

Performance Engineering

Case study: “Jacobi” stencil

The basics in two dimensions (2D)

Layer condition in 2D

From 2D to 3D

OpenMP parallelization strategies and layer condition in 3D

NT stores

Prof. Dr. G. Wellein^(a,b) , Dr. G. Hager^(a)

^(a) Erlangen National High Performance Computing Center (NHR@FAU)

^(b) Department für Informatik

Friedrich-Alexander-Universität Erlangen-Nürnberg, Sommersemester 2026



Stencil schemes


- Stencil schemes frequently occur in PDE solvers on regular lattice structures
- Basically it is a sparse matrix vector multiply (spMVM) embedded in an iterative scheme (outer loop)
- but the **regular access structure** allows for **matrix free coding**

```
do iter = 1, maxit
```

```
  Perform sweep over regular grid:  $y \leftarrow x$ 
```

```
  Swap  $y \leftrightarrow x$ 
```

```
enddo
```



- Complexity of implementation and performance depends on
 - update scheme, e.g. Jacobi-type, Gauss-Seidel-type,...
 - spatial (stencil) extent, e.g. 7-pt or 25-pt in 3D,...
 - ...

Case study: Jacobi stencil


The basics in two dimensions (2D)

Layer condition in 2D

From 2D to 3D

OpenMP parallelization strategies and layer condition in 3D

NT stores




Jacobi-type 5-pt stencil in 2D

```
do k=1,kmax
do j=1,jmax
  y(j,k) = const * ( x(j-1,k) + x(j+1,k) &
                    + x(j,k-1) + x(j,k+1) )
enddo
enddo
```

sweep

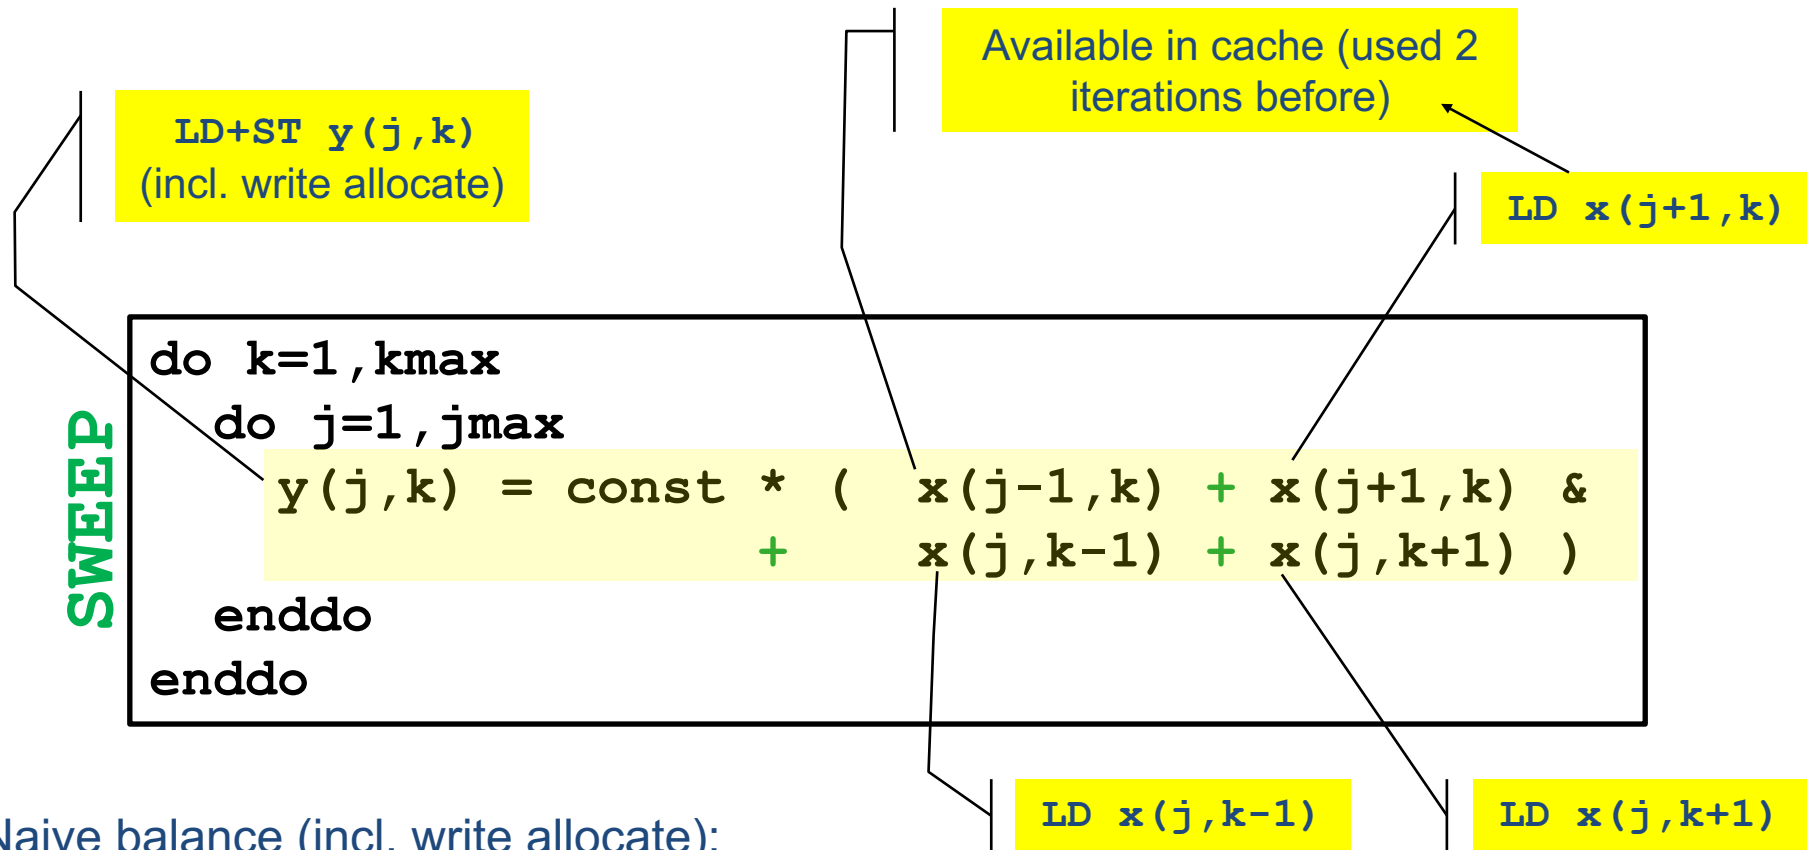
2 arrays

Lattice site update (LUP)



Appropriate performance metric:
“Lattice site updates per second” [LUP/s]
(here: Multiply by 4 FLOP/LUP to get FLOP/s rate)

Jacobi 5-pt stencil in 2D: data transfer analysis

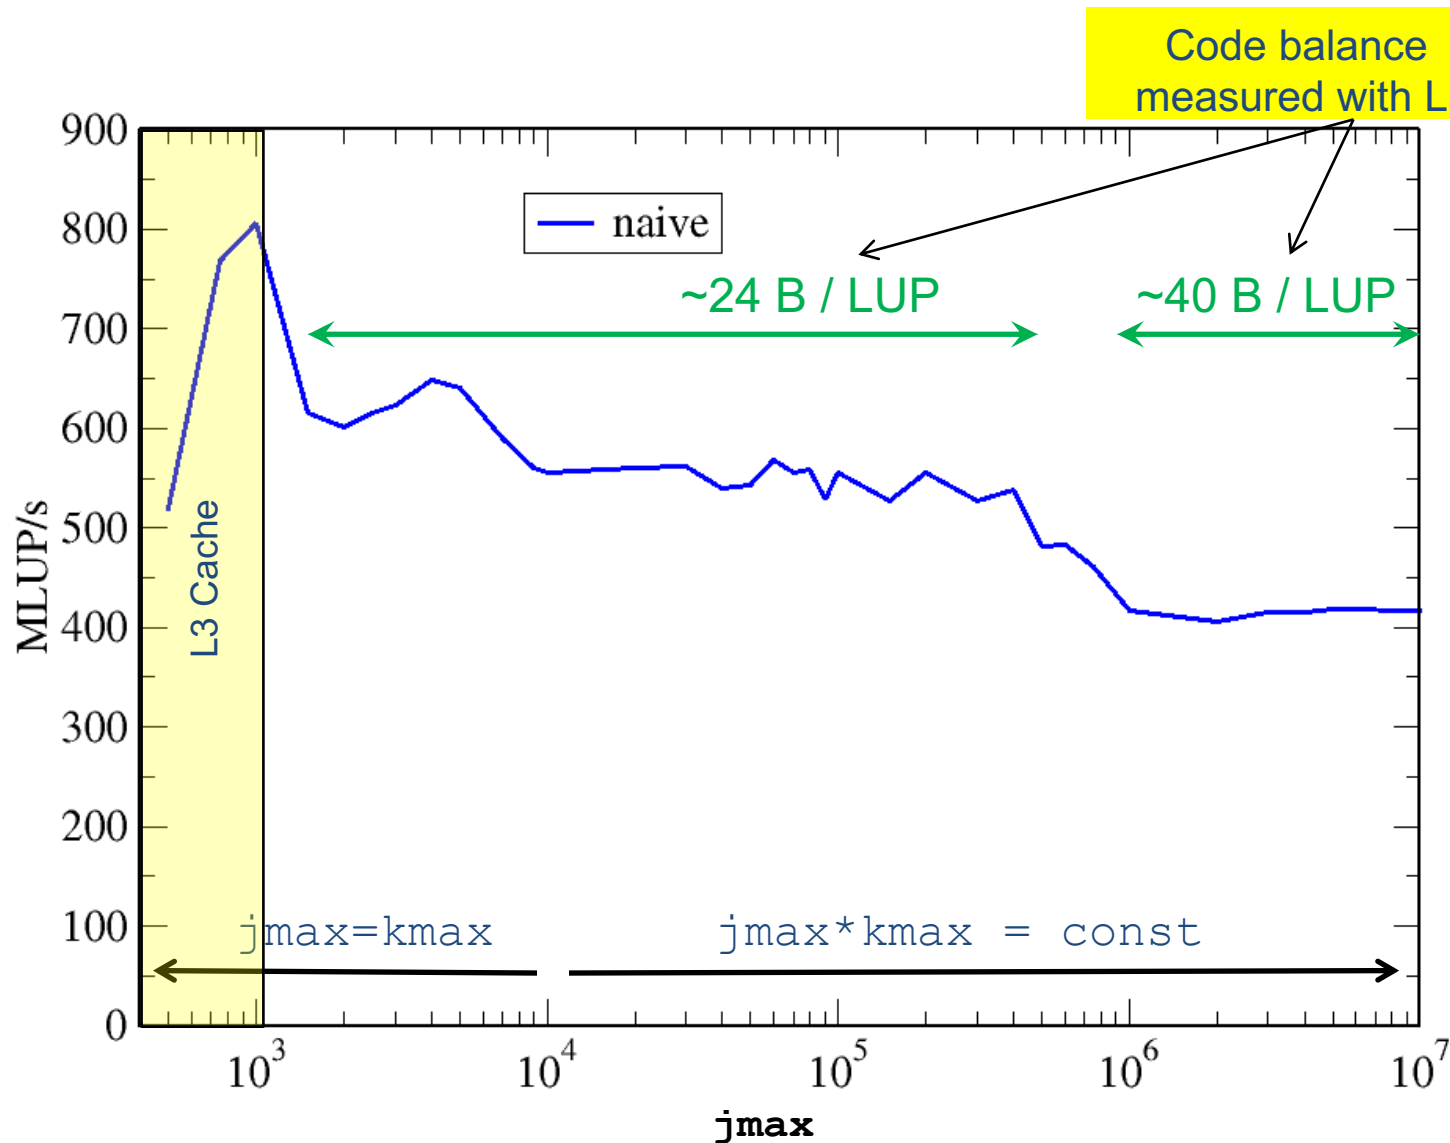


Naive balance (incl. write allocate):

$\mathbf{x}(\ :, \ :)$: 3 LD +
 $\mathbf{y}(\ :, \ :)$: 1 ST+ 1LD

→ $B_C = 5 \text{ Words} / \text{LUP}$ → $B_C = 40 \text{ B} / \text{LUP}$ (assuming double precision)

Jacobi 5-pt stencil in 2D: Single core performance



Questions:

1. How to achieve 24 B/LUP also for large j_{\max} ?
2. How to sustain $>600 \text{ MLUP/s}$ for $j_{\max} > 10^4$?
3. Why 24 B/LUP anyway???

Intel Compiler: ifort V13.1

Intel Xeon E5-2690 v2
("IvyBridge"@3 GHz)

Case study: Jacobi stencil


The basics in two dimensions (2D)

Layer condition in 2D

From 2D to 3D

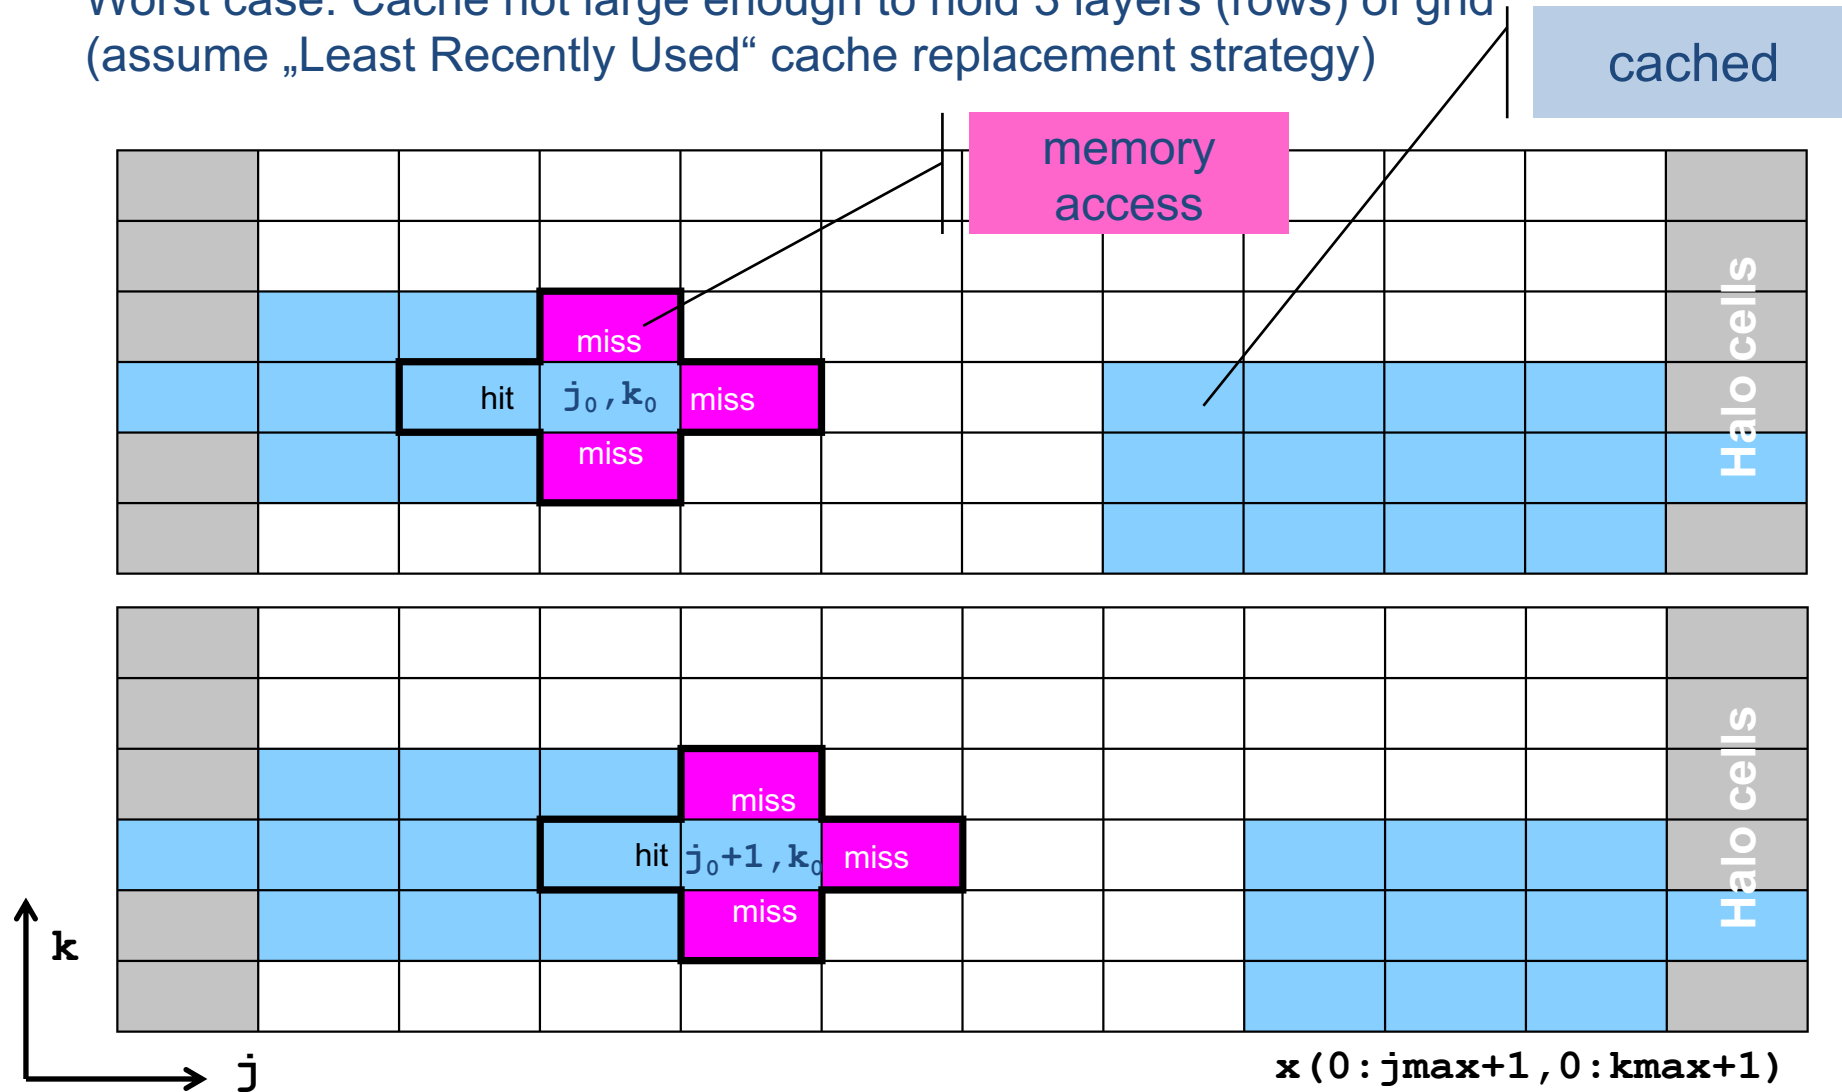
OpenMP parallelization strategies and layer condition in 3D

NT stores




Analyzing the data flow

Worst case: Cache not large enough to hold 3 layers (rows) of grid
(assume „Least Recently Used“ cache replacement strategy)



Analyzing the data flow

Worst case: Cache not large enough to hold 3 layers (rows) of grid
(+assume „Least Recently Used“ replacement strategy)




Analyzing the data flow



Reduce inner (j-) loop dimension successively



Best case: 3 „layers“ of grid fit into the cache!



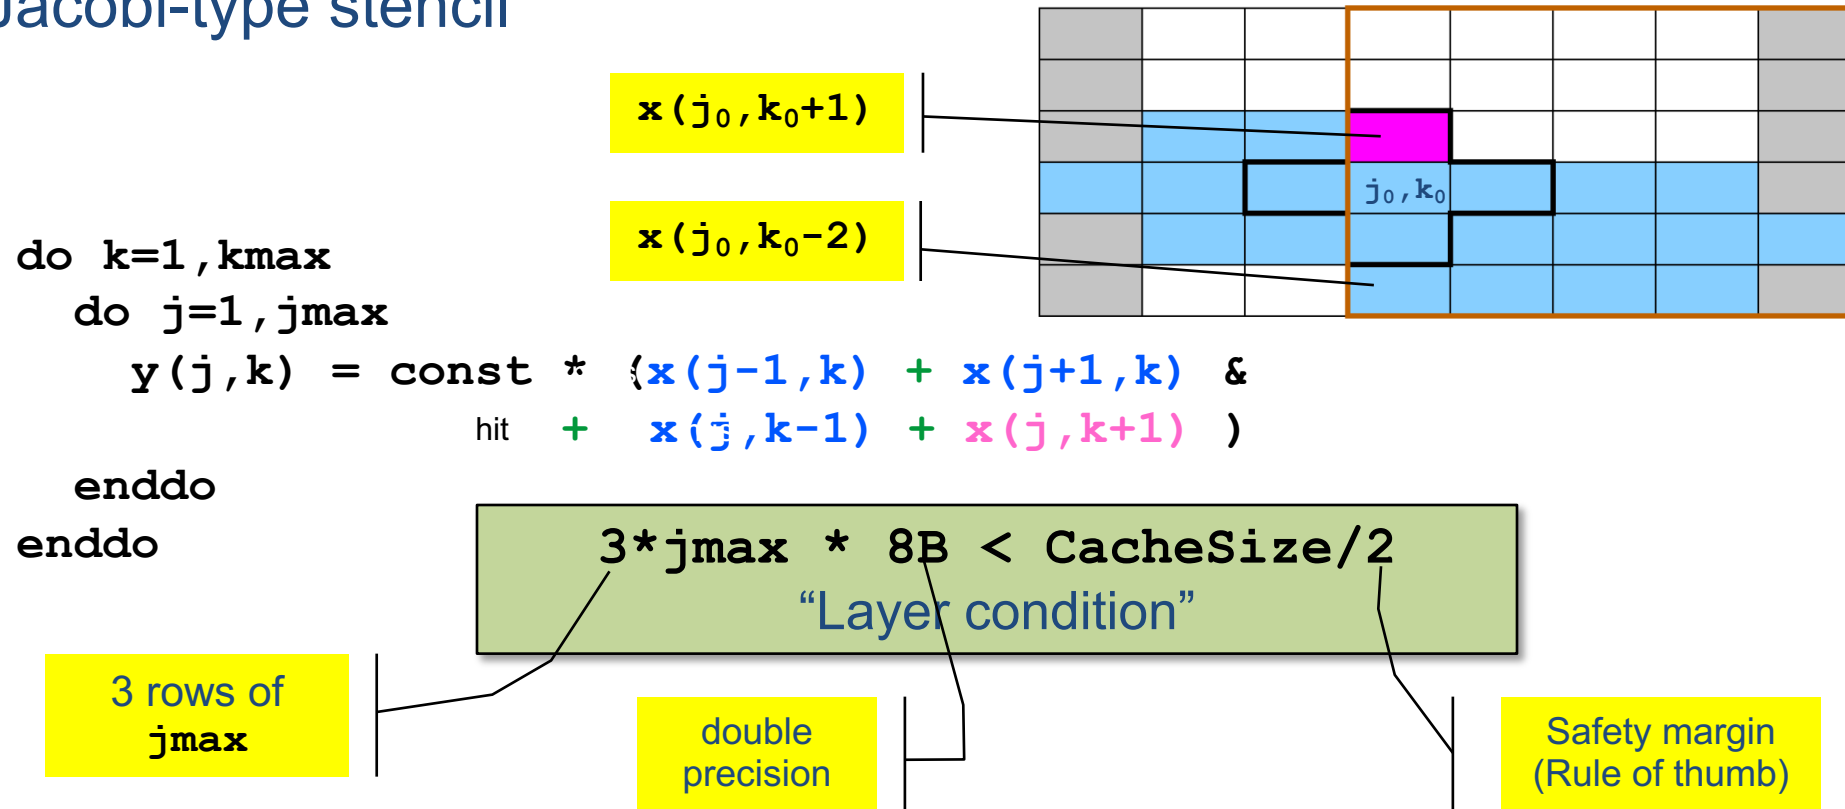
$x(0:j_{\max}1+1, 0:k_{\max}+1)$



$x(0:j_{\max}2+1, 0:k_{\max}+1)$

Analyzing the data flow: Layer condition

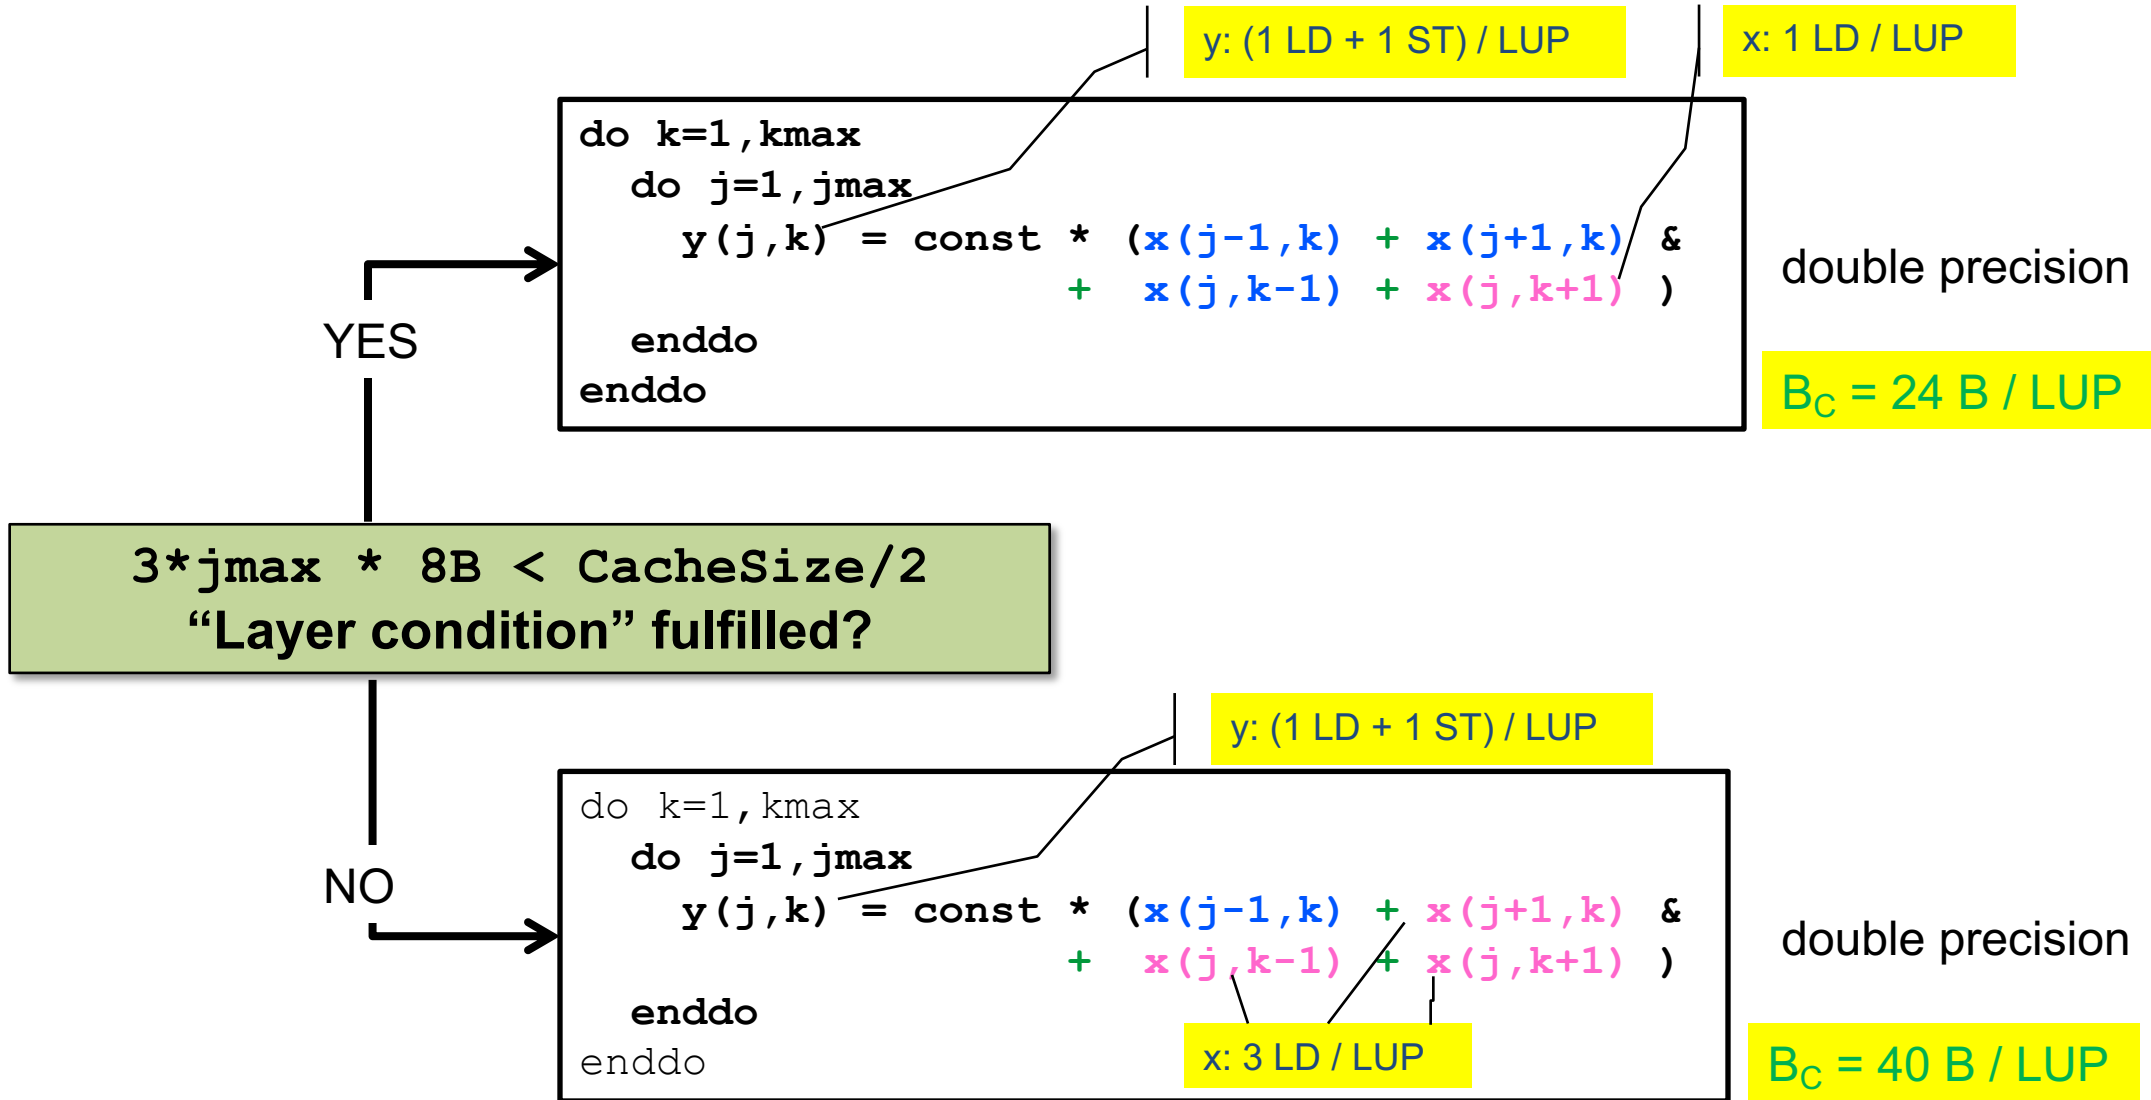
2D 5-pt Jacobi-type stencil



Layer condition:

- No impact of outer loop length (k_{\max})
- No strict guideline (cache associativity – data traffic for $y()$ not included)
- Needs to be adapted for other stencils, e.g., in 3D, long-range, multi-array,...

Analyzing the data flow: Layer condition (2D 5-pt Jacobi)



Analyzing the data flow: Layer condition (2D 5-pt Jacobi)

- How to establish the layer condition for all domain sizes ?
- Idea: **Spatial blocking**
 - Reuse elements of $\mathbf{x}()$ as long as they stay in cache
 - Sweep can be executed in any order, e.g., compute blocks in j-direction

→ “Spatial Blocking” of j-loop:

```
do jb=1, jmax, jbblock !           Assume jmax is multiple of jbblock
  do k=1, kmax
    do j= jb, (jb+jbblock-1) ! Length of inner loop: jbblock
      y(j,k) = const * (x(j-1,k) + x(j+1,k) &
                          + x(j,k-1) + x(j,k+1) )
    enddo
  enddo
enddo
```

New layer condition (blocking)

$$3 * \mathbf{jbblock} * 8B < \text{CacheSize} / 2$$

→ Determine for given **CacheSize** an appropriate **jbblock** value:

$$\mathbf{jbblock} < \text{CacheSize} / 48 B$$