Name: ____________________________

Directions: Work only on this sheet (on both sides, if needed); do not turn in any supplementary sheets of paper. There is actually plenty of room for your answers, as long as you organize yourself BEFORE starting writing. In order to get full credit, SHOW YOUR WORK.

1. This problem concerns the second machine-repair model in our SimPy examples. Recall that there are two machines, which sometimes break down. Up time is exponentially distributed with rate $\lambda_b = 1.0$, and repair time is exponentially distributed with rate $\lambda_r = 0.5$. There is only one repairperson. This is a Markov chain, with state i meaning that i machines are up. The long-run state probabilities are $\pi_0 = 1/(1 + \frac{\lambda_r}{\lambda_b} + \frac{\lambda_r^2}{2\lambda_b^2})$, $\pi_1 = \frac{\lambda_r}{\lambda_b}\pi_0$ and $\pi_2 = \frac{\lambda_r^2}{2\lambda_b^2}\pi_0$.

(a) (10) Suppose we wish to find $\nu = E(U_2 - U_1)$, where U_1 is a time when a breakdown of one machine occurs when the other is already down and U_2 is the next time that both machines are up. Citing line numbers, show how to modify the program MachRep2.py to determine ν via simulation.

(b) (10) Find ν in (a) analytically. (Give a numerical answer, not an expression with letters.)

(c) (10) Consider the following similar model: When both machines are up, one acts simply as a spare and thus cannot break down. But there are two repairpersons, so if both machines are down, both can be repaired simultaneously. Using only the information above and material from our course, find the long-run proportion of the time that both machines are down. Your answer must be purely in terms of λ_b and λ_r and constants. (The symbol π must not appear in your answer.) Make sure to justify your answer fully and clearly.

2. Consider the following program which uses simulation to find the probability of getting three heads in five tosses of a coin:

```
1    import random,math
2    rand = random.Random(12345)
3    simn = 10000
4    count3 = 0
5    for i in range(simn):
6       heads = 0
7       for toss in range(5):
8          if rand.uniform(0.0,1.0) < 0.5:
9             heads += 1
10      if heads == 3:
11         count3 += 1
12   print float(count3)/simn
```

(a) (10) Write one or more lines of code to replace line 12, which will print out an approximate 95% confidence interval for P(3 heads in 5 tosses).

(b) (10) Of course, that 95% is just an approximate figure. But we can use simulation to find the exact confidence level, since we know the actual value of P(3 heads in 5 tosses), 0.3125. Show what code to add to do this. (Be sure to state the line number(s) at which your code is to be added.)

3. This problem concerns the cell-phone analysis from our printed lecture notes.

(a) (10) Show the balance equation corresponding to transitions in and out of state n-g.

(b) (10) Find the per-channel utilization, u. Express your answer in terms of the λ_i, μ_i and π_i.

(c) (10) Find the long-run proportion of accepted calls which are handoff calls. Express your answer in terms of the λ_i, μ_i and π_i.

4. (10) Suppose in a certain queuing system jobs come from k different categories. Category i comprises a fraction α_i of all jobs, and within that category, the mean and variance of the service time are μ_i and σ_i^2, respectively. Let S denote the service time and C denote the job class. Find $Var(S)$ in terms of the μ_i and σ_i^2. You must make use of special material in our course; long derivations will get only partial credit.

5. (10) Suppose a disk will store backup files. We place the first file in the first track on the disk, then the second file right after the first in the same track, etc. Occasionally we will run out of room on a track, and the file we are placing at the time must be split between this track and the next. The amount of room X taken up by a file (a continuous random variable in this model) is uniformly distributed between 0 and 3 tracks. Find the long-run proportion of tracks which have the property that contain only one file. (The file may extend onto other tracks as well.)

Selected Solutions:

1.b $ET_{02} = \frac{1}{0.5} + ET_{12}$. $ET_{12} = \frac{1}{1.5} + \frac{1}{1.5} ET_{02}$, etc.

1.c The Markov chain for the new model is the reversed version of the original one. Since the latter is a birth/death process, it is reversible. Thus the two chains have the same π_i structure, with the roles of λ_b and λ_r interchanged.

2.b

```
import random,math

rand = random.Random(12345)
numberin = 0
simn = 1000
for rep in range(10000):
   count3 = 0
   for i in range(simn):
      heads = 0
      for toss in range(5):
         if rand.uniform(0.0,1.0) < 0.5:
            heads += 1
      if heads == 3:
         count3 += 1
   ph = float(count3)/simn
   radius = 1.96 * math.sqrt(ph*(1.0-ph)/simn)
   if math.fabs(ph-0.3125) < radius:
      numberin += 1
print numberin/10000.0
```

3.a

$$\pi_{n-g}[\lambda_2 + (n-g)\mu] = \pi_{n-g+1} \cdot (n-g+1)\mu + \pi_{n-g-1}\lambda$$

3.b During the time in state i, the utilization of the system is i/n. Thus the long-run utilization is $\sum_{i=0}^{n} \pi_i \frac{i}{n}$.

3.c Calls internal to the cell are generated at the rate λ_1, and are accepted if the system is in state $i \leq n-g-1$. Thus the number of internal calls accepted per unit time is

$$\lambda_1 \sum_{i=0}^{n-g-1} \pi_i$$

On the other hand, handoff calls come in at the rate λ_1, and are accepted if the system is in state $i \leq n-1$, so the number accepted per unit time is

$$\lambda_2 \sum_{i=0}^{n-1} \pi_i$$

Thus the proportion of accepted calls which are handoff calls is

$$\frac{\lambda_2 \sum_{i=0}^{n-1} \pi_i}{\lambda_1 \sum_{i=0}^{n-g-1} \pi_i + \lambda_2 \sum_{i=0}^{n-1} \pi_i}$$

4.

$$Var(S) = E[Var(S|C)] + Var[E(S|C)] = \sum_{i=1}^{k} \alpha_i \sigma_i^2 + \sum_{i=1}^{k} \alpha_i (\mu_i - \bar{\mu})^2$$

where $\bar{\mu} = \sum_{i=1}^{k} \alpha_i \mu_i$.

5. Think of the disk as consisting of a Very Long Line, with the end of one track being followed immediately by the beginning of the next track. The points at which files begin then form a renewal process, with "time" being distance along the Very Long Line. If we observe the disk at the end of the k_{th} track, this is observing at "time" k. That track consists entirely of one file if and only if the "age" A of the current file—i.e. the distance back to the beginning of that file—is greater than 1.0.

Our material on age and residual-life distribution showed that for large k the density of A is (approximately)

$$f_A(t) = \frac{1 - \frac{t}{3}}{1.5} = \frac{2}{3} - \frac{2}{9}t$$

using the fact that for a U(0,3) distribution the c.d.f. is t/3 and the mean is 1.5. So

$$P(A > 1) = \int_1^3 \left(\frac{2}{3} - \frac{2}{9}t \right) \, dt = \frac{4}{9}$$