Project 4, Due November 21 at midnight

The code in sparse_matvec.c is a sequential version of a sparse matrix-vector multiply. The matrix is sparse in that many of its elements are zero. Rather than representing all of these zeros which wastes storage, the code uses a representation called Compressed Row Storage (CRS), which only represents the nonzeros with auxiliary data structures to keep track of their location in the full matrix.

I provide:

Sparse input matrices which were generated from the MatrixMarket (see http://math.nist.gov/MatrixMarket/). The format for these is a sorted coordinate representation (row, col, value) and will need to be converted to CRS.

An implementation of dense matvec in CUDA.

You write:

A CUDA implementation