

Name: \_\_\_\_\_

Directions: **Work only on this sheet** (on both sides, if needed). **MAKE SURE TO COPY YOUR ANSWERS TO A SEPARATE SHEET FOR SENDING ME AN ELECTRONIC COPY LATER.**

**IMPORTANT NOTE:** If you believe that nothing needs to be placed into a blank, simply give **Nothing** as your answer in your file. If you do not answer at all, put 00 in your file.

1. (70) The following Hadoop code multiplies a vector  $\mathbf{x}$  by a (presumably very large) matrix  $\mathbf{a}$ . The input matrix has prepended to it a column of row numbers. The vector  $\mathbf{x}$  is input by executing code in a file  $\mathbf{x.py}$ .

So, if the input matrix is

```
0 1 2 0
1 5 8 -4
2 0 0 3
```

and the contents of  $\mathbf{x.py}$  are

```
x = [5,12,13]
```

then the final output will be

```
29
69
39
```

**Note:** No row/element numbers in the final output. Don't worry about leading blanks in the output.

Fill in the blanks. You may find the Python `len()` function useful; it returns the length of a Python list (array), so that for instance `len(x)` is 3 in the above example. Also, the `int()` function is like `atoi()` in C.

**axmap.py:**

```
1 #!/usr/bin/env python
2
3 from x import x # input x from file x.py
4 import sys
5
6 for line in sys.stdin:
7     tks = line.split()
8     rownum = tks[0]
9     row = tks[1:]
10    sum = 0
11    for i in range(BLANKa):
12        sum += BLANKb
13    print BLANKc
```

**axred.py:**

```
1 #!/usr/bin/env python
2
3 import sys
4
5 for line in sys.stdin:
6     line = line.strip()
7     tks = line.split('\t')
8     print BLANKd
```

2. (30) Fill in the blanks in the Snow code below, which finds the unique elements of an array in parallel. The built-in R function `unique()` works like this:

```
1 > x <- sample(1:8,10,replace=T)
2 > x
3 [1] 4 7 3 1 1 2 3 3 2 8
4 > unique(x)
5 [1] 4 7 3 1 2 8
```

Code:

```
1 # not claimed efficient, and
2 # no guarantee of ordering in result
3
4 parunique <- function(cls,x) {
5     parts <- clusterSplit(cls,1:length(x))
6     xparts <- lapply(parts,function(part) x[part])
7     tmp <- clusterApply(cls,xparts,BLANKa)
8     tmp <- Reduce(BLANKb)
9     BLANKc
10 }
```

## Solutions:

### 1.

#### axmap.py:

```
1  #!/usr/bin/env python
2
3  from x import x # input x from file x.py
4  import sys
5
6  for line in sys.stdin:
7      tks = line.split()
8      rownum = tks[0]
9      row = tks[1:]
10     sum = 0
11     for i in range(len(row)):
12         sum += int(row[i]) * x[i]
13     print '%s\t%s' % (rownum, sum)
```

#### axred.py:

```
1  #!/usr/bin/env python
2
3  import sys
4
5  for line in sys.stdin:
6      line = line.strip()
7      tks = line.split('\t')
8      print tks[1]
```

### 2.

```
1  # not claimed efficient, and no guarantee of ordering in result
2
3  parunique <- function(cls,x) {
4      parts <- clusterSplit(cls,1:length(x))
5      xparts <- lapply(parts,function(part) x[part])
6      tmp <- clusterApply(cls,xparts,unique)
7      tmp <- Reduce(c,tmp)
8      unique(tmp)
9  }
```