Homework 3: due by 11:55pm on Friday, February 22, 2013 (Total: 100 points)

Instructor: Vivek Sarkar

All homeworks should be submitted in a directory named hw_3 using the turn-in script. In case of problems using the script, you should email a zip file containing the directory to `comp322-staff@mailman.rice.edu` **before the deadline. See course wiki for late submission penalties.**

IMPORTANT: The programming assignment in this homework is more challenging than in past homeworks. It is critical that you start early on the programming assignment to meet the deadline.

Honor Code Policy: All submitted homeworks are expected to be the result of your individual effort. You are free to discuss course material and approaches to problems with your other classmates, the teaching assistants and the professor, but you should never misrepresent someone elses work as your own. If you use any material from external sources, you must provide proper attribution.

1 Written Assignments (25 points total)

Please submit your solutions to the written assignments in either a plain text file named `hw_3_written.txt` *or a PDF file named* `hw_3_written.pdf` *in the hw_3 directory.*

1.1 Future Tasks and Data-Driven Futures

1. (10 points) *Summarize the similarities and differences between futures and data-driven futures in HJ.* In your summary, you should state if it is possible to create a race condition or a deadlock when accessing the value in the future container when using either construct.

2. (15 points) Consider the HJ code fragment below that operates on an array of DataDrivenFutures. (Note that there are no `get()` operations in this example, so the only purpose of the `put()` operation is to synchronize with `await` clauses. For simplicity, we just use an empty string `""` as the object being put into a DataDrivenFuture.)

   Is it possible for any instance of the async statement in line 6 to be indefinitely blocked on its await clause? If so, explain how. If not, explain why not.

```
1.      DataDrivenFuture[] A = new DataDrivenFuture[n];
2.        for (int i = 0 ; i < n ; i++) {
3.            A[i] = new DataDrivenFuture();
4.        }
5.        for (int i = n-1 ; i >= 1 ; i--) {
6.            async await(A[i-1]) {
7.                . . .
8.                A[i].put();
9.            } // async
10.     } // for
11.     A[0].put();
```

2 Programming Assignment (75 points)

2.1 Pairwise Sequence Alignment

In this homework, we will focus on the *pairwise sequence alignment* problem in evolutionary and molecular biology, and how parallelism can help in solving this problem. (This homework is adapted from Homework 7 from the Spring 2011 offering of COMP 182 by Prof. Luay Nakhleh.)

Let X and Y be two sequences over alphabet Σ (for DNA sequences, $\Sigma = \{A, C, T, G\}$). An *alignment* of X and Y is two sequences X' and Y' over the alphabet $\Sigma \cup \{-\}$, where X' is formed from X by adding only dashes to it, and Y' is formed from Y by adding only dashes to it, such that

1 $|X'| = |Y'|$ *i.e.,* X' and Y' have the same size,

2 there does not exist an i such that $X'[i] = Y'[i] = -$, and

This alignment is also referred to as *global pairwise alignment* (as opposed to *local pairwise alignment*, which is used to align selected regions of sequences X and Y).

Sequence alignment helps biologists make inferences about the evolutionary relationship between two DNA sequences. Aligning two sequences amounts to "reverse engineering" the evolutionary process that acted upon the two sequences and modified them so that their characters and their lengths differ. As an example, one possible alignment of the two sequences $X = ACCT$ and $Y = TACGGT$ is as follows:

$$\begin{array}{cccccccc} X' & = & - & A & C & - & C & T \\ Y' & = & T & A & C & G & G & T \end{array}$$

As you may imagine, there may be multiple alignments for the same pair of sequences. For example, a trivial alternate alignment for X and Y is as follows:

$$\begin{array}{llllllllllll} X'' & = A & C & C & T & - & - & - & - & - & - \\ Y'' & = - & - & - & - & T & A & C & G & G & T \end{array}$$

2.2 Scoring in Pairwise Sequence Alignment: Optimality Criterion

As discussed above, a number of alignments exist for a given pair of sequences; therefore, we define a *scoring scheme* that gives higher scores to "better" alignments. Once the scoring scheme is defined, we seek an alignment with the highest score (among all feasible alignments). For DNA, a scoring scheme is given by a 5×5 matrix M, where for $p, q \in \{A, C, T, G\}$, $M_{p,q}$ specifies the score for aligning p in sequence X' with q in sequence Y', $M_{p,-}$ denotes the penalty for aligning p in sequence X' with a dash in sequence Y', and $M_{-,q}$ denotes the penalty for aligning q in sequence Y' with a dash in sequence X'. Assuming $|X'| = |Y'| = k$, the score of the alignment is

$$\sum_{i=1}^{k} M_{X'[i],Y'[i]}. \tag{1}$$

For this assignment, we will assume the following scoring scheme: $M_{p,p} = 5$, $M_{p,q} = 2$ (for $p \neq q$), $M_{p,-} = -2$ and $M_{-,q} = -4$.

For this scoring scheme. the score of the (X', Y') alignment in Section 2.1 is

$$M_{-,T} + M_{A,A} + M_{C,C} + M_{-,G} + M_{C,G} + M_{T,T} = (-4) + 5 + 5 + (-4) + 2 + 5 = 9$$

and the score of the (X'', Y'') alignment is $4 \times M_{p,-} + 6 \times M_{-,q} = -32$.

2.3 Sequential Algorithm to compute the Optimal Scoring for Pairwise Sequence Alignment

In this problem, we introduce a sequential dynamic programming algorithm (called the Smith-Waterman algorithm) to compute the Optimal Scoring for Pairwise Sequence Alignment. For two sequences X and Y of lengths m and n, respectively, denote by $S[i,j]$, $0 \le i \le m$ and $0 \le j \le n$, the score of the best alignment of the first i characters of X with the first j characters of Y. The boundary values are, $S[i,0] = i * M_{p,-}$ and $S[0,j] = j * M_{-,p}$. It has been shown that this optimal scoring can be defined as follows $\forall i,j \ge 1$:

$$S[i,j] = \max \begin{cases} S[i-1,j-1] + M_{X[i],Y[j]} \\ S[i-1,j] + M_{X[i],-} \\ S[i,j-1] + M_{-,Y[j]} \end{cases} . \tag{2}$$

The above definition directly leads to a sequential dynamic programming algorithm that can be implemented as shown in Listing 1. Assume that the input sequences are represented as Java strings, and the scoring matrix, S, is represented as a 2-dimensional array of size (*X.length()+1*) $\times$ (*Y.length()+1*). After the algorithm terminates, the final score is available in *S[X.length()][Y.length()]*.

The dependence structure of the iterations in Listing 1 is shown in Figure 1. The cells in the figure correspond to $S[i,j]$ values, and the arrows show the dependences among the $S[i,j]$ computations.

```
1  for (point[i,j] : [1:X.length(),1:Y.length()] ){
2     char xChar = X.charAt(i-1);
3     char YChar = Y.charAt(j-1);
4     int diagScore = S[i-1][j-1] + M[charMap(xChar)][charMap(YChar)];
5     int topColScore = S[i-1][j] + M[charMap(xChar)][0];
6     int leftRowScore = S[i][j-1] + M[0][charMap(YChar)];
7     S[i][j] = Math.max(diagScore, Math.max(leftRowScore, topColScore));
8    }
9  }
10 int finalScore = S[X.length()][Y.length()];
```

Listing 1: Sequential implementation of Smith-Waterman Algorithm for Optimal Scoring for Pairwise Sequence Alignment

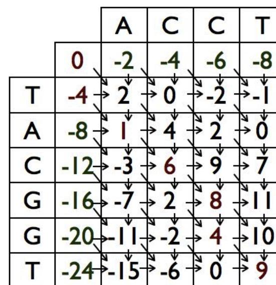

Figure 1: Dependences in Pairwise Sequence Alignment

This homework focuses on computing the optimal score for pairwise sequence alignment, not on the alignment itself. Though a biologist is ultimately interested in seeing the alignment, there are many applications where the score alone is of interest. For example, in multiple sequence alignment, the most commonly used approach is called progressive alignment, where an evolutionary tree is first built based on the scores of pairwise alignments, and then the tree is used as a guide for doing the multiple sequence alignment. In this case, the pairwise alignments are performed solely for the sake of obtaining scores, and the alignments themselves are not needed. However, it is important to compute the scores as quickly as possible when exploring alignments of large DNA sequences.

2.4 Your Assignment: Parallel Optimal Scoring for Pairwise Sequence Alignment

Your assignment is to design and implement parallel algorithms for optimal scoring for pairwise sequence alignment. You are free to use whichever HJ runtime system you choose (*e.g.,* work-sharing vs. work-stealing) that supports the constructs that you need and delivers good performance for your code.

We have provided a sequential implementation of the algorithm in `SeqScoring.hj` that you can use as a starting point. Your homework deliverables are as follows. *To facilitate automated testing of your programs, both programs that you submit must take exactly two command-line arguments,* string1 *and* string2*, as in* `SeqScoring.hj`*. Please use the provided templates for the files,* `UsefulParScoring.hj` *and* `SparseParScoring.hj`*.*

1. **Useful parallelism on Sugar compute nodes (35 points)** Create a new parallel version of `SeqScoring.hj` that is designed to achieve the smallest execution time using 8 cores on a dedicated Sugar compute node, and call it `UsefulParScoring.hj`.

   For this part of the assignment, we recommend first debugging your solution on small strings for correctness (which can be done on any platform), and then evaluating the performance of your implementation with pairs of strings of length $O(10^4)$ on dedicated Sugar compute nodes (as explained in Lab 4).

   Since each Sugar compute node has 16GB of memory, it is recommended that you increase the maximum heap size to 8GB by using the -mx option when running HJ programs on a Sugar compute node: "hj -mx 8000m -places 1:8 UsefulParScoring *string1 string2* ".

2. **Sparse memory version and useful parallelism on Sugar compute nodes (40 points)**

   The sequential algorithm outlined in Listing 1 and `SeqScoring.hj` allocates and uses a two-dimensional matrix which requires $O(n^2)$ space when processing strings of size $O(n)$. The goal of this part of the assignment is to create a *sparse* memory version of the program that can process strings of length $O(10^5)$ or greater by using space that's less than $O(n^2)$. The key idea to think about is what data really needs to be retained as the computation advances. For example, in the sequential version, row 1 of the S matrix can be freed (set to null and garbage-collected, or reused elsewhere) when the computation reaches row 3, since computation of row 3 only needs row 2 and not row 1.

   You will need to design and implement an analogous approach to reducing the space requirements of the parallel version. This will require reworking the data structure for matrix S, and may even require using a different algorithm from `UsefulParScoring.hj`. Call this version `SparseParScoring.hj`. As before, we recommend first debugging your solution on small strings for correctness, and then testing with pairs of strings of length $O(10^5)$ on dedicated Sugar compute nodes. These runs may last a few minutes, so be sure to run these computations only on Sugar compute nodes, and not on the login node. Also, there may be be some impact of Java's garbage collection (GC) on the performance you observe. Please contact a teaching staff member if you believe that GC overheads are interfering with your performance measurements.

3. [**Homework report** You should submit a brief report summarizing the design of your parallel algorithms in `UsefulParScoring.hj` and `SparseParScoring.hj`, explaining why you believe that each implementation is correct and data-race-free.

   Your report should also include the following measurements for `UsefulParScoring.hj` and `SparseParScoring.hj`:

   (a) Execution time of `SeqScoring.hj` and `UsefulParScoring.hj` on a Sugar compute node with inputs of length (approximately) 10,000. `SeqScoring.hj` can be executed with the default runtime options. `UsefulParScoring.hj` should be executed with the "-places 1:8" option to run with 8 workers (so as to use all 8 cores).

   (b) Execution time of sequential and parallel versions of `SparseParScoring.hj` with inputs of length (approximately) 100,000. As before, the sequential version should be executed without the default options options, and the parallel version with the "-places 1:8" option. The sequential HJ version can often (but not always) be obtained by removing all parallel keywords (async, finish, etc.)

2.5 Generation of Test Data

You are welcome to generate any test data that you choose to debug your programs. Just keep in mind that they need to be strings of characters in $\{A, C, T, G\}$.

This is optional, but if you are interested in generating pairs of DNA sequences under realistic models of evolution, you can use a free web service available at `http://bibiserv.techfak.uni-bielefeld.de/rose/submission.html`. You should use the following options for this web service (with default values for everything else):

- *Sequence type:* Select DNA.
- *Number of sequences:* Enter 2.
- *Average length of sequences:* Enter whatever length you need e.g., 10, 10000, 100000, etc.
- *Average pairwise distance:* This parameter impacts the number of gaps that you are likely to see in the alignment (a larger alignment will lead to more gaps). We recommend entering 250 for sequences of length $O(10^5)$.