

Solving a system with a lower-triangular matrix

$$L \cdot x = b$$

Theory describes the algorithm as:

for $i = 1 \dots N$ // Foreach row

$x_i \leftarrow b_i$ // Start with right-hand side

(*)
for $j = 1 \dots (i-1)$

$L_{ij} x_j \leftarrow x_i - L_{ij} x_j$ // Subtract all lower-indexed x_j 's (move them to RHS)

$x_i \leftarrow x_i / L_{ii}$ // Divide by diagonal elements

→ we will next consider some useful variations of this "baseline" algorithm

→ "In-Place" Forward Substitution

Applicable when the solution to $Lx=b$ is allowed to "overwrite" the vector b .

Stated another way... let's assume that initially \underline{x} contains the right-hand side, and then is overwritten in-place by the solution. This is then written:

$$\begin{aligned} & \text{for } i = 1 \dots N \quad // \text{ Assume } \underline{x} \text{ contains } \underline{b} \\ & \quad \text{for } j = 1 \dots (i-1) \quad \text{at start} \\ & \quad \quad L_{ij} x_j \leftarrow x_j - L_{ij} x_j \\ & \quad \quad x_i \leftarrow x_i / L_{ii} \end{aligned}$$

Of course, the statement of these algorithms implicitly presumes that L is a dense triangular matrix. How would we express this with CSR matrices?

forward substitution - CSR version

(Note: We will switch to 0-based indexing to match our implementation)

The trick is to replace the loop marked by $(*)$ with a traversal of the sparse row!

for $i = 0 \dots (N-1)$

$x[i] \leftarrow b[i]$ \leftarrow skip this for in-place

for $k = \text{rowOffsets}[i] \dots \text{rowOffsets}[i+1] - 2$

$j \leftarrow \text{columnIndices}[k]$

$x[i] \leftarrow x[i] - \text{values}[k] * x[j]$

$x[i] \leftarrow x[i] / \text{values}[\text{offsets}[i+1] - 1]$

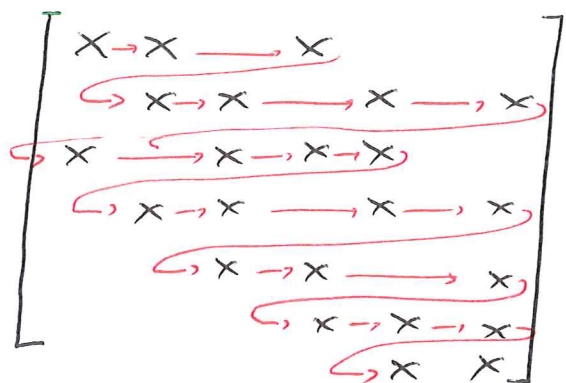
Stop 1-entry before diagonal!
This loop might be empty!

Notes:

→ Primary contributor to the cost: reading the matrix L

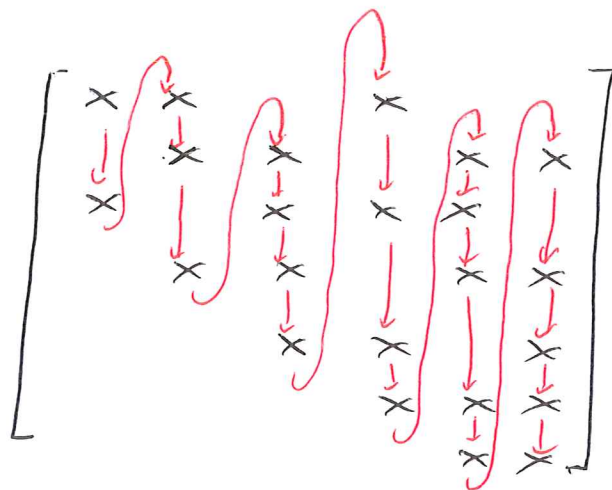
→ Algorithm "feels" very serial ... more on this later

Challenge: What if the matrix was given in the CSC (compressed sparse column) format instead?



CSR

row offsets
column Indices
values



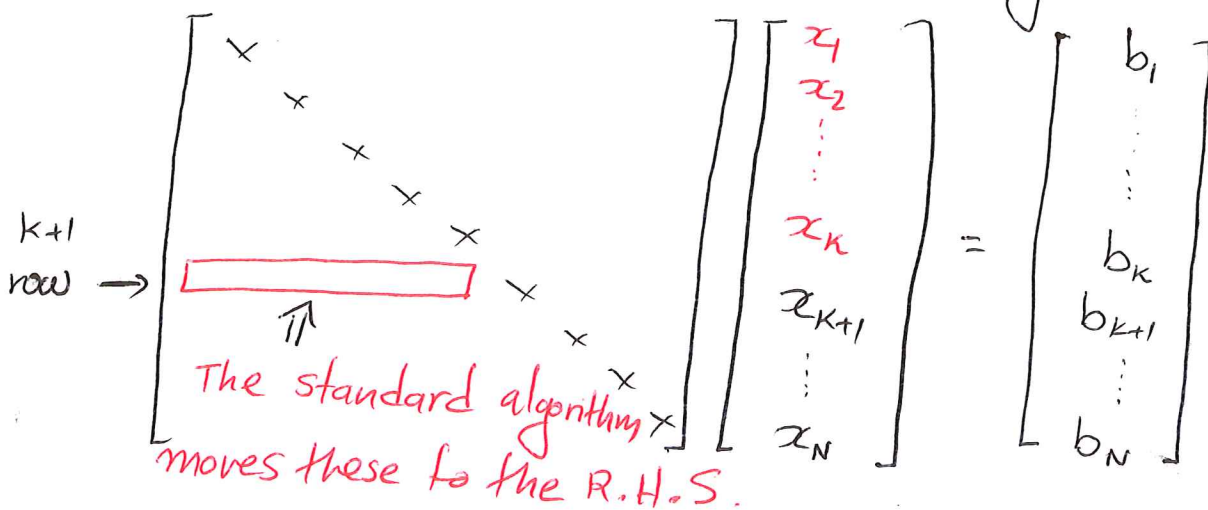
CSC

column offsets
row Indices
values


VS

Key insight for algorithm transformation:

At any stage of forward sub, we have some finalized x 's (x_1, x_2, \dots, x_k) and some yet-to compute.



An alternative view of the $(k+1)$ th step:



as soon as x_k is finalized, move all of these to the R.H.S. instead!

Mathematical Statement (pseudocode)

$x \leftarrow b$ // copy entire $b \rightarrow x$ (unless in-place)

for $j = 1 \dots N$

$x_j \leftarrow x_j / L_{jj}$

for $i = (j+1) \dots N$

$x_i \leftarrow x_i - L_{ij} x_j$

(This yields the final x_j !
I can do this because we eliminated all entries to the left of x_j , in the j -th equation!)

//
this eliminates the column shown here

CSC version - Sparse

CS639 - Feb 27

1 p. 6

$x \leftarrow b$ // copy entire array - skip for in-place

for $j = 0 \dots N-1$

$x[j] \leftarrow x[j] / \text{values}[\text{columnOffsets}[j+1] - 1]$

for $k = \text{columnOffsets}[j] + 1 \dots \text{columnOffsets}[j+1] - 1$
 \leftarrow skip diagonal

$i \leftarrow \text{rowIndices}[k]$

$x[i] \leftarrow x[i] - \text{values}[k] * x[j]$