

Folding, but not your laundry

This problem is based on the processes and chemistry used by chains of amino acids to "fold" into proteins. The structure of the problem is from Project Euler Problem 300 [1], you should become familiar with their description before proceeding.

1) Design a parallel algorithm which, given an arbitrary string of H and P elements, finds the conformation with the maximum number of H+H bonds. Note that the algorithm can, although doesn't have to be, embarrassingly parallel; that is it can use the parallel HPC resource to examine each potential arrangement individually in a parameter space search approach rather than using an algorithm that is parallel for each particular conformation.

2) Implement that algorithm using MPI, OpenMP, Hadoop or any other library/framework which provides support for the parallelism you require. Hybrid approaches which combine two or more of those tools are also permitted.

3) Test and refine your implementation using user supplied data sets that conform to the input specifications found in [1].

4) Run your program using the test string provided below, your software's output should closely resemble this format (note, this is an example only):

Input string: HHPPHHHPHHPH

Number of H+H bonds in best conformation: 9

Graphical representation (ASCII) of the best conformation:

```
PHH
HHHP
HHH
PP
```

Turn-in a write-up which describes your approach in detail and your complete source code. Partial credit will be given for partial solutions, that is a reasonable description of an approach/algorithm with a partial implementation of the solution. Source code without a detailed explanation will not earn any credit.

5) Test data: HHHPHPPHHPHHPH

[1] <http://projecteuler.net/index.php?section=problems&id=300>