Homework 1
Due Tuesday, March 21, 1995

Question 1

You are given a hierarchical shared-memory machine like the SGI 4D/340. Assume that the processors have a cache line size of 32 bytes. Now consider the following simple loop:

```
for i = 0 to 16
   for j = 0 to 255 {
      A[j] = do_something(A[j]);
   }
```

(i) Under what conditions would it be better to use a dynamically-scheduled loop?
(ii) Under what conditions would it be better to use a statically-scheduled loop?
(iii) For a dynamically-scheduled inner loop, how many iterations should a PE pick each time if A is an array of 4-byte integers?

Question 2

You are writing an image processing code, where the image is represented as a 2D array of pixels. The basic iteration in this computation consists of:

```
for i = 1 to 1024
   for j = 1 to 1024
      newA[i,j] = (A[i,j-1] + A[i-1,j] + A[i,j+1] + A[i+1,j])/4;
```

Assume `A` is a matrix of 4-byte single precision floats stored in row-major order (i.e `A[i,j]` and `A[i,j+1]` are at consecutive addresses in memory). `A` starts at memory location 0. You are writing this code for a 32 processor machine. Each PE has a 32 Kbyte direct-mapped cache and the cache lines are 64 bytes.

(a) In the first decomposition you try, you divide the matrix into sets of 32 rows each. What is the computation to communication ratio that you expect? State any assumptions that you make.

(b) Next you try a column-oriented decomposition. Each PE gets a contiguous set of 32 columns. What do you expect the computation to communication ratio to be in this case? State any assumptions that you make.

(c) Suppose the matrix `A` started at memory location 32 rather than address 0. If we use the same decomposition as in (b), do you expect this to change the computation to communication ratio? If yes, will it increase or decrease and why? If not, why not?

Question 3

You have been given the job of creating a word count program for a major book publisher. You will be working on a shared-memory multiprocessor with 32 processing elements. Your only stated interface is “get_words” which returns an array of the book’s next 1000 words to be counted.

The main work each processor will do should look like this:

```
while(get_words(word)) {
   for (i=0; i<1000; i++) {
      if word[i] is in list
         increment its count
      else
         add word to list
}

/* Once all words have been logged,
   the list should printed out */
```

Using pseudocode, create a detailed description of the control flow and data structures that you will use for this program. Your method should attempt to minimize space, synchronization overhead, and memory latency. This problem allows you a great deal of flexibility, so be sure to clearly state all assumptions and design decisions.