

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



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



From 2D to 3D


- 2D



- Towards 3D understanding
- Picture can be considered as 2D cut of 3D domain for (new) fixed i -coordinate:

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

From 2D to 3D



- $\mathbf{x}(0:\mathbf{i}\mathbf{max}+1, 0:\mathbf{j}\mathbf{max}+1, 0:\mathbf{k}\mathbf{max}+1)$ – Assume $\mathbf{i}$ -direction contiguous in main memory (Fortran notation)
- Stay at 2D picture and consider one cell of $\mathbf{j}$ - $\mathbf{k}$ plane as a contiguous slab of elements in $\mathbf{i}$ -direction: $\mathbf{x}(0:\mathbf{i}\mathbf{max}, \mathbf{j}, \mathbf{k})$

Layer condition: From 2D 5-pt to 3D 7-pt Jacobi-type stencil

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

```
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
```

double precision

$$B_C = 24 B / \text{LUP}$$


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

double precision

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

$$B_C = 24 B / \text{LUP}$$

3D 7-pt Jacobi-type Stencil (sequential)



Question: Does parallelization/multi-threading change the layer condition?

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 Stencil – OpenMP parallelization (I)

```
!$OMP PARALLEL DO SCHEDULE(STATIC)
do k=1,kmax
  do j=1,jmax
    do i=1,imax
      y(i,j,k) = 1/6. * (x(i-1,j,k)      +x(i+1,j,k) &
                        + x(i,j-1,k)      +x(i,j+1,k)
                        + x(i,j,k-1)      +x(i,j,k+1) )
    enddo
  enddo
enddo
```

Basic guideline:
Parallelize outermost loop

SCHEDULE(STATIC)
Equally large chunks in k-direction

Jacobi Stencil – OpenMP parallelization (I)

Equally large chunks
per thread in k-direction



“Layer condition” to be
fulfilled by each thread

“Layer condition”:

$$n_{\text{threads}} * 3 * j_{\text{max}} * i_{\text{max}} * 8B < CS/2$$

Example:


- STATIC schedule; $k=12$
- 4 threads $\rightarrow 4 * 3$ Layers

Thread 0

Thread 1

Thread 2

Thread 3



Note: Threads do not need to be in sync –
they only sync at the end of the sweep

Jacobi Stencil – OpenMP parallelization (I)

Impact of parallelization on **CacheSize**: core-local vs. shared caches

- CS_t is the available cache size for holding data when using **nthreads** (=cores)

“Layer condition” (LC): $nthreads * 3 * jmax * imax * 8B < CS_t / 2$

- If CS is **shared cache** (e.g. L3): $CS_t = \text{constant}$
→ LC more stringent at higher thread counts

L3-CacheSize=25 MB

1 thread: $imax=jmax < 720$

10 threads: $imax=jmax < 230$

- If CS is **core-local cache** (e.g. L2): $CS_t = \text{constant} * nthreads$
→ LC independent of thread count

L2-CacheSize=256 KB

1 thread: $imax=jmax < 70$

10 threads: $imax=jmax < 70$

Jacobi Stencil – OpenMP parallelization (I)

“Layer condition”: $nthreads * 3 * jmax * imax * 8B < CS_t / 2$

```
!$OMP PARALLEL DO SCHEDULE (STATIC)
```

```
do k=1,kmax
```

```
do j=1,jmax
```

```
do i=1,imax
```

```
  y(i,j,k) = 1/6. * (x(i-1,j,k) + x(i+1,j,k) &  
                    + x(i,j-1,k) + x(i,j+1,k) &  
                    + x(i,j,k-1) + x(i,j,k+1) )
```

```
enddo
```

```
enddo
```

```
enddo
```

double precision


$$B_C = 24 B / LUP$$

Intel® Xeon® Processor E5-2690 v2
10 cores@3 GHz
CacheSize = 25 MB (L3; shared)
 $b_S = 48 \text{ GB/s}$

Roofline model:
 $P = b_S / B_C^{mem}$


Best performance:
 $P = 2000 \text{ MLUP/s}$

Jacobi Stencil – OpenMP parallelization (I)



10 threads: performance drops around **imax=230**

1 thread: Layer condition OK – but can not saturate bandwidth



Validation: Measured data traffic from main memory [Bytes/LUP]

Jacobi Stencil – OpenMP parallelization (I)

Original “Layer condition” does not hold: $\text{nthreads} * 3 * j_{\text{max}} * i_{\text{max}} * 8B > CS_t / 2$

```
!$OMP PARALLEL DO SCHEDULE (STATIC)
do k=1,kmax
  do j=1,jmax
    do i=1,imax
      y(i,j,k) = 1/6.      * (x(i-1,j,k)      +x(i+1,j,k) &
                          + x(i,j-1,k)      +x(i,j+1,k)
                          + x(i,j,k-1)      +x(i,j,k+1) )
    enddo
  enddo
enddo
```

But assume: $\text{nthreads} * 3 * i_{\text{max}} * 8B < CS / 2$

(8+8) B/LUP for $y()$ (ST+WA)
+ 8 B/LUP for $x(i,j,k+1)$
+ 8 B/LUP for $x(i,j+1,k)$
+ 8 B/LUP for $x(i,j,k-1)$
→ $B_c = 40 \text{ B/LUP}$

→ $P = 1200 \text{ MLUP/s}$

Jacobi Stencil – OpenMP & spatial blocking

```
do jb=1,jmax,jblock ! Assume jmax is multiple of jblock
```

```
!$OMP PARALLEL DO SCHEDULE(STATIC)
```

```
do k=1,kmax
```

```
do j=jb, (jb+jblock-1) ! Loop length jblock
```

```
do i=1,imax
```

```
  y(i,j,k) = 1/6. * (x(i-1,j,k) +x(i+1,j,k) &  
                    + x(i,j-1,k) +x(i,j+1,k)  
                    + x(i,j,k-1) +x(i,j,k+1))
```

```
enddo
```

```
enddo
```

```
enddo
```

```
enddo
```


“Layer condition” (j-Blocking))
$$nthreads * 3 * jblock * imax * 8B < CS_t / 2$$

Ensure layer condition by choosing **jblock** appropriately (Cubic Domains):

$$jblock < CS_t / (imax * nthreads * 48B)$$

→ Back to prediction of $P = 2000 \text{ MLUP/s}$

Jacobi Stencil – OpenMP & spatial blocking



Impact of blocking factor **jblock**


Layer condition estimates appropriate **jblock**:

$$\mathbf{jblock} < CS_t / (2 * nthreads * 3 * imax * 8B)$$

$CS_t = 25$ MB
 $nthreads = 10$
 $imax = 400$

$$\mathbf{jblock} < 130$$

jblock=32,...,64
useful choices




Impact of **jblock**: Data traffic from L3 & main memory


Layer condition estimates appropriate **jblock**:

$$\mathbf{jblock} < CS_t / (2 * nthreads * 3 * imax * 8B)$$

$$\mathbf{jblock} < 130$$



Layer condition a bit too optimistic



Jacobi Stencil – what about j-loop parallelization

```
!$OMP PARALLEL PRIVATE(k)
```

```
do k=1,kmax
```

```
!$OMP DO SCHEDULE(STATIC)
```

```
do j=1,jmax
```

```
do i=1,imax
```

```
  y(i,j,k) = 1/6. * (x(i-1,j,k) +x(i+1,j,k) &  
                    + x(i,j-1,k) +x(i,j+1,k) +x(i,j,k-1) +x(i,j,k+1) )
```

```
enddo
```

```
enddo
```

```
!$OMP END DO
```

```
enddo
```

```
!$OMP END PARALLEL
```

All threads work on the same 3 k-layers at the same time


“Layer condition” (j-loop parallel; schedule=static)
 $3*jmax*imax*8B < CS_t/2$

Layer condition is independent of number of threads ($imax < 720$ for test system)


Impact of parallelization overhead?!

- 1 barrier each $jmax * imax$ LUPs
- Costs / barrier: approx. 1 500 cycles
- Computational time / barrier: $(jmax * imax \text{ LUPs}) / \text{maxMLUPs}$
e.g. $jmax=imax=200$ & $\text{maxMLUPs}=2000$ MLUPs $\rightarrow 20 * 10^{-6} \text{ s}$
 $\rightarrow 60\,000$ cycles

Jacobi Stencil – what about j-loop parallelization



Performance will drop at domain sizes ~500,...,700



Validation: Measured data traffic from main memory [Bytes/LUP]

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 Stencil – can we further improve?



Jacobi Stencil – Blocking + NT-stores



Conclusions from the Jacobi example

- Data transfer analysis is key to structured performance modeling and optimization
 - Avoiding slow data paths == re-establishing the most favorable layer condition
 - Optimal blocking factor for spatial blocking can be estimated
 - Be guided by the cache size the layer condition → does not depend on outer loop!
 - Blocking: overhead at block boundaries → Largest cache first choice
 - No need for exhaustive scan of “optimization space”
 - Performance optimization by re-establishing the layer condition was guided and validated by (roofline-like) model
 - Modeling and optimizing “Jacobi” stencil is blueprint for many stencil(-like) kernels.
-

Open topics – layer condition: 3D (outer loop parallel)

Derived for
OMP_SCHEDULE=STATIC
→ Does it change for other
scheduling strategies?

$$\text{nthreads} * 3 * \underbrace{\text{jmax} * \text{imax}}_{\text{}} * 8B < CS_t / 2$$

What about “square blocking”:
 $\text{jmax} \rightarrow \text{jblock} \ \&$
 $\text{imax} \rightarrow \text{iblock} \ ?$

What about i-blocking:
 $\text{imax} \rightarrow \text{iblock} \ ?$

OpenMP parallelization and blocking for a shared cache


Layer conditions make for interesting effects

- Less and less shared cache available per thread as #threads goes up
- LC may break “along the way”

■ Solutions


1. Choose small enough block or domain size
 - Layers either small enough to fit in core-private caches *or*
 - Shared cache big enough to hold all layers for all threads
2. Adaptive blocking for shared cache:

$$jblock = \frac{C}{\#threads \times 48 \text{ byte}}$$



Open topics – layer condition & outer loop parallel w/ **STATIC,1** schedule

```
!$OMP PARALLEL DO SCHEDULE(STATIC,1)
do k=1,kmax
  do j=1,jmax
    do i=1,imax
      y(i,j,k) = 1/6. * (x(i-1,j,k)    +x(i+1,j,k) &
                        + x(i,j-1,k)    +x(i,j+1,k)
                        + x(i,j,k-1)    +x(i,j,k+1) )
    enddo
  enddo
enddo
```



What is the relevant layer condition here???


Case study: Parallelizing a Gauss-Seidel Solver



3D matrix-free Gauss-Seidel smoother

- Matrix-free iterative solver for $Ax = b$

$$x_i^{(k)} = \frac{1}{a_{ii}} \left(- \sum_{j=1}^{i-1} a_{ij} x_j^{(k)} - \sum_{j=i+1}^n a_{ij} x_j^{(k-1)} + b_i \right)$$



- Here used for Dirichlet boundary value (PDE) problem $\Delta x = 0$


```
for(it=0; it<itmax; ++it) { // or convergence check
    for(k=1; k<kmax-1; ++k) {
        for(j=1; j<jmax-1; ++j) {
            for(i=1; i<imax-1; ++i) {
                x[k][j][i] = ( x[k][j][i-1] + x[k][j][i+1]
                    + x[k][j-1][i] + x[k][j+1][i]
                    + x[k-1][j][i] + x[k+1][j][i] ) / 6.0;
            }
        }
    }
}
```

“new data” “old data”

OpenMP parallelization?

- Naïve OpenMP loop parallelization impossible due to loop-carried dependency on all spatial loop levels


```
for (t=0; t<itmax; ++t) {  
    for (k=1; k<kmax-1; ++k) {  
        for (j=1; j<jmax-1; ++j) {  
            for (i=1; i<imax-1; ++i) {  
                x[k][j][i] = ( x[k][j][i-1] + x[k][j][i+1]  
                            + x[k][j-1][i] + x[k][j+1][i]  
                            + x[k-1][j][i] + x[k+1][j][i]) / 6.0;  
            }  
        }  
    }  
}
```



- Can we solve this in parallel but still keep the dependencies intact?


Idea: wavefront parallelization

- Parallelization approach
 - Middle (j) loop is parallel
 - Outer dimension: wavefront scheme
- Block can be updated if bottom neighbor (same threadID) and left neighbor (threadID-1) are done
- W_i : independent blocks, “wavefronts”
- After each wavefront: **synchronization** to maintain ordering




Wavefront parallelization

- Wind-up phase
 - Not all threads active
 - Each wavefront (W_i) is executed by i threads concurrently



Wavefront parallelization

- “Full pipeline”: All threads active



- Wind-down phase starts after T_0 has completed its k loop (not shown)

Wavefront parallelization with OpenMP in 3D

```
#pragma omp parallel private(nth,istart,iend,tid,kk,it,k,i) {
    nth = omp_get_num_threads();
    tid = omp_get_thread_num();
    jstart= (jmax-2)/nth * tid + 1;
    jend  = (tid==nth-1 ? jmax-2 : jstart+(jmax-2)/nth-1);
    for(t=0; t<itmax; ++t) {
        for(k=1; k<kmax-1+nth-1; ++k) {
            kk = k - tid;
            if(kk>=1 && kk<kmax-1) {
                for(j=jstart; j<=jend; ++j) {
                    for(i=1; i<kmax-1; ++i) {
                        x[kk][j][i] = ( x[kk][j][i-1] + x[kk][j][i+1]
                                      + x[kk][j-1][i] + x[kk][j+1][i]
                                      + x[kk-1][j][i] + x[kk+1][j][i])/6.0;
                    }
                }
            }
        }
        #pragma omp barrier
    }
}
```

Chop j loop into
nthreads chunks

Wind-up/-down

Wavefront sync

Wavefront parallelization – open questions

- **Global barrier** per middle loop sweep (i.e., $k_{\max}-2$ barriers overall)
 - Remedy?

Is there a **global performance limit**, i.e. **minimum balance**?

- Minimum data traffic:
 - Update whole array $\mathbf{x}(0:i_{\max}+1, 0:j_{\max}+1, 0:k_{\max}+1)$ once
 - 16 byte per element (read & write)
 - Minimum data traffic: $16*i_{\max}*j_{\max}*k_{\max}$ byte
- Work:
 - $i_{\max}*j_{\max}*k_{\max}$ Lattice Site Updates (LUPs)

→ $B_C = 16$ byte/LUP