

For any $0 \leq i < n$, every index $R[i] \leq k < R[i + 1]$ corresponds to a non-zero entry at row i and column $C[k]$ with value $V[k]$.

For instance, the matrix A used as an example above would be described in a C data structure as

```
int R[7] = { 0, 3, 5, 7, 9, 9, 10 };
int C[10] = { 0, 1, 2, 1, 2, 4, 5, 0, 2, 4 };
int V[10] = { 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 };
```

Write a method that on input n , R , C , V , and x outputs the product Ax . You may use the above matrix as an example.

- (4) In practice, it is convenient to assemble a matrix first in coordinate format and then convert it into compressed sparse row format for faster matrix multiplication. Write an algorithm that transforms a matrix in coordinate format into a matrix in compressed sparse row format.
- (5) A sparse matrix A arising in numerical partial differential equations may have k non-zero entries per row. Suppose that integers/indices are saved with 4 bytes each and that floating-point values occupy 8 bytes. If matrix has size $N \times N$, how much memory does it occupy for each of the previous three formats? Give an example with $N = 10^3, 10^6, 10^9$ and state with decimal (or binary) prefixes (Kilo, Mega, Tera, ...).

Remark 1

The sparse matrix formats require some overhead to describe the non-zero structure of the matrix. This overhead pays off in many applications because the number of non-zero entries is magnitudes smaller than the number of total entries. For example, for $N \times N$ matrices in finite element methods, the number of non-zero entries grows only linearly in N instead of quadratically.

Different sparse matrix formats have different advantages and disadvantages. The coordinate format occupies a bit more memory than the compressed sparse rows format, but it is rather easy to append (or erase) matrix entries. A matrix in compressed sparse rows format is less easy to modify, but the format is rather compact.

A major advantage of the compressed sparse row format is that the matrix-vector product can be parallelized over the rows, which gives a considerable speed up on many-core CPUs.