

MAQAO Hands-on exercises

LProf: Lightweight Generic Profiler

CQA: Code Quality Analyzer

ONE View: Performance View Aggregator

Setup

Copy handson material

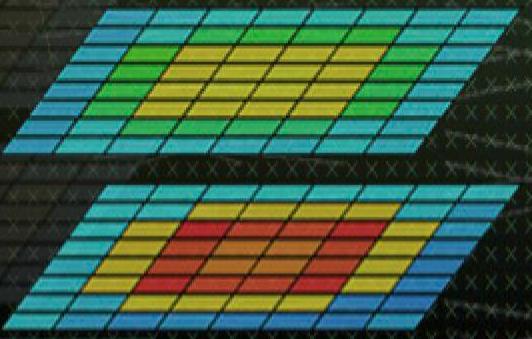
```
$> cp /home/cv991149/tutorial/MAQAO_HANDSON.tar.gz $WORK
```

Inflate the archive at the root of your home folder

```
$> cd $WORK  
$> tar zxvf MAQAO_HANDSON.tar.gz  
$> cd MAQAO_HANDSON
```

Load MAQAO

```
$> module load UNITE  
UNITE loaded  
$> module load maqao/2.2.1  
maqao/2.2.1 loaded  
$> maqao --version  
maqao 2.2.1 - ...
```



MAQAO LProf Hands-on exercises

Updating jobscript

Edit the jobscrip to add invocation of the profiler

```
$> vim $WORK/MAQAO_HANDSON/lprof/run_lprof.lsf
```

```
...
module load intel openmpi
module list
module load UNITE
module load maqao/2.2.1
...
cd ${WORK}/NPB3.3-MZ-MPI/bin
$MPIEXEC $FLAGS_MPI_BATCH maqao lprof bt-mz_C.8
```

Launching job and viewing results

Launch the job

```
$> bsub < $WORK/MAQAO_HANDSON/lprof/run_lprof.lsf
```

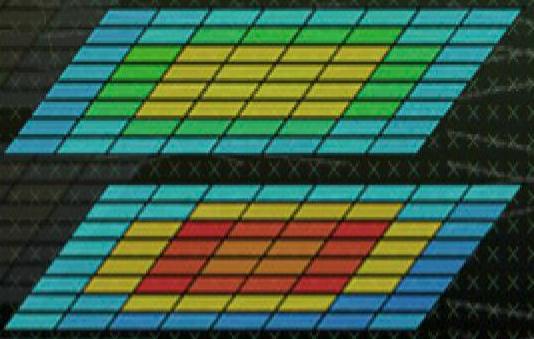
The experiment directory is generated in the execution directory (`(${WORK}/NPB3.3-MZ-MPI/bin)`). By default, it is named `maqao_<timestep>` (with `<timestep>` in the form of mmDDYYYYHHMM).

A sample experiment directory is located in

`MAQAO_HANDSON/lprof/sample_xpdir`

Visualize the results:

```
$> cd ${WORK}/NPB3.3-MZ-MPI/bin  
$> maqao lprof xp=<xpdir> -df  
$> maqao lprof xp=<xpdir> -df -dt  
$> maqao lprof xp=<xpdir> -dl  
$> maqao lprof xp=<xpdir> -df -cv=full
```



MAQAO / CQA Hands-on exercises

Setup

CQA can be directly executed on a login node because it uses static analysis

Login to CLAIX

```
$> ssh -Y <your_login>@login.hpc.itc.rwth-aachen.de
```

Load MAQAO environment

```
$> module load maqao
```

Load a recent GCC compiler

```
$> module load gcc/6.3
```

Switch to CQA handson folder

```
$> cd $WORK/MAQAO_HANDSON/cqa/matmul
```

Matrix Multiply code

```
void kernel0 (int n,
              float a[n][n],
              float b[n][n],
              float c[n][n]) {
    int i, j, k;

    for (i=0; i<n; i++)
        for (j=0; j<n; j++) {
            c[i][j] = 0.0f;
            for (k=0; k<n; k++)
                c[i][j] += a[i][k] * b[k][j];
        }
}
```

“Naïve” dense matrix multiply implementation in C

Compiling, running and analyzing kernel0 in -O3

Compiling and executing kernel 0

```
$> make OPTFLAGS=-O3 KERNEL=0  
$> ./matmul 100 1000  
Cycles per FMA: 2.20
```

Analyzing with CQA

```
$> maqao cqa matmul fct-loops=kernel0
```

NB: the usual way to use CQA consists in finding IDs of hot loops with the MAQAO profiler and forwarding them to CQA (loop=17,42...).

To simplify this hands-on, we will bypass profiling and directly request CQA to analyze all innermost loops in functions (max 2-3 loops/function for this hands-on).

For generating and viewing CQA HTML output

```
$> maqao cqa matmul fct-loops=kernel0 of=html  
$> firefox cqa_html/index.html
```

CQA output for kernel0 (from the “gain” confidence level)

Vectorization

(...) By fully vectorizing your loop,
you can lower the cost of an iteration
from 3.00 to 0.38 cycles (8.00x
speedup). (...)

- Remove inter-iterations dependences from your loop and make it unit-stride.

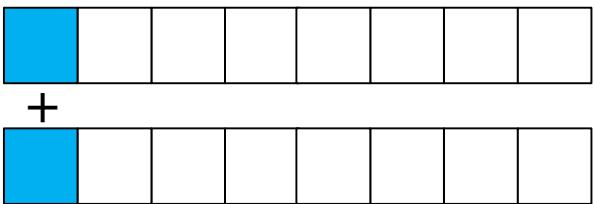
- * If your arrays have 2 or more dimensions, **check whether elements are accessed contiguously and, otherwise, try to permute loops accordingly:**

C storage order is row-major: `for(i)
a[j][i] = b[j][i];` (slow, non stride 1)
=> `for(i) for(j) a[i][j] = b[i][j];`
(fast, stride 1)

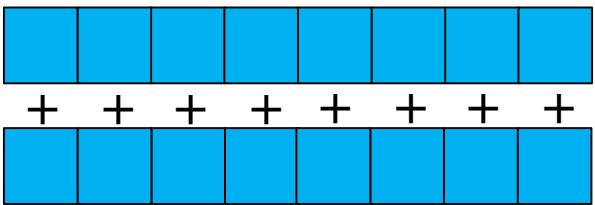
- * If your loop streams arrays of structures (AoS), try to use (...) SoA

Vectorization (summing elements):

VADDSS
(scalar)



VADDPS
(packed)



- Accesses are not contiguous => let us permute k and j loops
- No structures here...

CQA output for kernel0 (from the “gain” confidence level)

Code quality analysis

Source loop ending at line 10 in ...NDSON_test/CQA/matmul/kernel.c

It is composed of the loop 2

MAQAO binary loop id: 2

The loop is defined in /home/hpc/a2c06/lu23voj/MAQAO_HANDSON_test/CQA/matmul/kernel.c:9-10
In the binary file, the address of the loop is: 4009f0

8% of peak computational performance is used (0.67 out of 8.00 FLOP per cycle (GFLOPS @ 1GHz))

Gain **Potential gain** **Hints** **Experts only**

Vectorization status

Your loop is not vectorized (all SSE/AVX instructions are used in scalar mode).
Only 25% of vector length is used.

Vectorization

Your loop is processing FP elements but is NOT OR PARTIALLY VECTORIZED and could benefit from full vectorization.
By fully vectorizing your loop, you can lower the cost of an iteration from 3.00 to 0.75 cycles (4.00x speedup).

Since your execution units are vector units, only a fully vectorized loop can use their full power.

Proposed solution(s):

Two propositions:

- Try another compiler or update/tune your current one:
- Remove inter-iterations dependences from your loop and make it unit-stride.

* If your arrays have 2 or more dimensions, check whether elements are accessed contiguously and, otherwise, try to permute loops accordingly:
C storage order is row-major: for(i) for(j) a[j][i] = b[j][i]; (slow, non stride 1) => for(i) for(j) a[i][j] = b[i][j]; (fast, stride 1)

* If your loop streams arrays of structures (AoS), try to use structures of arrays instead (SoA):
for(l) a[l].x = b[l].x; (slow, non stride 1) => for(l) a.x[l] = b.x[l]; (fast, stride 1)

Bottlenecks

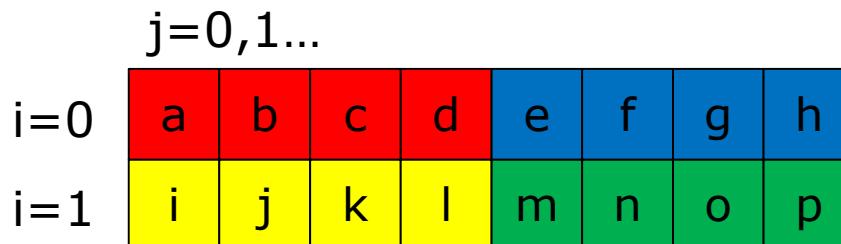
Detected a non usual bottleneck.

Proposed solution(s):

- Pass to your compiler a micro-architecture specialization option:
* use march=native.

Impact of loop permutation on data access

Logical mapping



Efficient vectorization +
prefetching

Physical mapping

(C stor. order: row-major)



```
for (j=0; j<n; j++)  
  for (i=0; i<n; i++)  
    f(a[i][j]);
```



```
for (i=0; i<n; i++)  
  for (j=0; j<n; j++)  
    f(a[i][j]);
```



Removing inter-iteration dependences and getting stride 1 by permuting loops on j and k

```
void kernell1 (int n,
               float a[n][n],
               float b[n][n],
               float c[n][n]) {
    int i, j, k;

    for (i=0; i<n; i++) {
        for (j=0; j<n; j++)
            c[i][j] = 0.0f;

        for (k=0; k<n; k++)
            for (j=0; j<n; j++)
                c[i][j] += a[i][k] * b[k][j];
    }
}
```

Kernel1: loop interchange

Compiling and executing kernel 1

```
$> make clean  
$> make OPTFLAGS=-O3 KERNEL=1  
$> ./matmul 100 1000  
Cycles per FMA: 0.56
```

Analyzing with CQA

```
$> maqao cqa matmul fct-loops=kernel1
```

For generating and viewing CQA HTML output

```
$> maqao cqa matmul fct-loops=kernel1 of=html  
$> firefox cqa_html/index.html
```

CQA output for kernel1

Vectorization status

Your loop is fully vectorized (...)
but on 50% vector length.

Vectorization

- Pass to your compiler a micro-
architecture specialization option:
 * **use march=native**
- Use vector aligned instructions...

FMA

Presence of both ADD/SUB and MUL
operations.

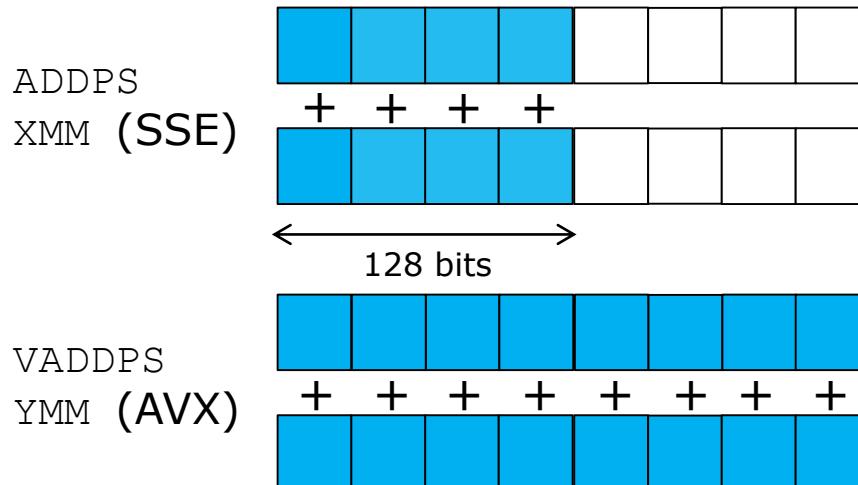
- Pass to your compiler a micro-
architecture specialization option...
- Try to change order in which...

- Let's add **-march=native** to OPTFLAGS

Impacts of architecture specialization: vectorization and FMA

- Vectorization
 - SSE instructions (SIMD 128 bits) used on a processor supporting AVX ones (SIMD 256 bits)
 - => 50% efficiency loss

- FMA
 - Fused Multiply-Add ($A + BC$)
 - Intel architectures: supported on MIC/KNC and Xeon starting from Haswell



```
# A = A + BC  
  
VMULPS <B>, <C>, %XMM0  
VADDPS <A>, %XMM0, <A>  
# can be replaced with  
something like:  
VFMADD312PS <B>, <C>, <A>
```

Kernel1 + **-march=native**

Compiling and executing kernel 1 with additional flag

```
$> make clean  
$> make OPTFLAGS="-O3 -march=native" KERNEL=1  
$> ./matmul 100 1000  
Cycles per FMA: 0.38
```

Analyzing with CQA

```
$> maqao cqa matmul fct-loops=kernel1 --confidence-  
levels=gain,hint
```

For generating and viewing CQA HTML output

```
$> maqao cqa matmul fct-loops=kernel1 of=html  
$> firefox cqa_html/index.html
```

CQA output for kernel1 (using “gain” and “hint” conf. levels)

Vectorization status

Your loop is fully vectorized (...)

Vector unaligned load/store...

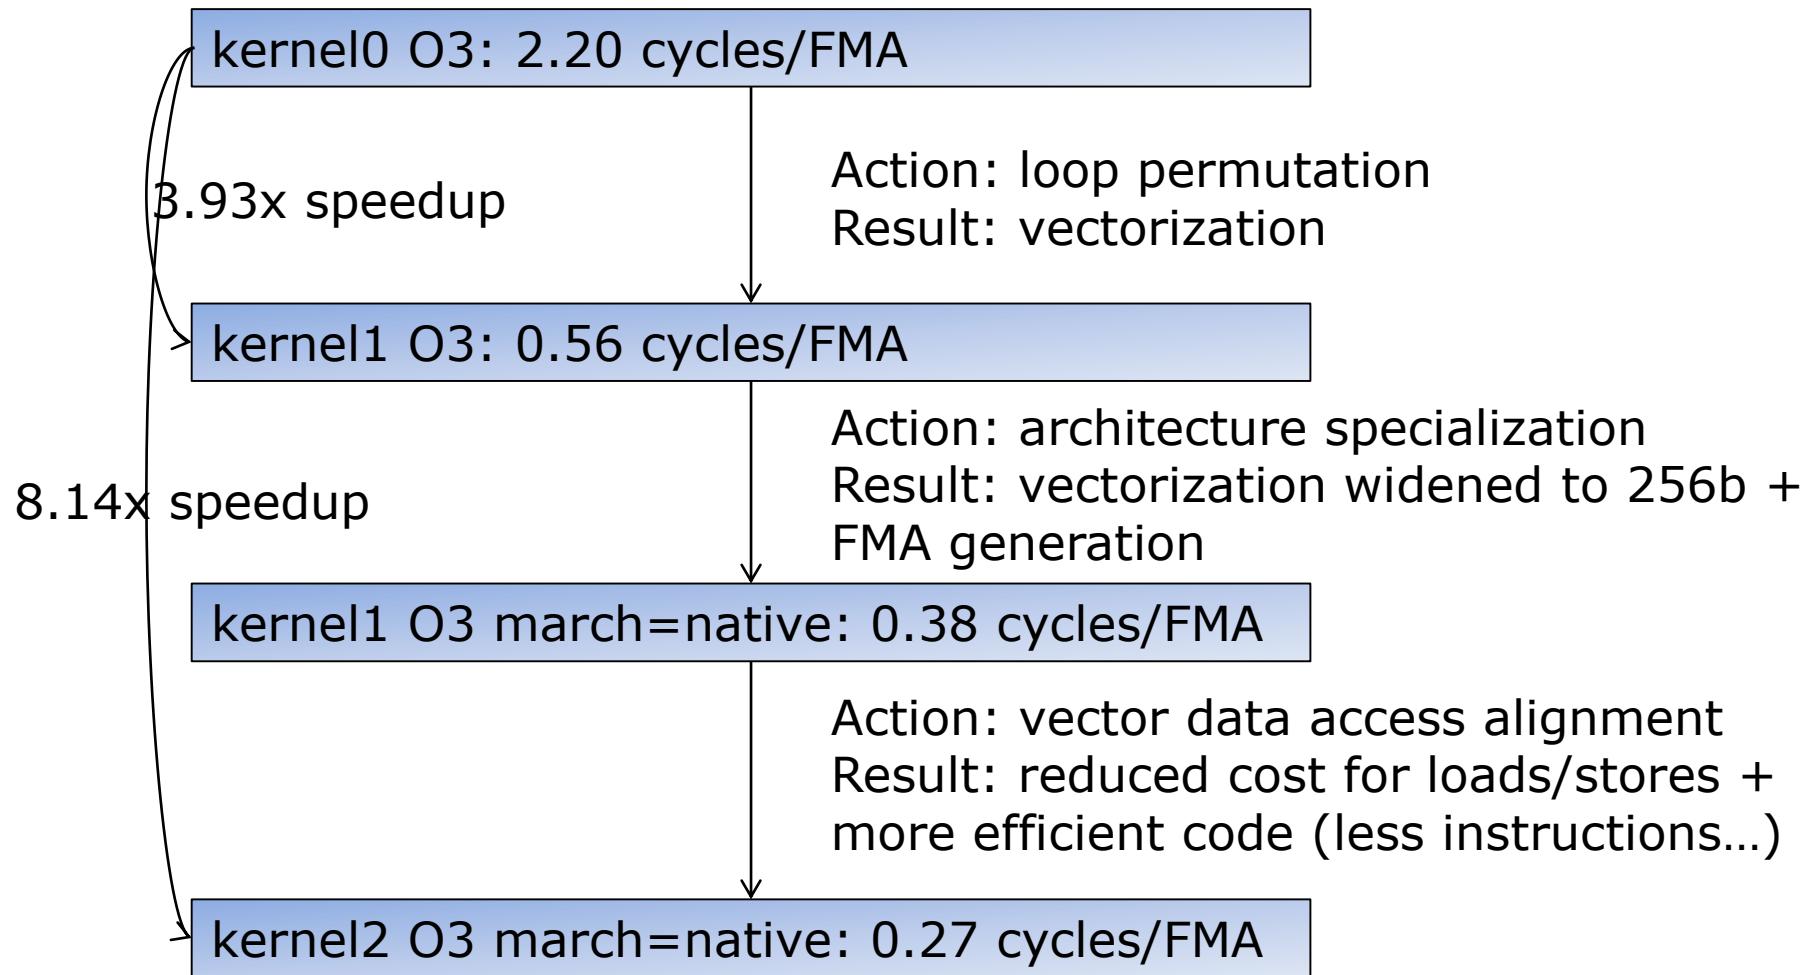
- Use vector aligned instructions:
 1) align your arrays on 32 bytes
boundaries,
 2) inform your compiler that your
arrays are vector aligned: use the
 _builtin_assume_aligned built-in

- Let's switch to the next proposal: vector aligned instructions

kernel2: assuming aligned vector accesses

```
$> make clean
$> make OPTFLAGS="-O3 -march=native" KERNEL=2
$> ./matmul 100 1000
Cannot call kernel2 on matrices with size%8 != 0 (data non
aligned on 32B boundaries)
Aborted
$> ./matmul 104 1000
Cycles per FMA: 0.27
```

Summary of optimizations and gains



Hydro example

Switch to the other CQA handson folder

```
$> cd $WORK/MAQAO_HANDSON/cqa/hydro
```

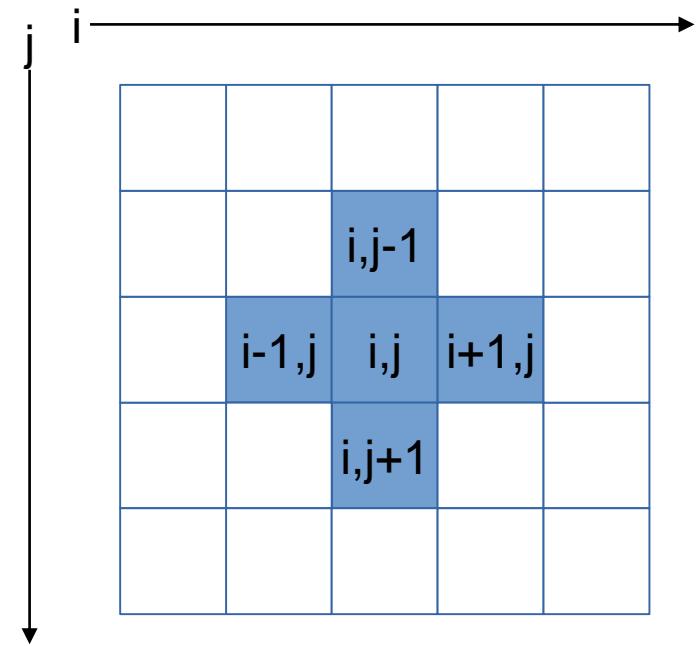
Hydro code

```
int build_index (int i, int j, int grid_size)
{
    return (i + (grid_size + 2) * j);
}

void linearSolver0 (...) {
    int i, j, k;

    for (k=0; k<20; k++)
        for (i=1; i<=grid_size; i++)
            for (j=1; j<=grid_size; j++)
                x[build_index(i, j, grid_size)] =
(a * ( x[build_index(i-1, j, grid_size)] +
        x[build_index(i+1, j, grid_size)] +
        x[build_index(i, j-1, grid_size)] +
        x[build_index(i, j+1, grid_size)] )
        + x0[build_index(i, j, grid_size)])
    ) / c;
}
```

Iterative linear system solver
using the Gauss-Siedel
relaxation technique.
« Stencil » code



Compiling, running and analyzing kernel0 (icc -O3 -xHost)

params: <grid size> <# repetitions>

```
$> make KERNEL=0
$> ./hydro 250 50
Cycles per element for solvers: 2064.14
$> maqao lprof xp=sx -- ./hydro 250 50
$> maqao lprof xp=sx -dl | head
#####
# Loop ID | Function Name | Source Info | Level | Time (%)
#####
# 144     | project      | 103,105@kernel.c | Innermost | 28.29
# 55      | c_densitySolver | 103,105@kernel.c | Innermost | 23.03
# 95      | c_velocitySolver | 103,105@kernel.c | Innermost | 21.71
# 88      | c_velocitySolver | 103,105@kernel.c | Innermost | 19.08
# 142     | project      | 371,374@kernel.c | Innermost | 1.32

$> maqao cqa hydro loop=144 --confidence-levels=gain,potential,hint
```

In this application the kernel routine, linearSolver, were inlined in caller functions. Moreover, there is here direct mapping between source and binary loop. Consequently the 4 highlighted loops are identical and only one needs analysis.

CQA output for kernel0

Composition and unrolling

It is composed of the loop 144
and is **not unrolled or unrolled with**
no peel/tail loop.

The analysis will be displayed for
the requested loops: 144

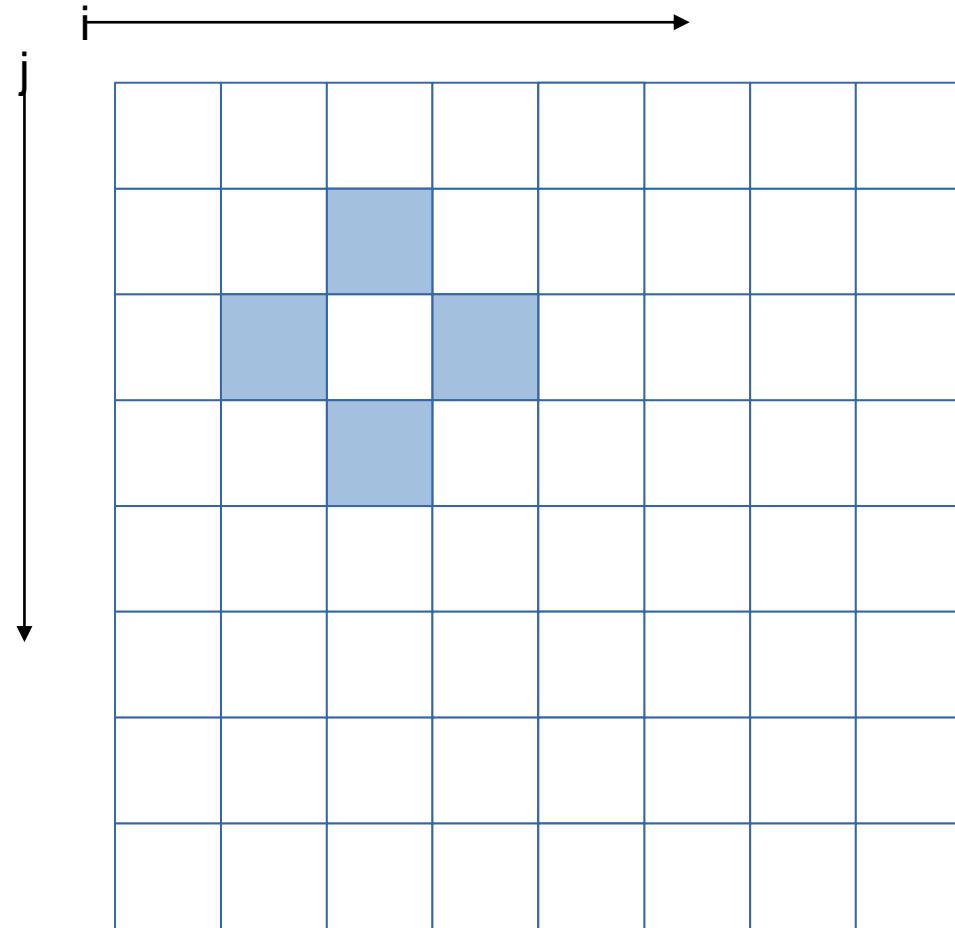
Unroll opportunity

Loop is potentially data access
bound.

**Unroll your loop if trip count is
significantly higher than target
unroll factor and if some data
references are common to consecutive
iterations...**

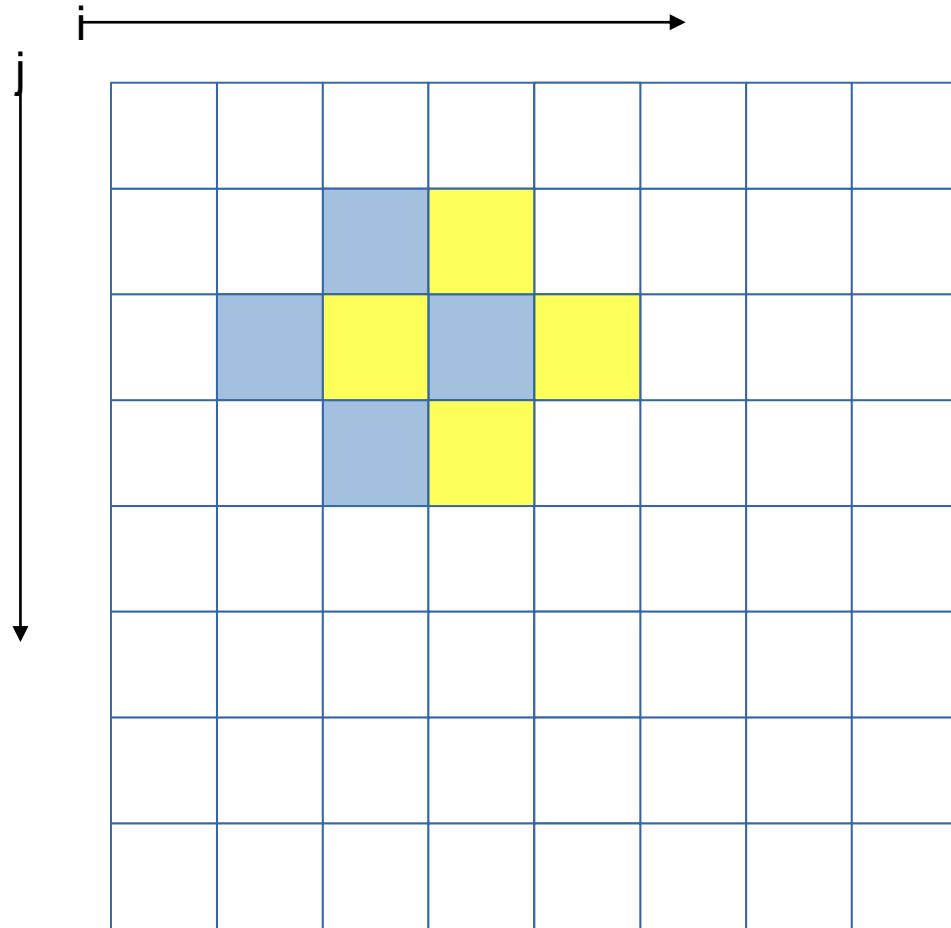
Unrolling is generally a good deal:
fast to apply and often provides
gain. Let's try to reuse data
references through unrolling

Memory references reuse : 4x4 unroll footprint on loads



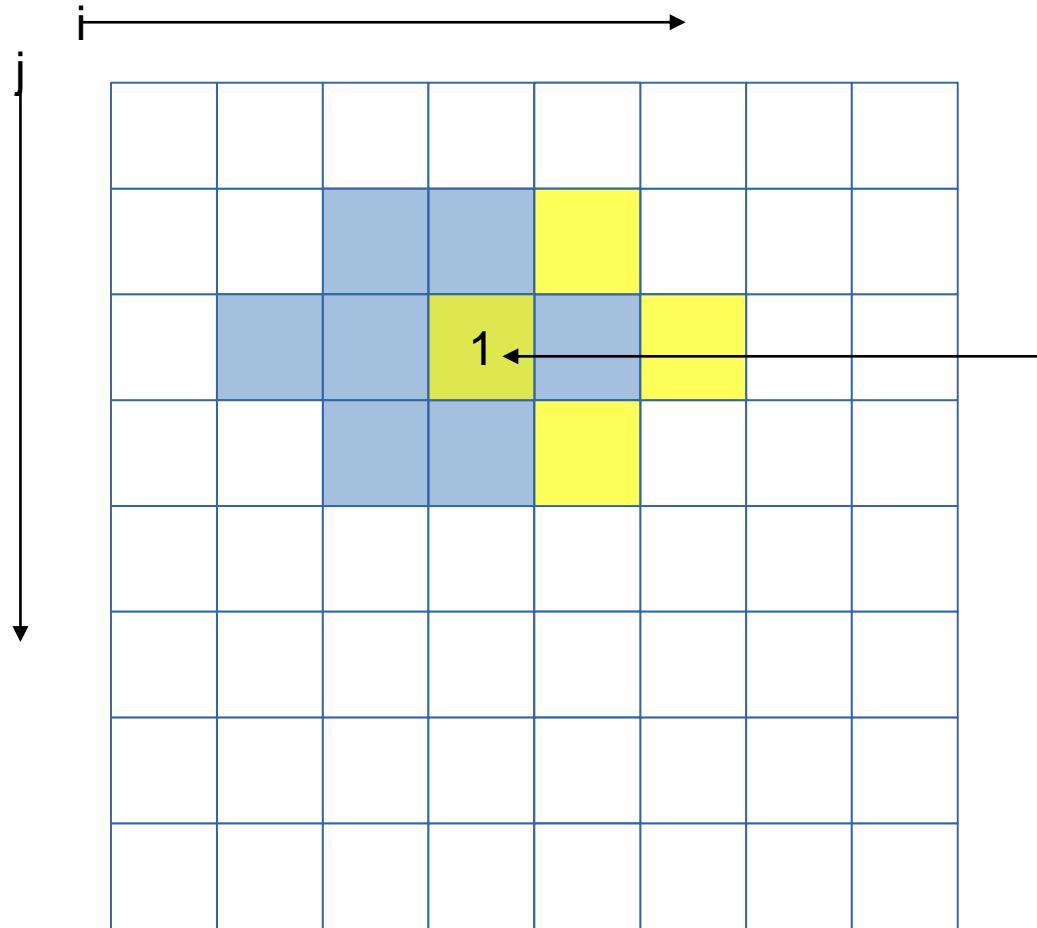
LINEAR_SOLVER($i+0, j+0$)

Memory references reuse : 4x4 unroll footprint on loads



LINEAR_SOLVER($i+0, j+0$)
LINEAR_SOLVER($i+1, j+0$)

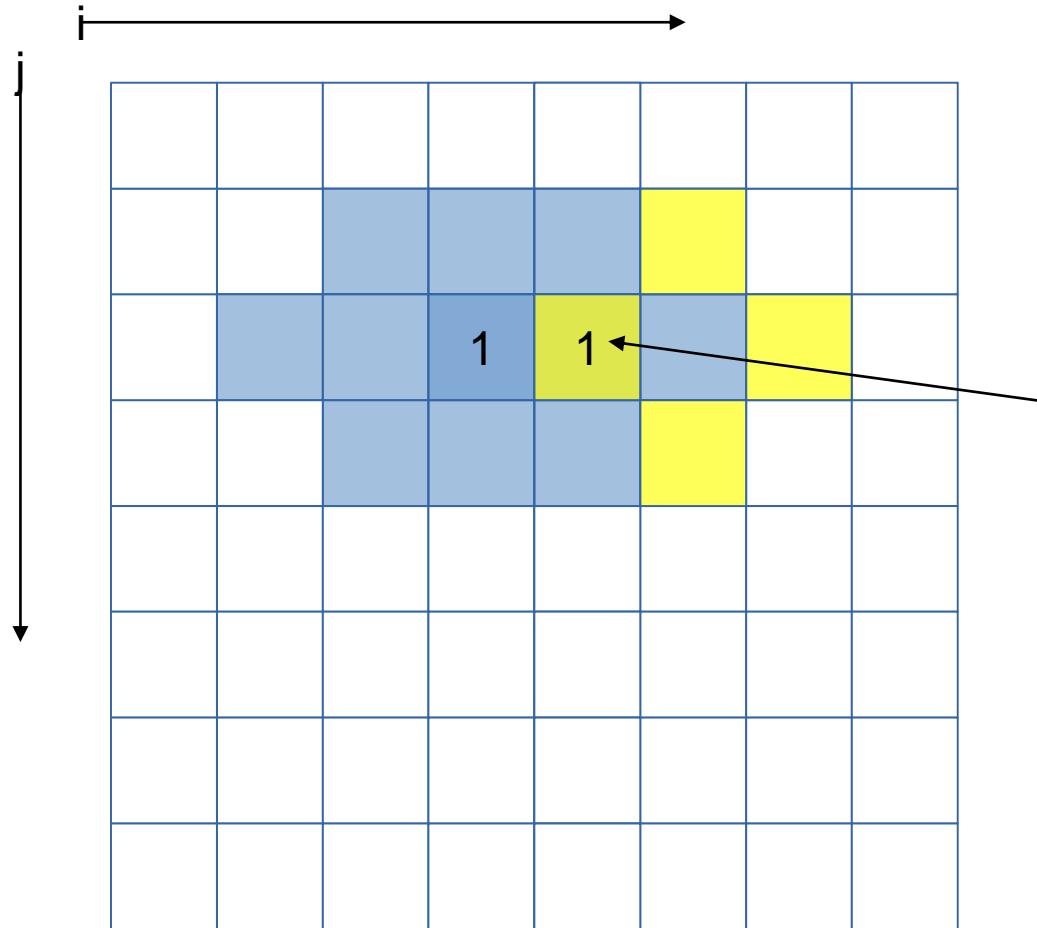
Memory references reuse : 4x4 unroll footprint on loads



LINEAR_SOLVER($i+0, j+0$)
LINEAR_SOLVER($i+1, j+0$)
LINEAR_SOLVER($i+2, j+0$)

1 reuse

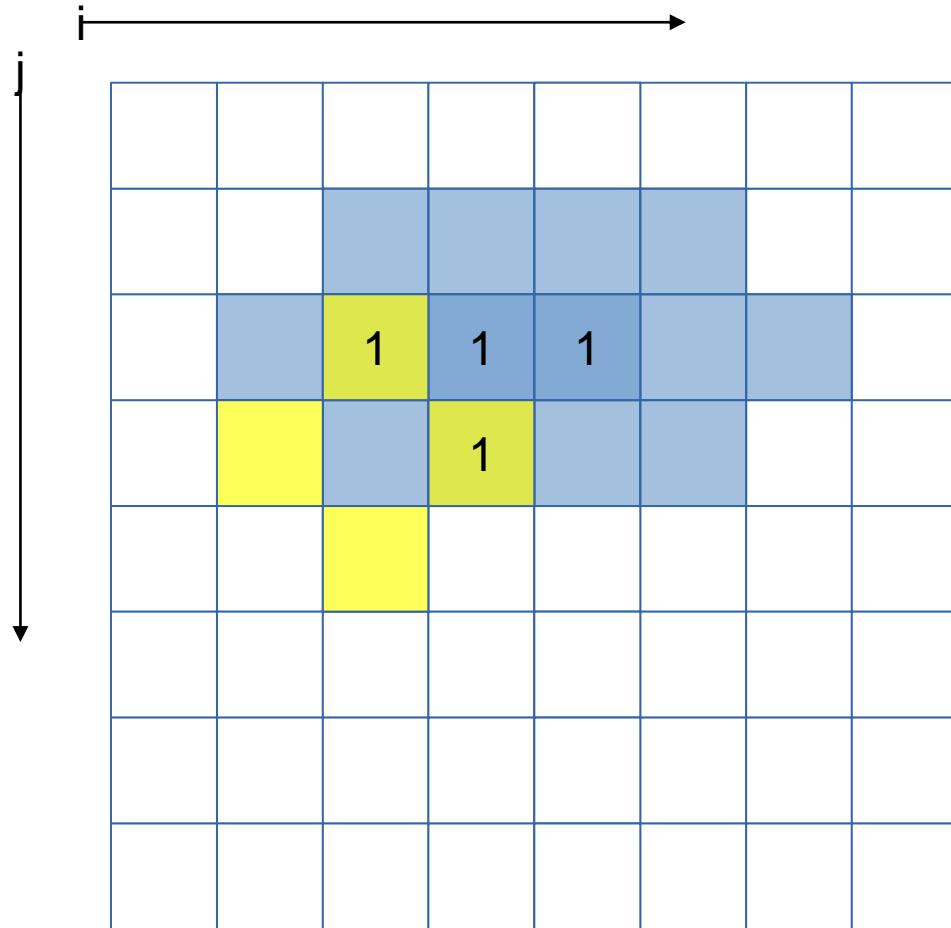
Memory references reuse : 4x4 unroll footprint on loads



LINEAR_SOLVER($i+0, j+0$)
LINEAR_SOLVER($i+1, j+0$)
LINEAR_SOLVER($i+2, j+0$)
LINEAR_SOLVER($i+3, j+0$)

2 reuses

Memory references reuse : 4x4 unroll footprint on loads



LINEAR_SOLVER($i+0, j+0$)

LINEAR_SOLVER($i+1, j+0$)

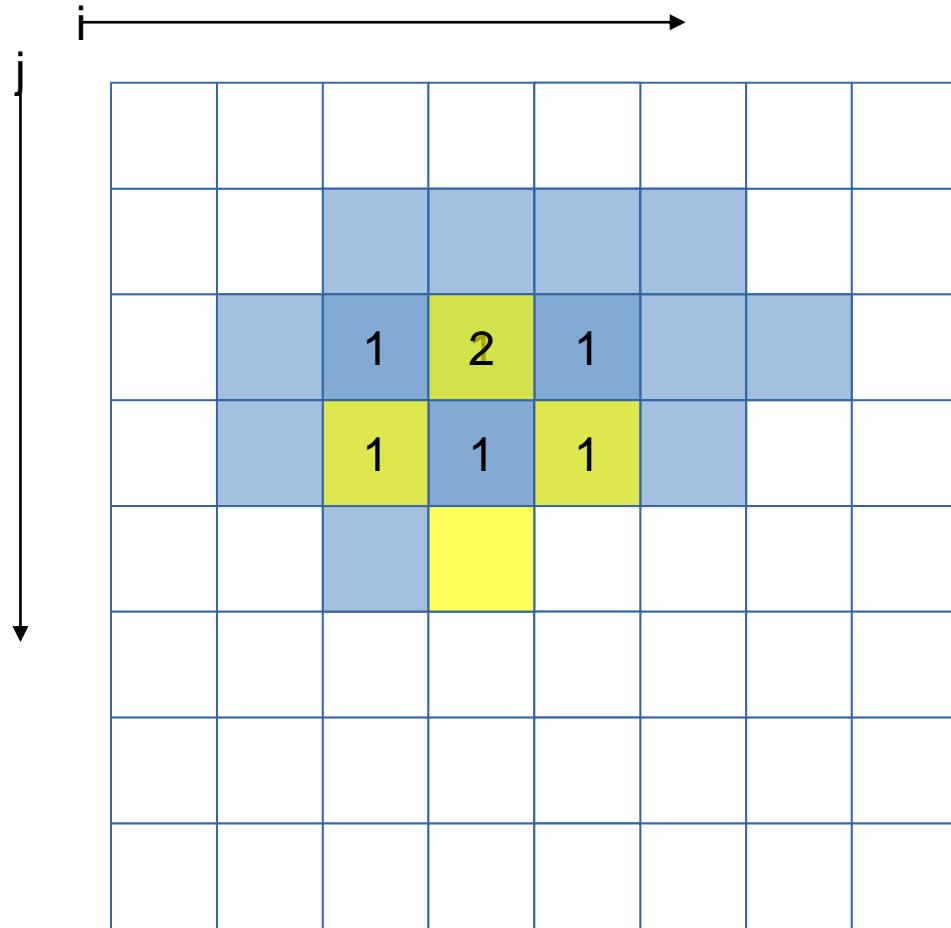
LINEAR_SOLVER($i+2, j+0$)

LINEAR_SOLVER($i+3, j+0$)

LINEAR_SOLVER($i+0, j+1$)

4 reuses

Memory references reuse : 4x4 unroll footprint on loads



LINEAR_SOLVER(i+0,j+0)

LINEAR_SOLVER(i+1,j+0)

LINEAR_SOLVER(i+2,j+0)

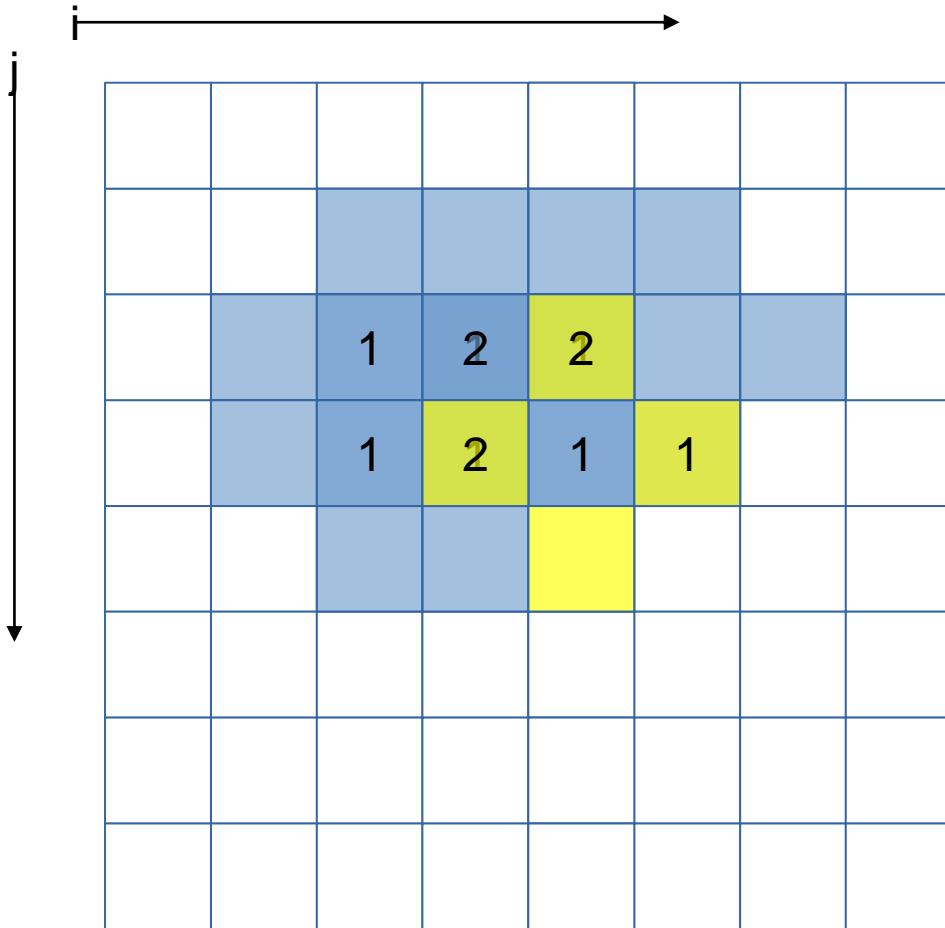
LINEAR_SOLVER(i+3,j+0)

LINEAR_SOLVER(i+0,j+1)

LINEAR_SOLVER(i+1,j+1)

7 reuses

Memory references reuse : 4x4 unroll footprint on loads

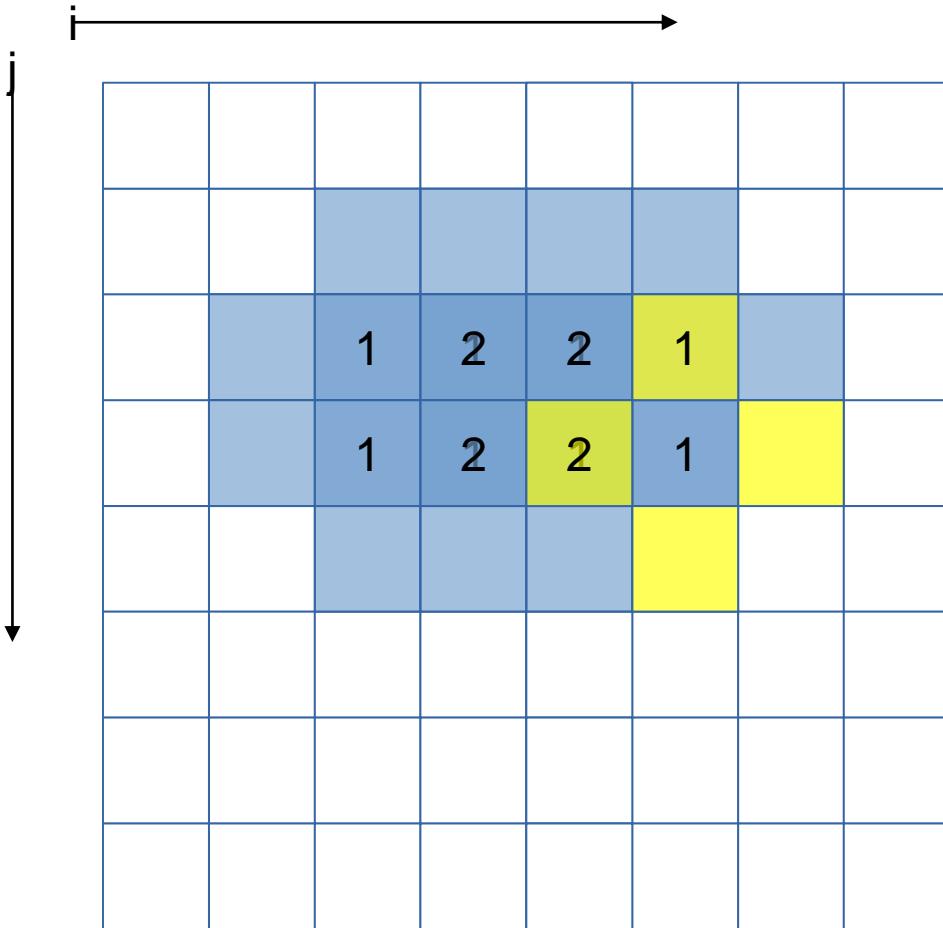


LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

LINEAR_SOLVER(i+0,j+1)
LINEAR_SOLVER(i+1,j+1)
LINEAR_SOLVER(i+2,j+1)

10 reuses

Memory references reuse : 4x4 unroll footprint on loads

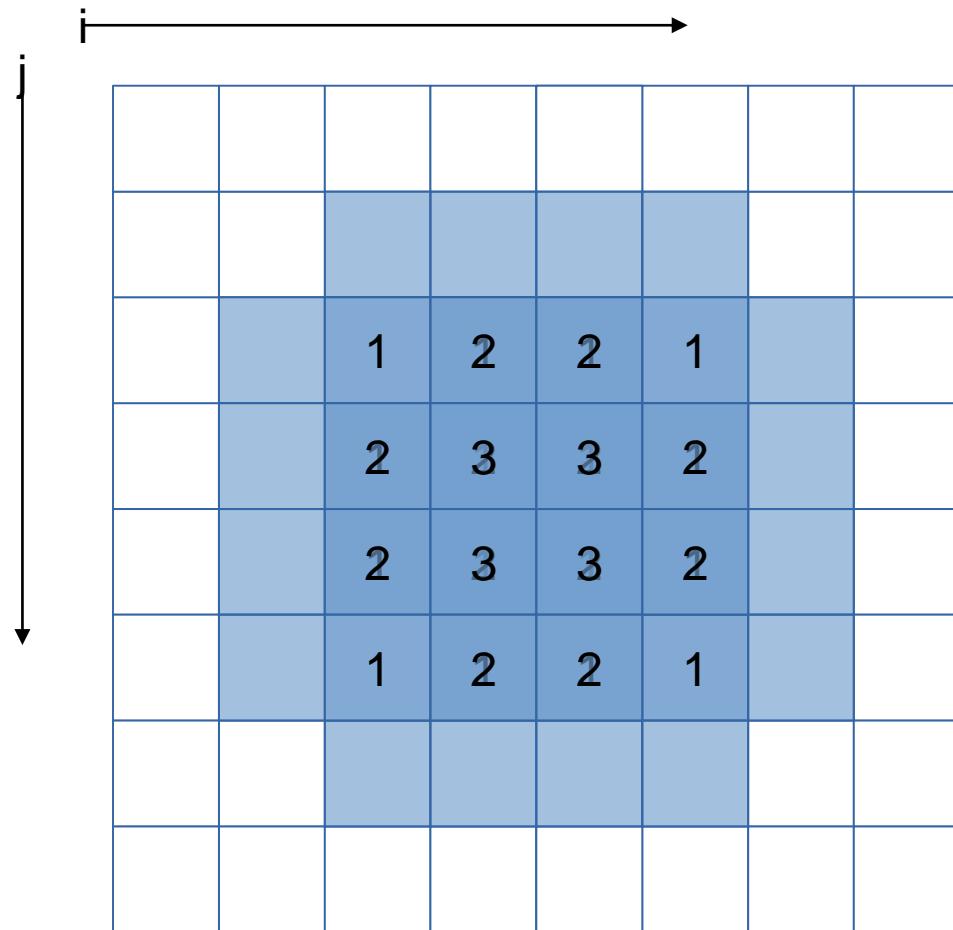


LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

LINEAR_SOLVER(i+0,j+1)
LINEAR_SOLVER(i+1,j+1)
LINEAR_SOLVER(i+2,j+1)
LINEAR_SOLVER(i+3,j+1)

12 reuses

Memory references reuse : 4x4 unroll footprint on loads



LINEAR_SOLVER($i+0-3, j+0$)

LINEAR_SOLVER($i+0-3, j+1$)

LINEAR_SOLVER($i+0-3, j+2$)

LINEAR_SOLVER($i+0-3, j+3$)

32 reuses

Impacts of memory reuse

- For the x array, instead of $4 \times 4 \times 4 = 64$ loads,
now only 32 (32 loads avoided by reuse)
- For the x0 array no reuse possible : 16 loads
- Total loads : 48 instead of 80

4x4 unroll

```
#define LINEARSOLVER(...) x[build_index(i, j, grid_size)] = ...  
  
void linearSolver2 (...) {  
    (...)  
  
    for (k=0; k<20; k++)  
        for (i=1; i<=grid_size-3; i+=4)  
            for (j=1; j<=grid_size-3; j+=4) {  
                LINEARSOLVER (... , i+0, j+0);  
                LINEARSOLVER (... , i+0, j+1);  
                LINEARSOLVER (... , i+0, j+2);  
                LINEARSOLVER (... , i+0, j+3);  
  
                LINEARSOLVER (... , i+1, j+0);  
                LINEARSOLVER (... , i+1, j+1);  
                LINEARSOLVER (... , i+1, j+2);  
                LINEARSOLVER (... , i+1, j+3);  
  
                LINEARSOLVER (... , i+2, j+0);  
                LINEARSOLVER (... , i+2, j+1);  
                LINEARSOLVER (... , i+2, j+2);  
                LINEARSOLVER (... , i+2, j+3);  
  
                LINEARSOLVER (... , i+3, j+0);  
                LINEARSOLVER (... , i+3, j+1);  
                LINEARSOLVER (... , i+3, j+2);  
                LINEARSOLVER (... , i+3, j+3);  
            }  
    }  
}
```

grid_size must now be multiple of 4. Or loop control must be adapted (much less readable) to handle leftover iterations

Kernel1

```
> make clean
> make KERNEL=1
> ./hydro 250 50
Cycles per element for solvers: 735.97
> maqao lprof xp=sx2 -- ./hydro 250 50
> maqao lprof xp=sx2 -dl | head
#####
# Loop ID | Function Name | Source Info | Level | Time (%)
#####
# 148     | linearSolver2 | 14,167@kernel.c | Innermost | 57.14
# 56      | c_densitySolver | 14,167@kernel.c | Innermost | 19.64
# 79      | c_velocitySolver | 14,283@kernel.c | Innermost | 3.57

> maqao cqa hydro loop=148 --confidence-levels=gain,potential,hint
```

Remark: less calls were unrolled since linearSolver is now much more bigger

CQA output for kernel1

Matching between your loop ...

The binary loop is composed of 96 FP
arithmetical operations:

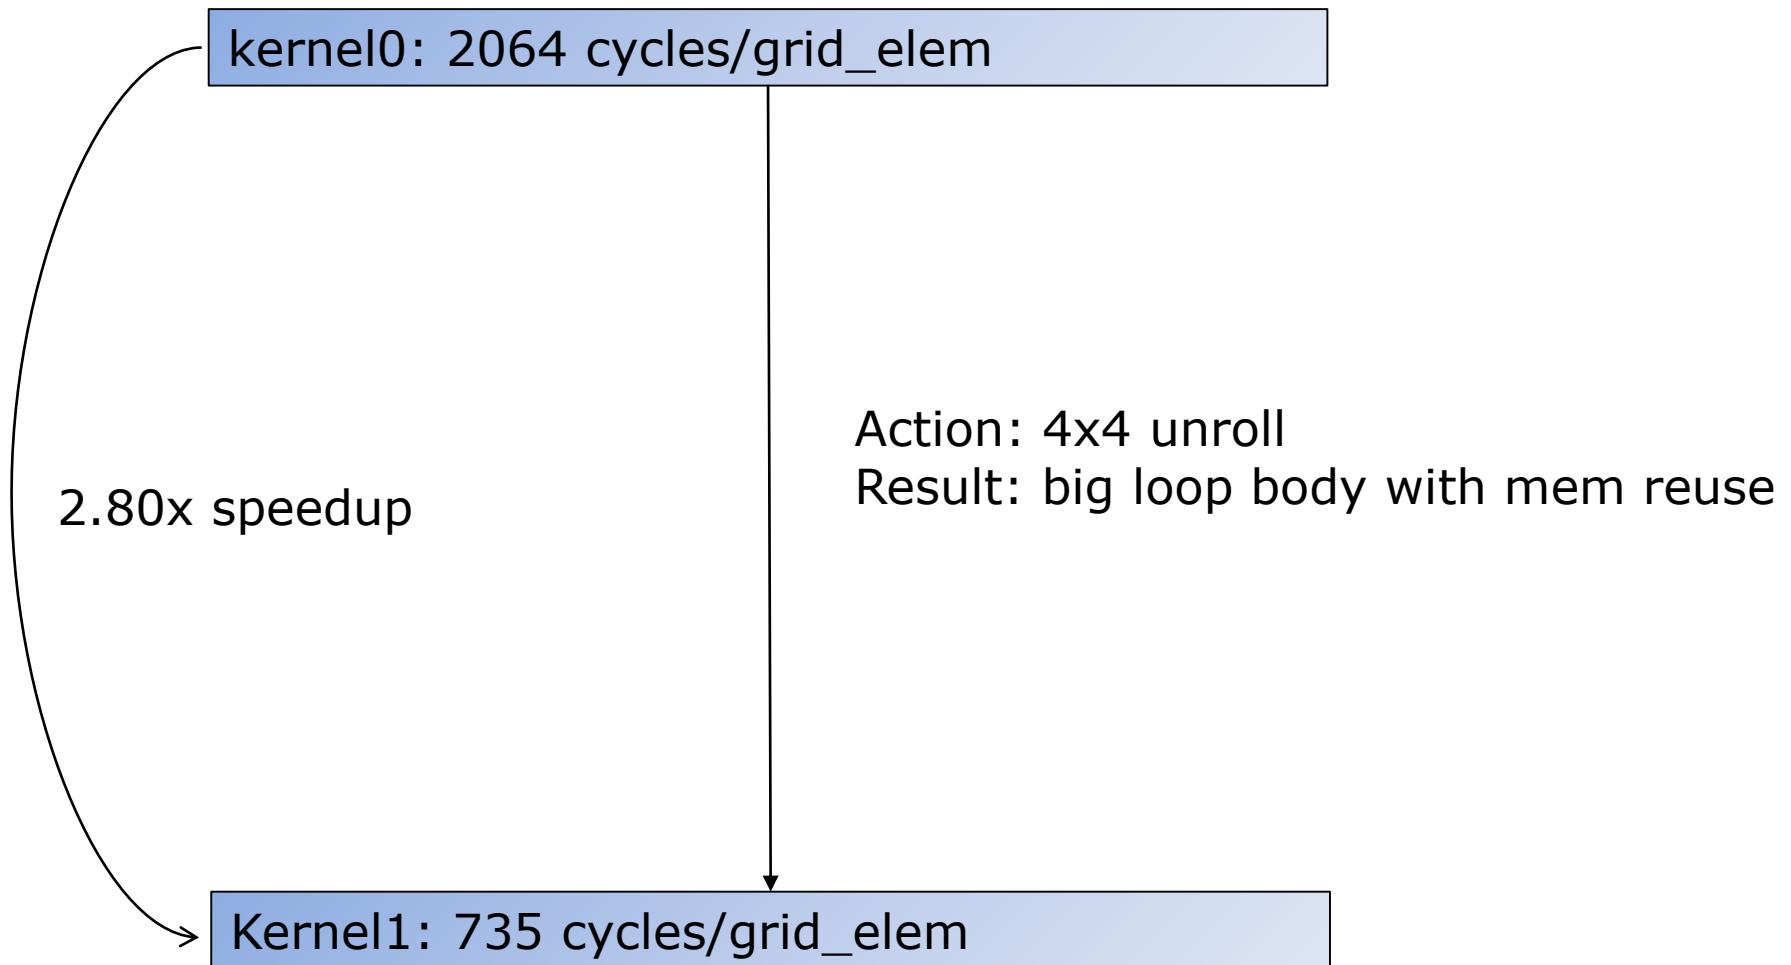
- 64: addition or subtraction
- 32: multiply

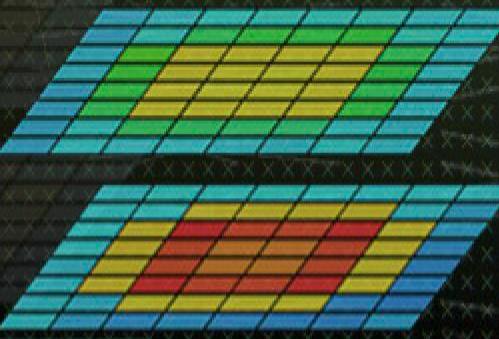
The binary loop is loading 272 bytes
(68 single precision FP elements).
The binary loop is storing 64 bytes
(16 single precision FP elements).

4x4 Unrolling was applied

Expected 48... But still better
than 80

Summary of optimizations and gains





MAQAO ONE View Demo

Setup

Build bt-mz_A.8

```
$> cd $WORK/NPB3.3-MZ-MPI  
$> make bt-mz CLASS=A NPROCS=8
```

Load MAQAO

```
$> module load UNITE maqao/2.2.1  
UNITE loaded  
maqao/2.2.1 loaded
```

Set number of threads

```
$> export OMP_NUM_THREADS=6
```

Prepare ONE View

Check mpi wrapper

```
$> echo $MPIEXEC  
/opt/MPI/bin/mpieexec
```

Edit the ONE View configuration file

```
$> vim $WORK/MAQAO_HANDSON/oneview/config_bt.lua
```

```
...  
binary          = "NPB3.3-MZ-MPI/bin/bt-mz_A.8"  
...  
mpi_command    = "/opt/MPI/bin/mpieexec -n 8 -m 4"  
...
```

Launching MAQAO ONE View and viewing results

Launch ONE View

```
$> cd ${WORK}
$> maqao oneview --create-report=one --
config=${WORK}/MAQAO_HANDSON/oneview/config_bt.lua --format=html
```

The experiment directory is generated in the current directory (`${WORK}`). By default, it is named `maqao_<timestamp>` (with `<timestamp>` in the form of `YYYY-mm-DD_HH_MM_ss`).

A sample experiment directory is located in

`MAQAO_HANDSON/oneview/sample_xpdir`

Results can be visualized using firefox

```
$> firefox <exp-dir>/RESULTS/one_html/index.html
```