

Erlangen Regional  
Computing Center



FRIEDRICH-ALEXANDER  
UNIVERSITÄT  
ERLANGEN-NÜRNBERG

# Case study: A Jacobi smoother

The basics in two dimensions

Layer conditions

Optimization by spatial blocking




- Stencil schemes frequently occur in PDE solvers on regular lattice structures
- The **regular access structure** allows for **matrix-free coding**

```
do iter = 1, max_iterations
```

```
    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 extent, e.g. 7-pt or 25-pt in 3D,...

# Jacobi-type 5-pt stencil in 2D

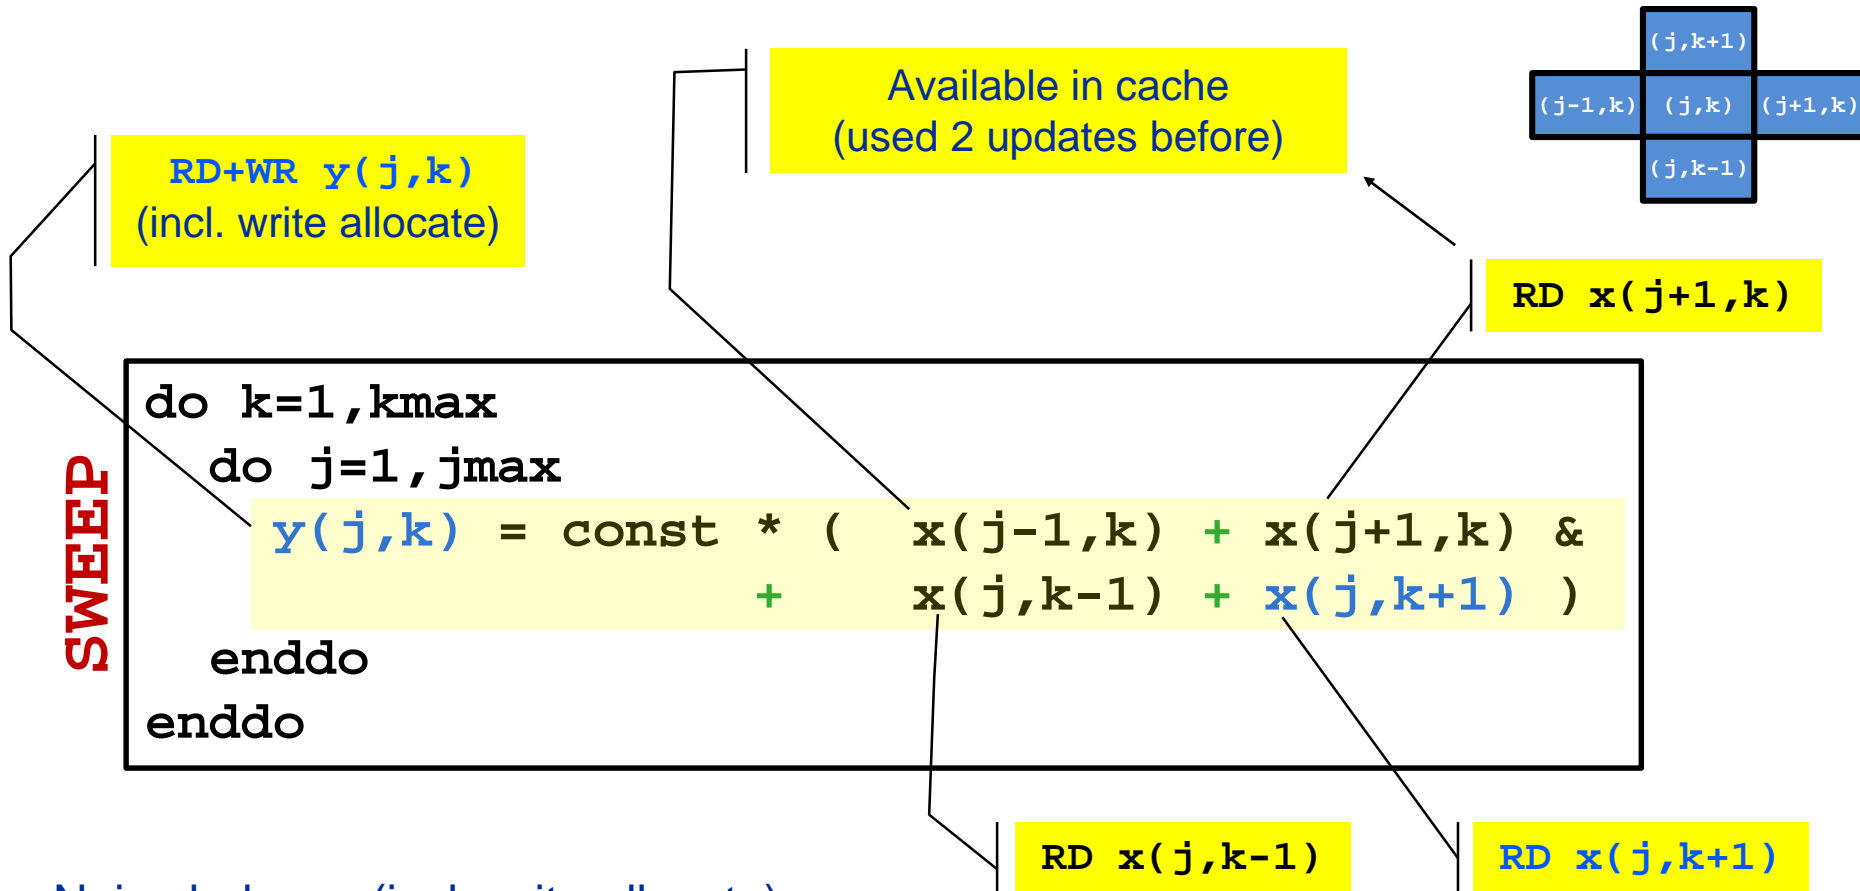
sweep

```
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
```

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)



Naive balance (incl. write allocate):

$x(:, :)$  : 3 RD +

$y(:, :)$  : 1 WR + 1 RD

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

# Case study: A Jacobi smoother

The basics in two dimensions

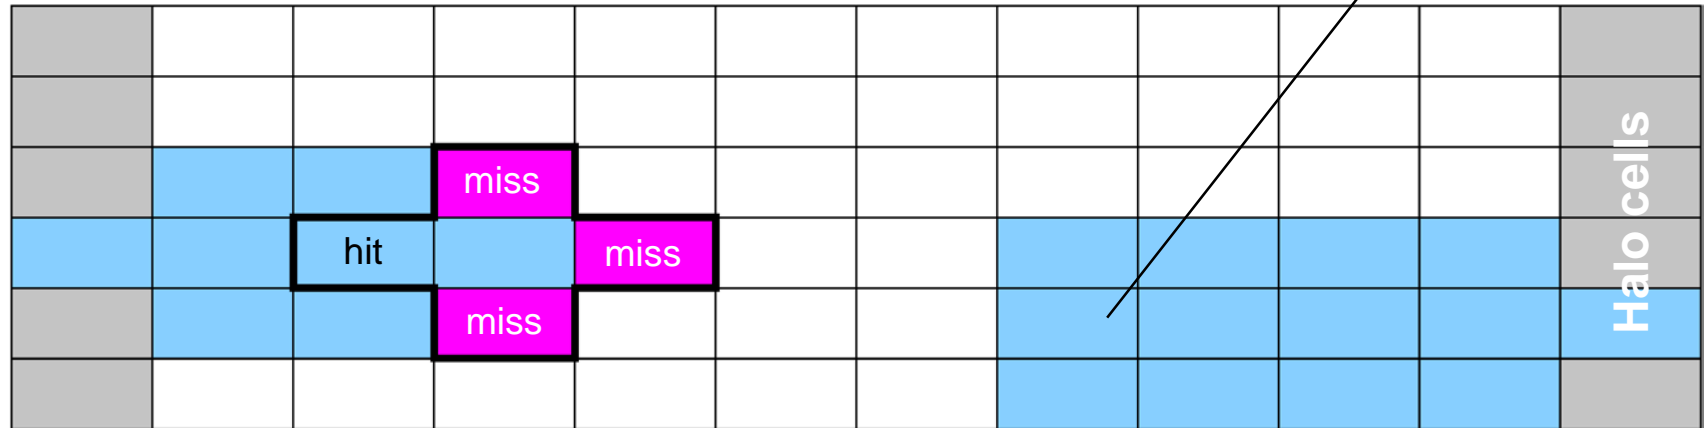
Layer conditions

Optimization by spatial blocking



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

cached

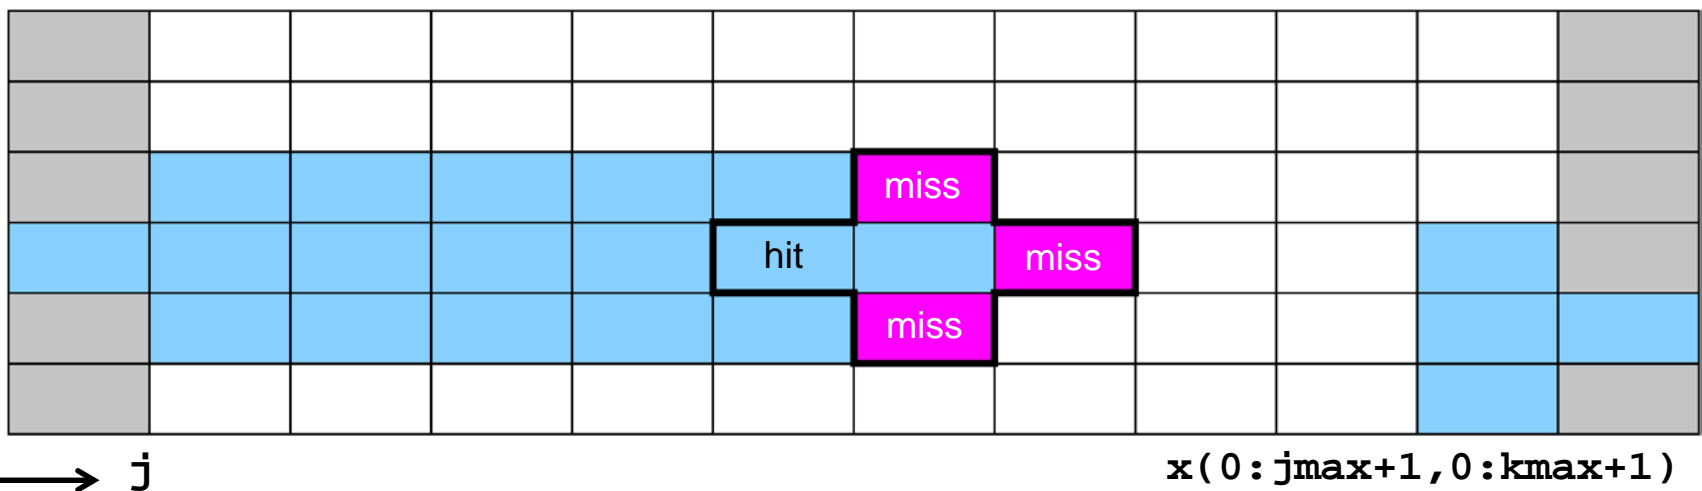



k

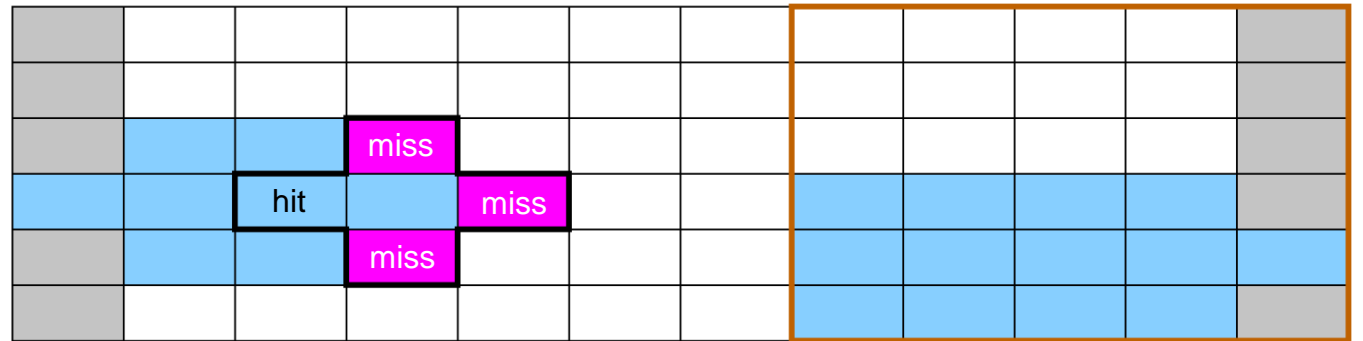
j

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

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





Reduce inner (j-) loop dimension successively



`x(0:jmax1+1, 0:kmax+1)`




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



`x(0:jmax2+1, 0:kmax+1)`





```
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
```

$$3 * j_{\max} * 8B < \text{CacheSize} / 2$$

“Layer condition”

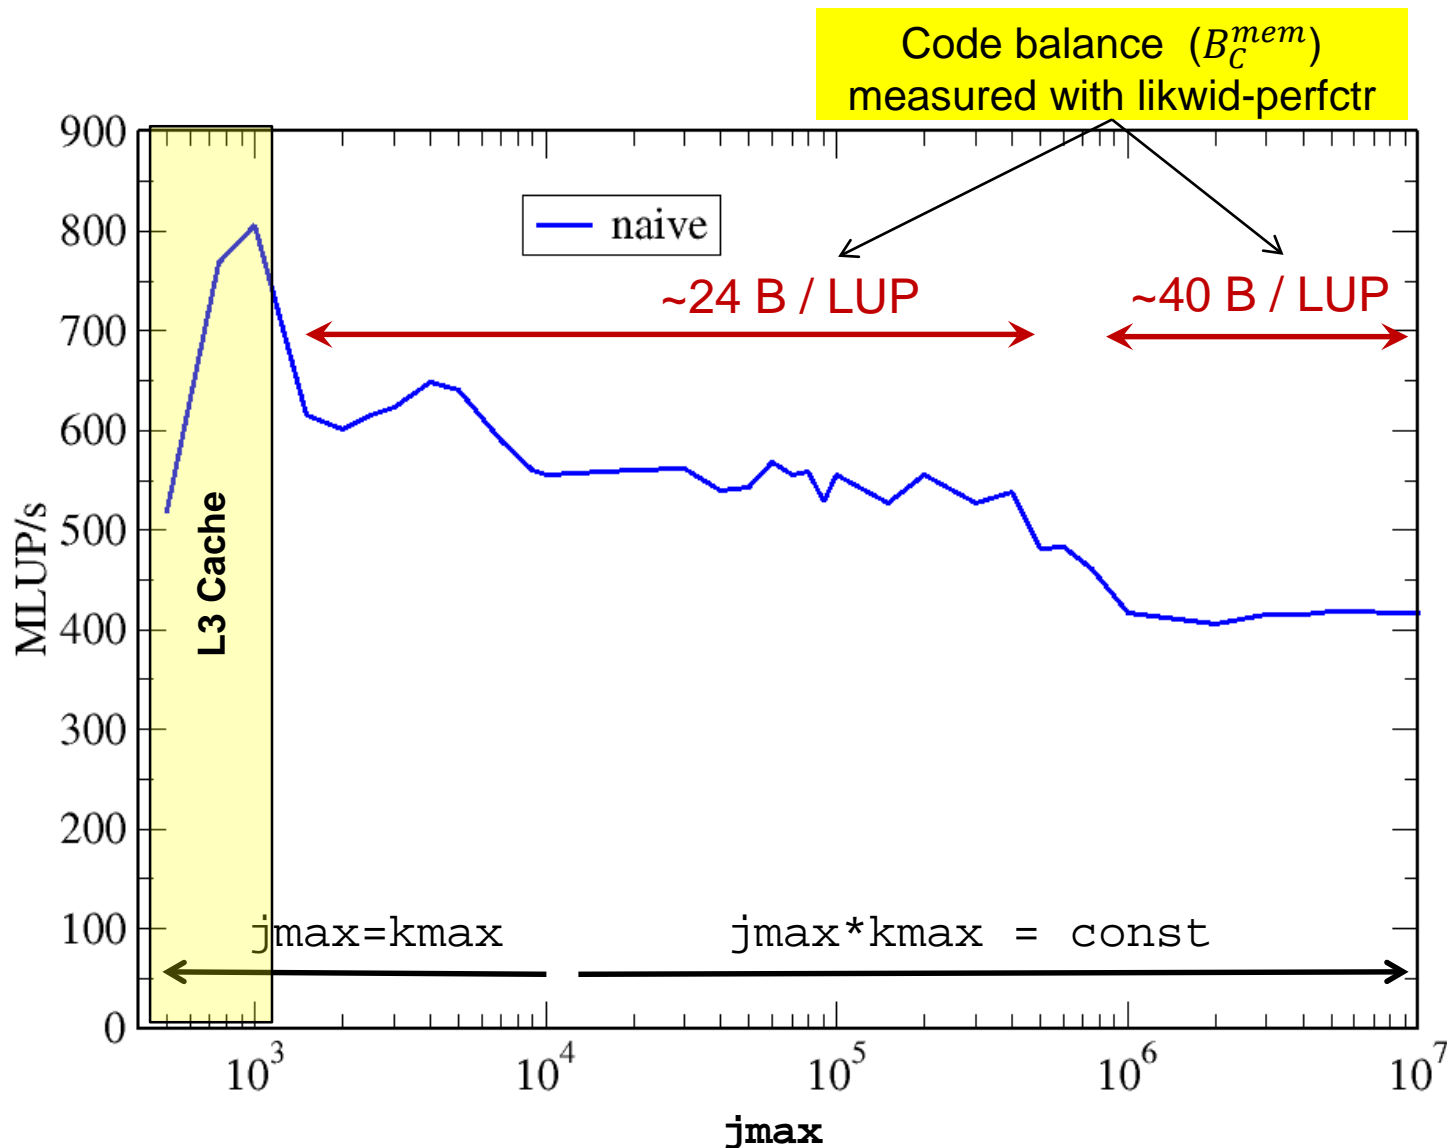
3 rows of  
 $j_{\max}$

double  
precision

Safety margin  
(Rule of thumb)

## Layer condition:


- Does not depend on 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., 3D 7-pt stencil)



Questions:

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

Intel Compiler  
ifort V13.1  
Intel Xeon E5-2690 v2  
("IvyBridge"@3 GHz)



Erlangen Regional  
Computing Center



# Case study: A Jacobi smoother

The basics in two dimensions

Layer conditions

Optimization by spatial blocking




## Why 24 byte/LUP?

- High-level view of sweep: read  $\mathbf{x}(\cdot)$ , write  $\mathbf{y}(\cdot)$   
→ if maximum reuse with sweep is possible, 24 byte/LUP should be achievable

But how to establish the layer condition for all domain sizes?


- Idea: **Spatial blocking**
  - Reuse elements of  $\mathbf{x}(\cdot)$  as long as they stay in cache
  - Main idea: Order of site updates does not matter  
→ “reduce inner dimension” by cutting the inner loop short



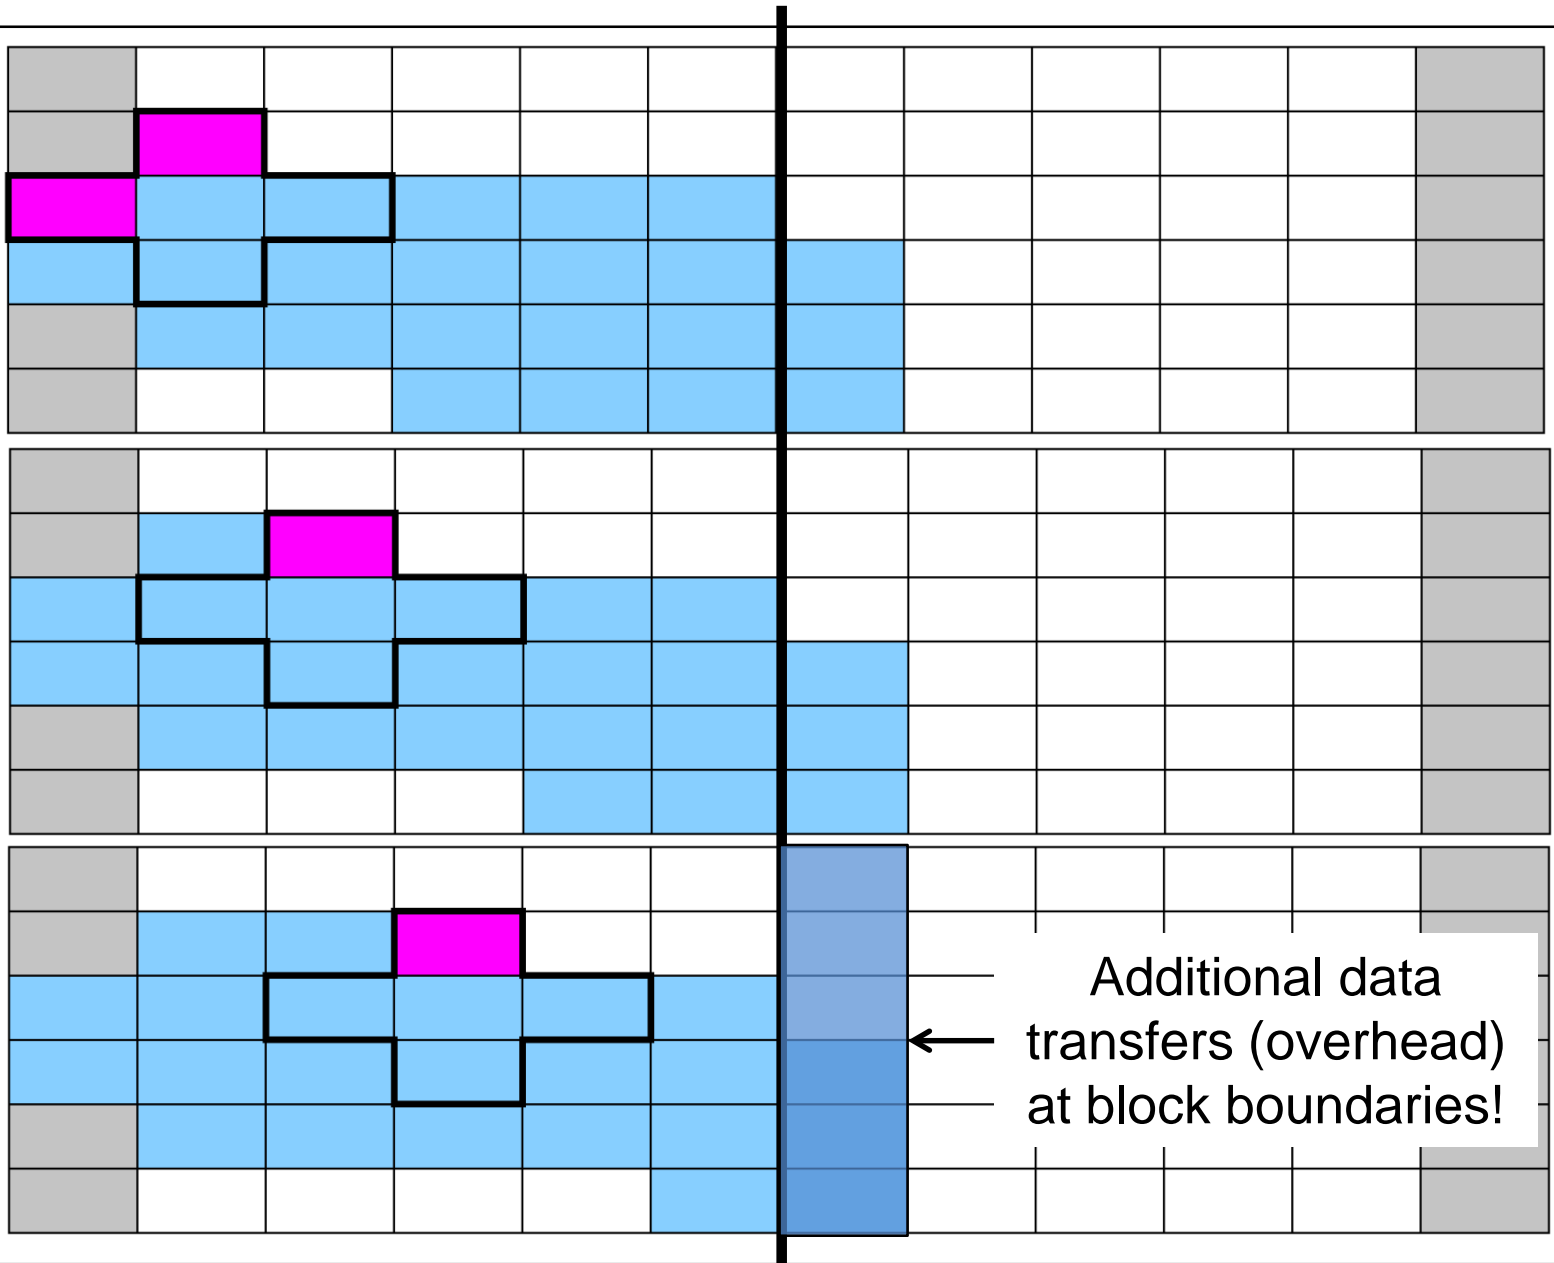
# Idea: Enable data reuse by blocking!

Split  
domain into  
subblocks:

e.g. block  
size = 5



# Establish the layer condition by blocking



→ “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
```

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

**New layer condition (blocking)**


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

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

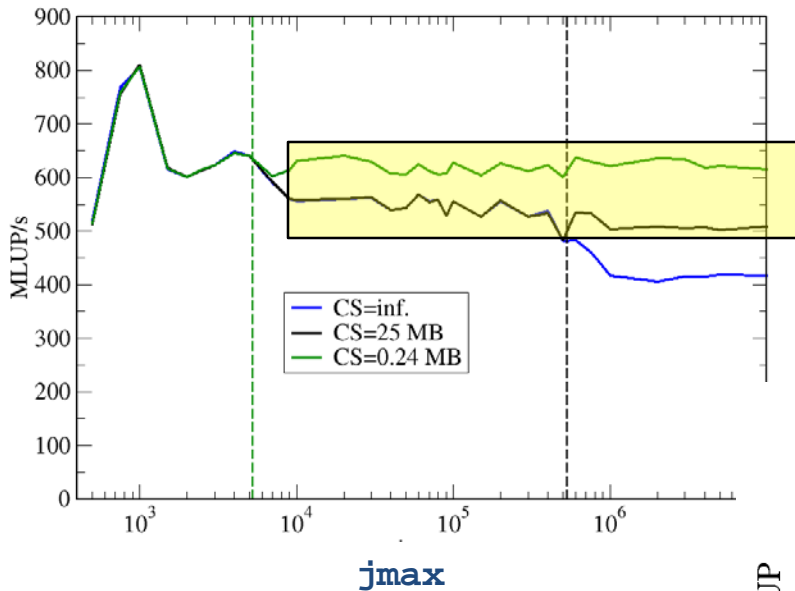
# Establish layer condition by spatial blocking

$$jblock < CacheSize / 48 B$$

Which cache to block for?

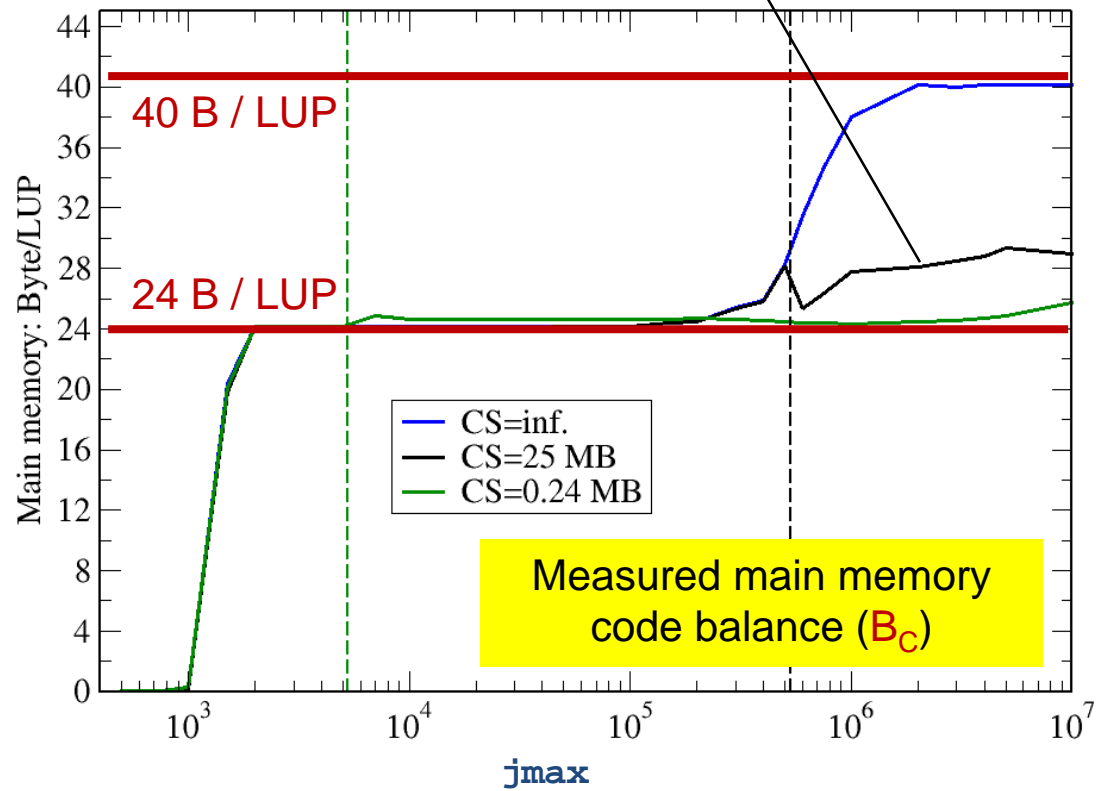






Main memory access is not reason for different performance (but L3 access is!)

Blocking factor (CS=25 MB) still a little too large




Measured main memory code balance ( $B_C$ )

Intel Compiler  
ifort V13.1  
Intel Xeon E5-2690 v2  
("IvyBridge"@3 GHz)

## Straightforward OpenMP work sharing:

```
do jb=1, jmax, jblock
!$OMP PARALLEL DO SCHEDULE(static)
do k=1, kmax
do j= jb, min(jb+jblock-1, jmax)
y(j,k) = const * (x(j-1,k) + x(j+1,k) &
+ x(j,k-1) + x(j,k+1) )
enddo
enddo
!$OMP END PARALLEL DO
enddo
```



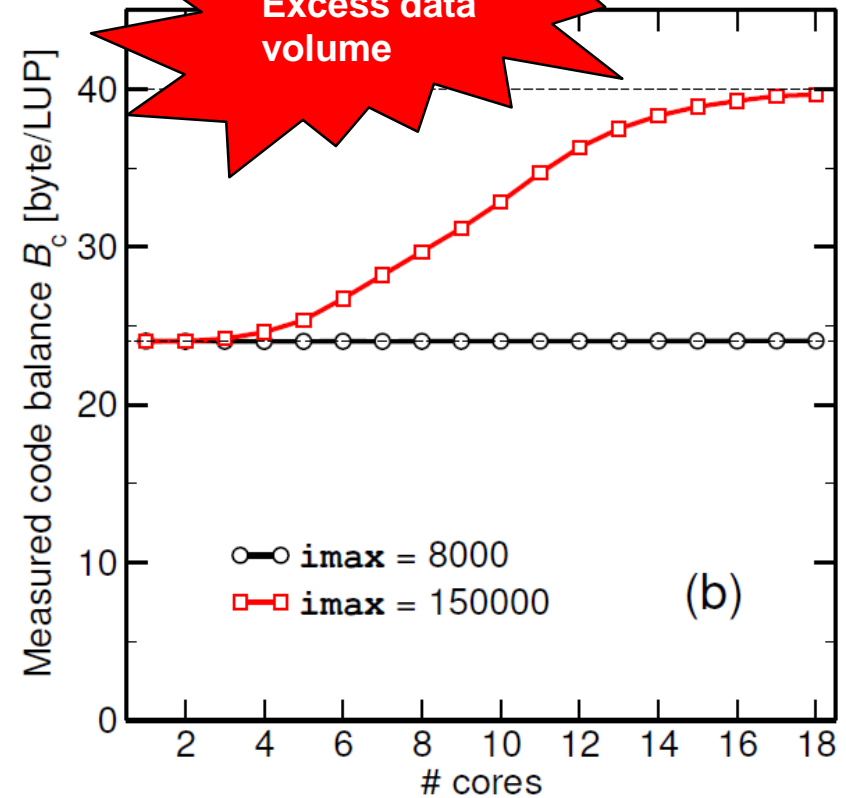

- Caveat: LC must be fulfilled per thread → shared cache causes smaller blocks!


**Layer condition:**

$$3 * jblock * 8B < CS_t / 2$$


Cache size available  
per thread

Example: 2D 5-point stencil on Intel Xeon Broadwell 18-core (non-CoD), 45 MiB of shared L3 cache







(a)



(b)



(c)



(d)

- a) Long-range  $r = 2$ : 5 layers ( $2r + 1$ )
- b) Long-range  $r = 2$  with gaps: 6 layers (2 per populated row)
- c) Asymmetric: 3 layers
- d) 2D box: 3 layers

- We have **made sense** of the memory-bound **performance** vs. problem size
  - “**Layer conditions**” lead to **predictions of code balance**
  - “**What part of the data comes from where**” is a crucial question
  - The model works only if the **bandwidth is “saturated”**
    - In-cache modeling is more involved
- **Avoiding slow data paths** == re-establishing the most favorable layer condition
- Improved code showed the **speedup predicted** by the model
- Optimal **blocking factor can be estimated**
  - Be guided by the cache size the **layer condition**
  - No need for exhaustive scan of “optimization space”
- **Food for thought**
  - Multi-dimensional loop blocking – would it make sense?
  - Can we choose a “better” OpenMP loop schedule?
  - What would change if we parallelized inner loops?