managers.

18. In this chapter, we described the trade-offs among all the managers in the operating system. Study a system to which you have access and, assuming you have sufficient funds to upgrade only one component for the system, explain which component you would choose to upgrade to improve overall system performance, and explain why. ANS: Answers to this question will vary for the class. Use this exercise to explore the concept of bottlenecks and how a single improvement in a system will not (perhaps cannot) eliminate all bottlenecks, but could make them more manageable or move the bottleneck to a region of the system that is more acceptable. For example, if the bottleneck is at the printer, adding more printers will move the bottleneck to the CPU where jobs will now wait for processing, instead of waiting in the print queue.

19. Perform a software inventory of your computer system to identify all applications resident on the computer. If you have more than a dozen applications, choose 12 of them and check each for patches that are available for your software. For each vendor, identify how many patches are available for your software, the number of critical patches, and what patch cycle you would recommend for that vendor's patches: annual, quarterly, monthly, or weekly updates. ANS: These answers will vary depending on each system. We suggest that this be an opportunity to stress the need to keep software patched correctly and the continuing threat to all data on the system when application and system software is vulnerable. Use this exercise as an opportunity to discuss the availability of patches and the appropriate patch cycle for different types of applications and system software.
20. As memory management algorithms grow more complex, the CPU overhead increases and overall performance can suffer. On the other hand, some operating systems perform remarkably better with additional memory. Explain in detail (using your own words) why some perform better with additional memory. Then explain in detail why some do not perform better. ANS: The answers to this question will depend on the operating system in use. For example the newer Windows operating systems can be configured to make use of extra storage available in a flash memory device to use it as extra memory space and thus improve performance. An operating system without room or capabilities to use additional storage as memory expansions would not benefit.
21. Conduct an inventory of your computer to discover how many processes are active when it is connected to the Internet. What is the total number of processes currently running? For which processes can you identify the application? How many processes are linked to applications that you cannot identify? Discuss how the unidentifiable processes pose a challenge regarding effective system administration and what you would do to address this challenge. ANS: For each system examined the answers to this question will be different. We hope students will use this exercise as an opportunity to discuss the multitude of software that runs on different computing devices and the vulnerabilities they present to the systems/network manager.
22. Compare and contrast throughput, turnaround time, and response time. Explain what each measures and how they are monitored. Which measurement is more important in

your current computing environment? Explain why. ANS: The three are all measurements of a computer system's efficiency. The following is taken from the book's discussion of Measurement Tools.

Throughput is a composite measure that indicates the productivity of the system as a whole; the term is often used by system managers. Throughput is usually measured under steady-state conditions and reflects quantities such as "the number of jobs processed per day" or "the number of online transactions handled per hour." Throughput can also be a measure of the volume of work handled by one unit of the computer system, an isolation that's useful when analysts are looking for bottlenecks in the system.

Turnaround time is the time from the submission of the job until its output is returned to the user. Whether in an online or batch context, this measure depends on both the workload being handled by the system at the time of the request and the type of job or request being submitted. Some requests, for instance, might be handled faster than others because they require fewer resources.

Response time is the interval required to process a user's request and is an important measure of system performance. It's measured from when the user presses the key to send the message until the system indicates receipt of the message.

23. Availability indicates the likelihood that a resource will be ready when a user needs it. Explain in your own words how the measurements to calculate this value are collected. ANS:

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

Therefore, to calculate the availability of a given device, we will need both the MTBF (mean time between failures) and MTTR (mean time to repair) for that device. The MTBF is given by the manufacturer. The MTTR, on the other hand, needs to be calculated and will depend on a variety of factors including the seriousness of the malfunction, the availability of technicians and parts, etc.