

# Performance Modeling of Stencil Codes

Holger Stengel  
Erlangen Regional Computing Center (RRZE)  
University of Erlangen-Nuremberg

ParCo Mini-Symposium: Performance Modeling and Engineering  
September 13, 2013  
Garching, Germany



- **Motivation**
- **ECM Performance Model**
- **Stencil Examples**
  - 2D Jacobi
  - 3D 25-point long-range
- **Conclusions**



- **Common goal: Improve code performance to save resources**
- **Setting up a performance model helps to**
  - Gain insight into hardware / software interaction
  - Predict performance behavior
  - Unveil performance limitations
  - Perform purposeful optimization



## Input:

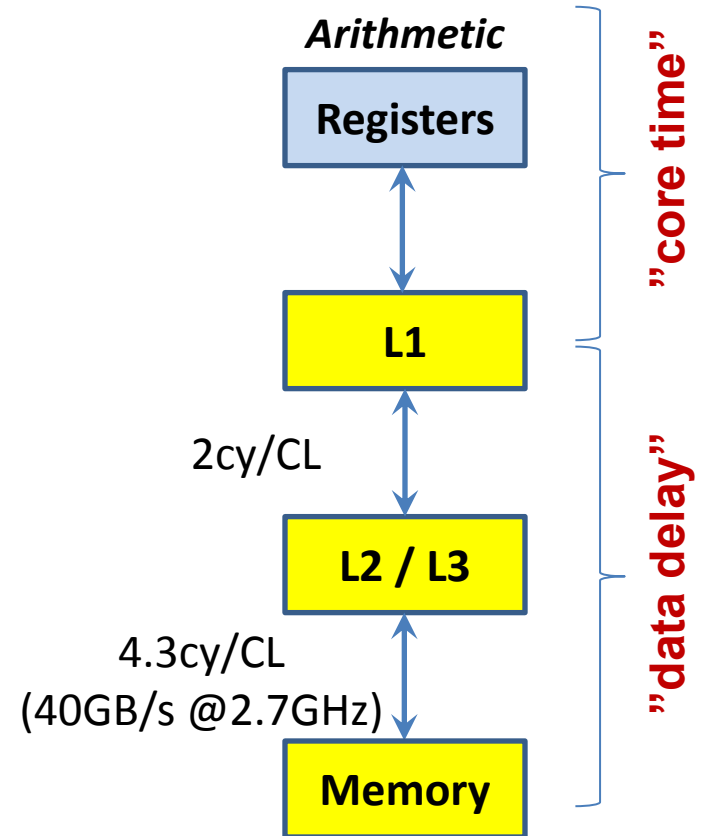
- Single-core **execution** time w/ data in L1 (e.g. from Intel IACA tool or manually)
- **Cache** line transfer between cache levels (processor handbook)
- Attainable **memory** bandwidth from (multi-) stream measurements
- Unit of work: cache line size (64B)

## Assumptions:

- Runtime contributions can **overlap**  
→ Upper and lower performance bounds
- **Scale** single-core performance until first bottleneck is hit

## Output:


- Predicts single-core performance
- Predicts scaling behavior / saturation point



\* Treibig & Hager 2010

[http://dx.doi.org/10.1007/978-3-642-14390-8\\_64](http://dx.doi.org/10.1007/978-3-642-14390-8_64)


## Example 1: 2D Jacobi in DP with SSE2 on SNB



# Example 1: 2D Jacobi in DP with SSE2 on SNB



```
1 // Jacobi 2D line update
2 for(int j=start; j<end; j++){
3     t1[i][j]= ( t0[i-1][j] +
4                 t0[i+1][j] +
5                 t0[i][j-1] +
6                 t0[i][j+1] ) * 0.25;
7 }
```



4-way unrolling  
→ 8 LUP / iteration



## Instruction count

- 13 LOAD
- 4 STORE
- 12 ADD
- 4 MUL

```
1 movups  (0xrbp,0xr15,8), 0xmm1
2 movups  16(0xrbp,0xr15,8), 0xmm3
3 movups  32(0xrbp,0xr15,8), 0xmm5
4 movups  48(0xrbp,0xr15,8), 0xmm7
5 addpd   (0xr9,0xr15,8), 0xmm1
6 addpd   16(0xr9,0xr15,8), 0xmm3
7 addpd   32(0xr9,0xr15,8), 0xmm5
8 addpd   48(0xr9,0xr15,8), 0xmm7
9 addpd   -8(0xr10,0xr15,8), 0xmm1
10 movups  8(0xr10,0xr15,8), 0xmm2
11 movups  24(0xr10,0xr15,8), 0xmm4
12 movups  40(0xr10,0xr15,8), 0xmm6
13 addpd   0xmm2, 0xmm3
14 addpd   0xmm4, 0xmm5
15 addpd   0xmm6, 0xmm7
16 addpd   0xmm2, 0xmm1
17 addpd   0xmm4, 0xmm3
18 addpd   0xmm6, 0xmm5
19 addpd   56(0xr10,0xr15,8), 0xmm7
20 mulpd   0xmm0, 0xmm1
21 mulpd   0xmm0, 0xmm3
22 mulpd   0xmm0, 0xmm5
23 mulpd   0xmm0, 0xmm7
24 movups  0xmm1, (0xr11,0xr15,8)
25 movups  0xmm3, 16(0xr11,0xr15,8)
26 movups  0xmm5, 32(0xr11,0xr15,8)
27 movups  0xmm7, 48(0xr11,0xr15,8)
```

# Example 1: 2D Jacobi in DP with SSE2 on SNB




## Processor characteristics (SSE instructions per cycle)

- 2 LOAD || (1 LOAD + 1 STORE)
- 1 ADD
- 1 MUL

## Code characteristics (SSE instructions per iteration)

- 13 LOAD
- 4 STORE
- 12 ADD
- 4 MUL



core execution:  
**12 cy**




## ■ Situation 1: Data set fits into L1 cache

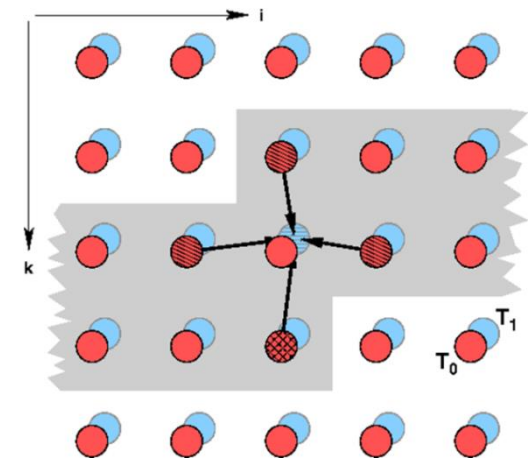
- ECM prediction:  
 $(8 \text{ LUP} / 12 \text{ cy}) * 3.5 \text{ GHz} = 2.3 \text{ GLUP/s}$
- Measurement: 2.2 GLUP/s

## ■ Situation 2: Data set fits into L2 cache (not into L1)

- 3 **additional** transfer streams from L2 to L1 (data delay)
- ECM prediction:  
 $(8 \text{ LUP} / (12+6) \text{ cy}) * 3.5 \text{ GHz} = 1.5 \text{ GLUP/s}$
- Measurement: 1.9 GLUP/s




Overlap?





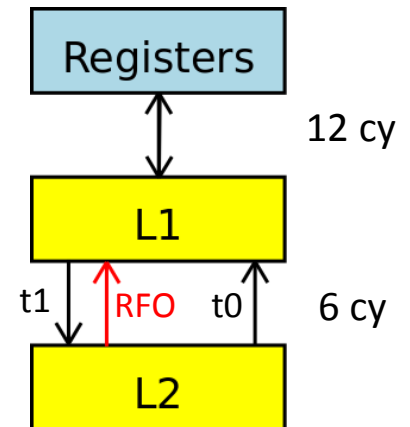
# Example 1: 2D Jacobi in DP with SSE2 on SNB



L1 „single ported“

→ no overlap during LD/ST

- ECM prediction **w/ overlap**:  
 $(8 \text{ LUP} / (8.5+6) \text{ cy}) * 3.5 \text{ GHz} = 1.9 \text{ GLUP/s}$
- Measurement: 1.9 GLUP/s



**“If the model fails, we learn something”**



---

## **Example 2: 3D long-range stencil in SP with AVX on SNB**



- 4 neighbors per direction
- Operations per update
  - 27 LOAD (25 V, 1 ROC, 1 U)
  - 1 STORE (U)
  - 26 ADD
  - 15 MUL
- Core time  
→ IACA

```
1 // 3D long-range line update (single precision)
2 for(i=4; i<nnx-4; i++) {
3     lap = coef0 * V(i,j,k)
4     + coef[1] * ( V(i+1,j ,k ) + V(i-1,j ,k ) )
5     + coef[1] * ( V(i ,j+1,k ) + V(i ,j-1,k ) )
6     + coef[1] * ( V(i ,j ,k+1) + V(i ,j ,k-1) )
7     + coef[2] * ( V(i+2,j ,k ) + V(i-2,j ,k ) )
8     + coef[2] * ( V(i ,j+2,k ) + V(i ,j-2,k ) )
9     + coef[2] * ( V(i ,j ,k+2) + V(i ,j ,k-2) )
10    + coef[3] * ( V(i+3,j ,k ) + V(i-3,j ,k ) )
11    + coef[3] * ( V(i ,j+3,k ) + V(i ,j-3,k ) )
12    + coef[3] * ( V(i ,j ,k+3) + V(i ,j ,k-3) )
13    + coef[4] * ( V(i+4,j ,k ) + V(i-4,j ,k ) )
14    + coef[4] * ( V(i ,j+4,k ) + V(i ,j-4,k ) )
15    + coef[4] * ( V(i ,j ,k+4) + V(i ,j ,k-4) );
16    U(i,j,k) = 2.f * V(i,j,k) - U(i,j,k) + ROC2(i,j,k) * lap;
17 }
```



- **Performs architecture-specific code analysis**
- **Prerequisite: Mark start and end of dominant work loop**
  - In high-level code (documented)
  - In assembly code (see `iacaMarks.h`)
    - Does not influence code optimization (e.g. vectorization)
    - Assembly loop might perform multiple updates per iteration (unrolling, SIMD)
- **Important reports (throughput mode):**
  - Block throughput: runtime of one loop iteration (→ core-time)
  - Throughput bottleneck: limiting resource for code execution
  - Port pressure: dominant pipeline port



```
[..]
Throughput Analysis Report
-----
Block Throughput: 34.25 Cycles          Throughput Bottleneck: FrontEnd

Port Binding In Cycles Per Iteration:
-----
| Port | 0 - DV | 1 | 2 - D | 3 - D | 4 | 5 |
-----
| Cycles | 19.0  0.0 | 26.0 | 31.5  30.5 | 31.5  30.5 | 2.0 | 24.0 |
-----
```

[..]

```
| Num Of |           Ports pressure in cycles           | |
| Uops   | 0 - DV | 1 | 2 - D | 3 - D | 4 | 5 | |
-----
| 1      |         |   | 0.5  0.5 | 0.5  0.5 |   |   | | vmovups xmm1, xmmword ptr [r12+r14*4+0x10]
| 1      |         |   | 0.5  0.5 | 0.5  0.5 |   |   | | vmovups xmm0, xmmword ptr [rsi+r14*4+0x10]
| 1      |         |   | 0.5  0.5 | 0.5  0.5 |   |   | | vmovups xmm13, xmmword ptr [r12+r14*4+0x14]
| 1      |         |   | 0.5  0.5 | 0.5  0.5 |   |   | | vmovups xmm14, xmmword ptr [r12+r14*4+0xc]
| 1      |         |   | 0.5  0.5 | 0.5  0.5 |   |   | | vmovups xmm12, xmmword ptr [rdi+r14*4+0x10]
| 1      |         |   | 0.5  0.5 | 0.5  0.5 |   |   | | mov r15, qword ptr [rsp+0x338]
| 2      |         |   | 0.5  0.5 | 0.5  0.5 |   | 1.0 | | vinsertf128 ymm10, ymm1, xmmword ptr [r12+r
| 1      | 1.0     |   |         |         |   |   | | vmulps ymm1, ymm7, ymm10
| 1      | 1.0     |   |         |         |   |   | | vmulps ymm15, ymm8, ymm10
| 2      |         |   | 0.5  0.5 | 0.5  0.5 |   | 1.0 | | vinsertf128 ymm11, ymm0, xmmword ptr [rsi+r
| 1      |         | 1.0 |         |         |   |   | | vsubps ymm1, ymm1, ymm11
[..]
```

→ AVX vectorization, no unrolling: One iteration updates 8 SP (float) elements

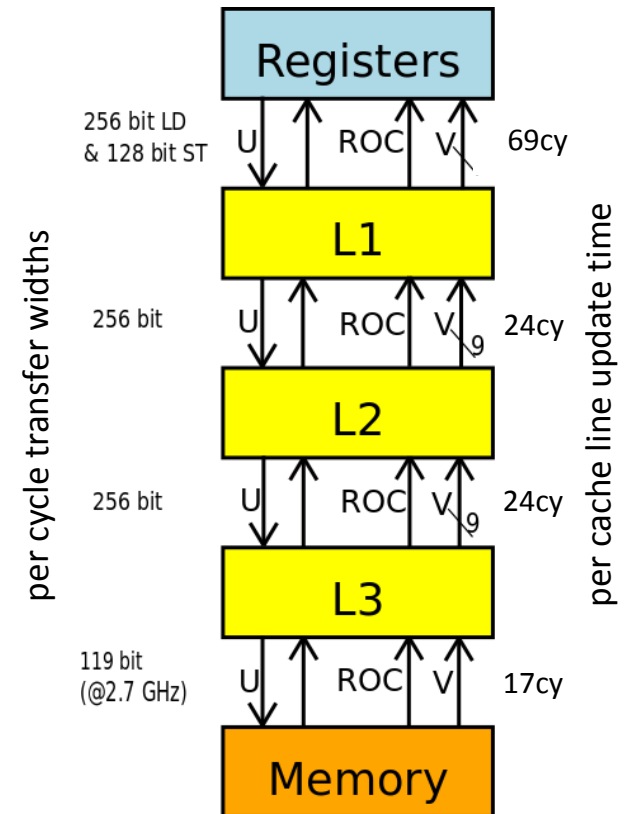


- 4 neighbors per direction
- Operations per update
  - 27 LOAD (25 V, 1 ROC, 1 U)
  - 1 STORE (U)
  - 26 ADD
  - 15 MUL
- Core time (IACA)
  - 34.25 cy / 8 LUP (SP)  
→ 69 cy / CL
  - LOAD dominated

```
1 // 3D long-range line update (single precision)
2 for(i=4; i<nnx-4; i++) {
3     lap = coef0 * V(i,j,k)
4         + coef[1] * ( V(i+1,j ,k ) + V(i-1,j ,k ) )
5         + coef[1] * ( V(i ,j+1,k ) + V(i ,j-1,k ) )
6         + coef[1] * ( V(i ,j ,k+1) + V(i ,j ,k-1) )
7         + coef[2] * ( V(i+2,j ,k ) + V(i-2,j ,k ) )
8         + coef[2] * ( V(i ,j+2,k ) + V(i ,j-2,k ) )
9         + coef[2] * ( V(i ,j ,k+2) + V(i ,j ,k-2) )
10        + coef[3] * ( V(i+3,j ,k ) + V(i-3,j ,k ) )
11        + coef[3] * ( V(i ,j+3,k ) + V(i ,j-3,k ) )
12        + coef[3] * ( V(i ,j ,k+3) + V(i ,j ,k-3) )
13        + coef[4] * ( V(i+4,j ,k ) + V(i-4,j ,k ) )
14        + coef[4] * ( V(i ,j+4,k ) + V(i ,j-4,k ) )
15        + coef[4] * ( V(i ,j ,k+4) + V(i ,j ,k-4) );
16    U(i,j,k) = 2.f * V(i,j,k) - U(i,j,k) + ROC2(i,j,k) * lap;
17 }
```




- **Data delay**
  - Spatial blocking for minimal traffic between L3 cache and memory
  - 8 additional streams for smaller caches
- **Single-core performance**
  - $2.7\text{GHz} / (134\text{cy} / 16\text{LUP}) = 322\text{MLUP/s}$
  - Measurement: 320MLUP/s
  - LOAD dominated → no overlap expected
- **Socket scaling (8 cores)**
  - $8 * 322\text{MLUP/s} = 2576\text{MLUP/s}$
  - Limit:  $40\text{GB/s} / 16\text{B/LUP} = 2500\text{MLUP/s}$ 
    - Saturation at 8 cores
  - Measurement:  $\sim 2200\text{MLUP/s}$  (88% max.)
- **Optimization possibilities**
  - Data transfer well optimized (blocking)
  - Reduce core time (LD); limit: ADD ( $2 * 26\text{cy}$ )
    - Possible speedup:  $69\text{cy} - 52\text{cy} = 17\text{cy}$  ( $\sim 15\%$ )





## 3D long-range stencil

1 socket, Sandy Bridge, 2.7GHz, intel64/13.1up03







*"There is nothing as practical as a good theory."*

Kurt Lewin (1890-1947)