ECS223a Parallel Algorithms
Homework 1

You are encouraged to talk to other people about these problems, but please **write up the solutions by yourself**. Cite any conversations you had with others, as well as books, papers or Web sites you consulted.

Always explain the answer in **your own words**; do not copy text from the book, other books, Web sites, your friends' homework, etc. Explain your solution as you would to someone who does not understand it, for instance to a beginning graduate student or an advanced undergraduate. Do not give several solutions to one problem; pick the best one.

Please type your homework. If you know LaTeX, use that. If not, you may type your answers in any word processing system and write in mathematical notation by hand as necessary. Include pictures if appropriate; you can draw in pictures by hand or include them in the file.

1. The polynomial
$$f(x) = a_n x^n + a_{n-1} x^{n-1} + \ldots + a_1 x + a_0$$
can be represented by its coefficient vector a. Describe how to evaluate the polynomial at a specific number x using scan operations.

2. A matrix is *sparse* if most of its entries are zero; so for an $n \times n$ matrix A, the number m of non-zero entires is $<< n^2$. Assume $n = O(m)$ (although not necessarily $\theta(m)$). Describe how to represent A using $O(m)$ space, so that the matrix-vector multiplication Av can be computed using $O(m)$ work and $O(\lg m)$ depth on a CREW PRAM. Assume the column vector v is given as an array of length n and will contain $O(n)$ non-zeros. Can you come up with an EREW algorithm? Feel free to do whatever pre-processing you like with the matrix, so long as the representation you use in the multiplication algorithm is still size $O(m)$.

3. Do Exercise 2.44 at the end of Chapter 2. Here is an example of input and output, for $n = 16$ (so there are four "colors", 0-3):
$$\begin{aligned} \text{colors} &= [2,0,0,1,3,3,2,3,1,2,0,1,3,0,2,0] \\ \text{values} &= [1,1,2,2,5,1,2,1,4,3,3,2,0,2,3,1] \\ \text{output} &= [1,1,3,2,5,6,3,7,6,6,6,9,7,8,9,9] \end{aligned}$$

4. Do Exercise 2.45 at the end of Chapter 2.

5. In the Modern GPU Library, selection is done using an algorithm related to radix sort; feel free to check out at least the beginning of the Web page about this, at `http://www.moderngpu.com/select/radixselect.html`. Give a theoretical work-depth bound for this algorithm. Discuss in what sense it might be preferable to the algorithm we studied.