MAQAO
Hands-on exercises
LRZ Cluster

LProf: lightweight generic profiler
LProf/MPI: Lightweight MPI oriented profiler
CQA: code quality analyzer

Setup

Copy handson material

```
> cp /home/hpc/a2c06/lu23bud/LRZ-VIHPSTW21/tools/maqao/MAQAO_HANDSON_LRZ.tar.xz $HOME
```

Untar the archive at the root of your HOME folder

```
> cd $HOME
> tar xf MAQAO_HANDSON_LRZ.tar.xz
> cd MAQAO_HANDSON
```

Copy MPI GUI

```
> scp –r MPI_GUI my_machine:
```

Add MAQAO path to your local path

```
> source ./scripts/env.sh
```

MAQAO LProf
Hands-on exercises

Jean-Baptiste LE-RESTE

Setup

Go to the Handson folder

```
> cd $HOME/MAQAO_HANDSON
```

Locate script and modify it as needed

```
> vim scripts/lprof_bt-mz_ompi.5P.2T.sh
```

Launch your job

```
> sbatch scripts/lprof_bt-mz_ompi.5P.2T.sh
```

Visualize the results

```
> ./scripts/display.sh (then follow the instructions displayed in the terminal)
```

Using MAQAO LProf

A copy of the output is located in output_examples/LProf folder in case you experience a problem

Now follow live demo/comments

MAQAO LProfMPI Hands-on exercises

Jean-Baptiste LE-RESTE

Setup

Go to the Handson folder

```
> cd $HOME/MAQAO_HANDSON
```

Locate script and modify it as needed

```
> vim scripts/lprof-mpi_bt-mz_intel.5P.2T.sh
```

Launch your job

```
> sbatch scripts/bt-mz_ompi.5P.2T_lprof_mpi.sh
```

Visualize the results

Copy/Paste the directory ./MPI_GUI and the file ./results/MPI_Profile.js locally.
Then open the MPI_GUI/res/MPI.html and load the MPI_Profile.js file.

MAQAO / CQA
Hands-on exercises

Emmanuel OSERET

Setup for UV2 (Westmere)

Login uv2

```
> ssh uv2
```

Load a recent GCC compiler

```
> module load gcc/5
```

Switch to CQA handson folder

```
> cd $HOME/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

```
> make OPTFLAGS=-O3 KERNEL=0
> ./matmul 100 1000
Cycles per FMA: 3.27
> maqao cqa matmul fct-loops=kernel0 [of=html]
```

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 requesting CQA to analyze all innermost loops in functions (max 2-3 loops/function for this hands-on).

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 0.75 cycles (4.00x
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 stide 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):

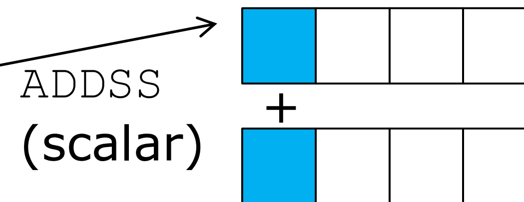

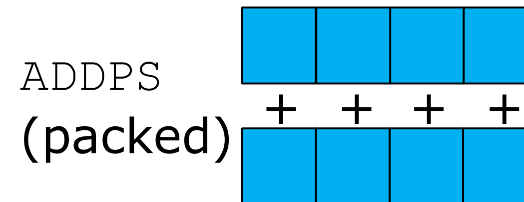

- Accesses are not contiguous => let’s permute k and j loops
- No structures here…

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(i) a[i].x = b[i].x; (slow, non stride 1) => for(i) a.x[i] = b.x[i]; (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

	j=0,1…							
i=0	a	b	c	d	e	f	g	h
i=1	i	j	k	l	m	n	o	p

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]);
```

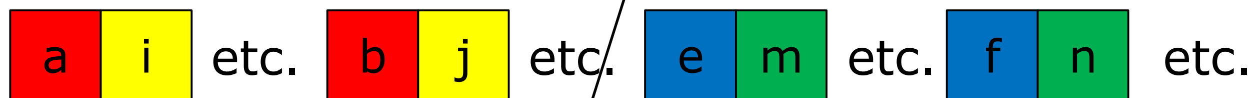


```
for (i=0; i<n; i++)
  for (j=0; j<n; j++)
    f(a[i][j]);
```

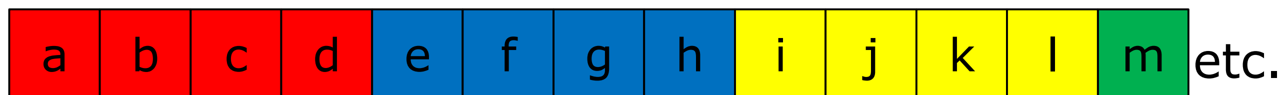

Removing inter-iteration dependences and getting stride 1 by permuting loops on j and k

```
void kernel1 (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

```
> make clean
> make OPTFLAGS=-O3 KERNEL=1
> ./matmul 100 1000
Cycles per FMA: 1.03
> maqao cqa matmul fct-loops=kernel1 --confidence-
levels=gain,potential,hint
```

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

```
Vectorization status
--------------------
Your loop is fully vectorized…

Vector unaligned load/store instructions
----------------------------------------
 - 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

Aligning vector accesses in driver + assuming them in kernel

```
int main (…) {
   (…)
#if KERNEL==2
  puts (« driver.c: Using
posix_memalign instead of malloc »);
  posix_memalign ((void **) &a, 32,
size_in_bytes);
  posix_memalign ((void **) &b, 32,
size_in_bytes);
  posix_memalign ((void **) &c, 32,
size_in_bytes);
#else
  a = malloc (size_in_bytes);
  b = malloc (size_in_bytes);
  c = malloc (size_in_bytes);
#endif
   (…)
}
```

```
void kernel2 (int n,
              float a[n][n],
              float b[n][n],
              float c[n][n]) {
  int i, j, k;

  for (i=0; i<n; i++) {
    float *ci =
__builtin_assume_aligned (c[i], 32);
    for (j=0; j<n; j++)
      ci[j] = 0.0f;
    for (k=0; k<n; k++) {
      float *bk =
__builtin_assume_aligned (b[k], 32);
      for (j=0; j<n; j++)
        ci[j] += a[i][k] * bk[j];
  }
}
```

kernel2: assuming aligned vector accesses

```
> make clean
> make OPTFLAGS=-O3 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: 0.92
```

Summary of optimizations and gains

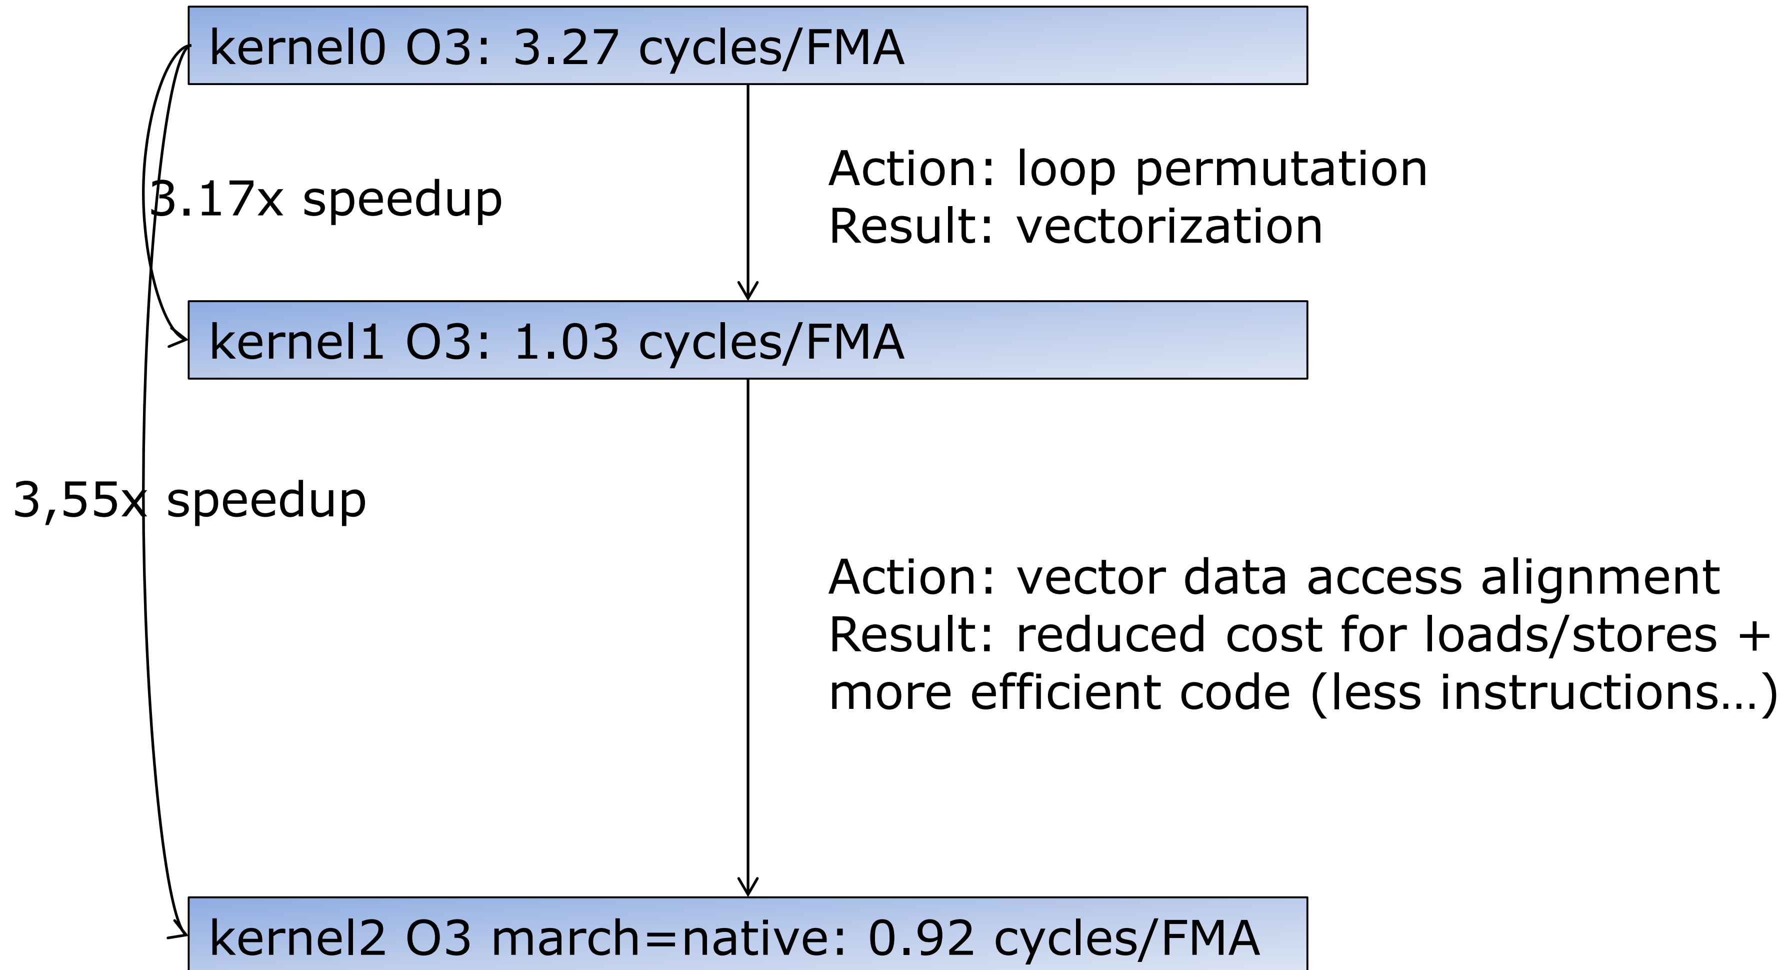

What if Haswell ?

CQA can cross-analyze to another micro-architecture

Compiling and analyzing kernel0 in –O3

```
> make OPTFLAGS=-O3 KERNEL=0
> maqao cqa matmul uarch=HASWELL fct-loops=kernel0
```

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). (…)
```

- 8x instead of 4x

kernel1: loop interchange

```
> make clean
> make OPTFLAGS=-O3 KERNEL=1
> maqao cqa matmul uarch=HASWELL fct-loops=kernel1
```

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…
```

- Westmere: 100% length…
- Let's add –march=haswell 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

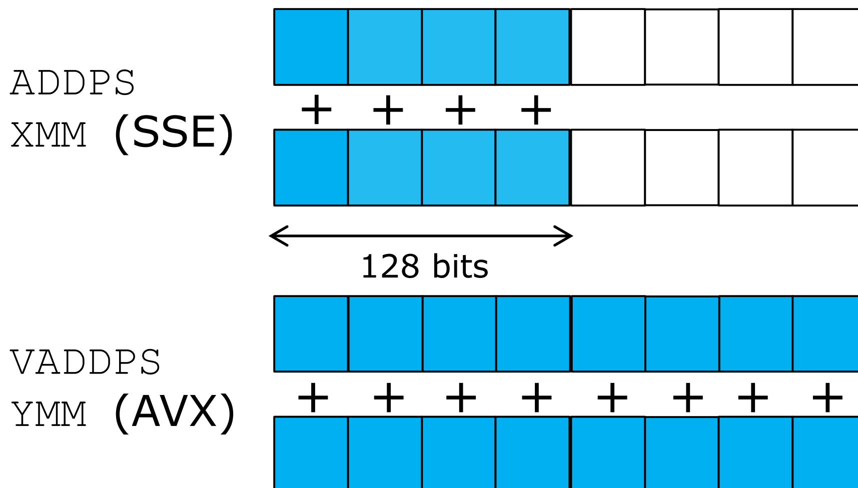

- FMA
  - Fused Multiply-Add (A+BC)
  - Intel architectures: supported on MIC/KNC and Xeon starting from Haswell

```
# A = A + BC

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

Kernel1 + -march=native

```
> make clean
> make OPTFLAGS="-O3 -march=haswell" KERNEL=1
> maqao cqa matmul uarch=HASWELL fct-loops=kernel1 --confidence-levels=gain,hint
```

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=haswell" KERNEL=2
```

Setup for an Haswell node

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

Login lxlogin5/6

```
> ssh <your_login>@lxlogin5.lrz.de
```

Load MAQAO environment

```
> module load maqao
```

Load a recent GCC compiler

```
> module load gcc/5
```

Switch to CQA handson folder

```
> cd $HOME/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

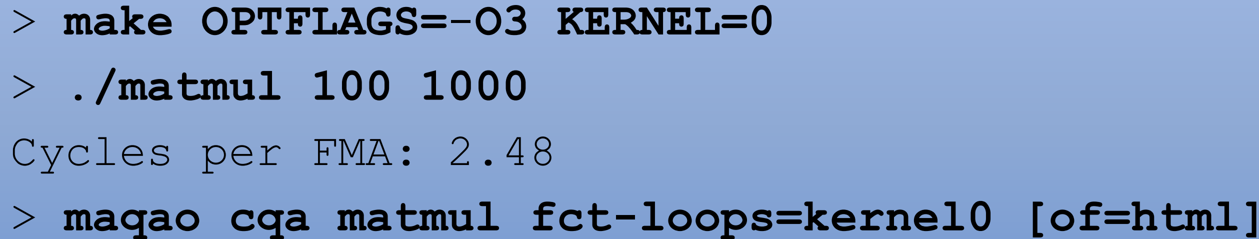

NB: the usual way to use CQA consists in finding IDs of hot loops with the MAQAO profiler and forwarding them to CQA (loop=17,42…).
To simplify this hands-on, we will bypass profiling and directly requesting CQA to analyze all innermost loops in functions (max 2-3 loops/function for this hands-on).

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 stide 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):

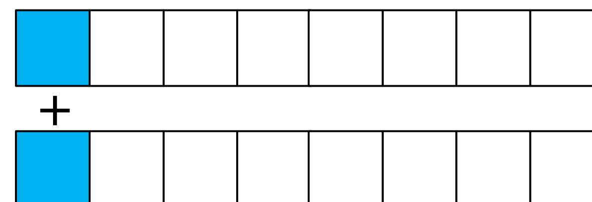

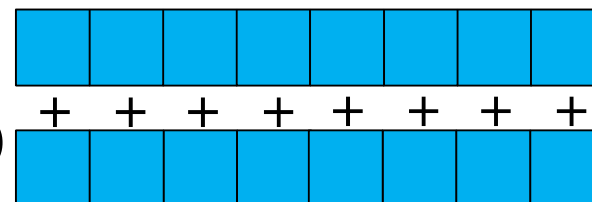

- Accesses are not contiguous => let’s permute k and j loops
- No structures here…

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(i) a[i].x = b[i].x; (slow, non stride 1) => for(i) a.x[i] = b.x[i]; (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

	j=0,1…							
i=0	a	b	c	d	e	f	g	h
i=1	i	j	k	l	m	n	o	p

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]);
```

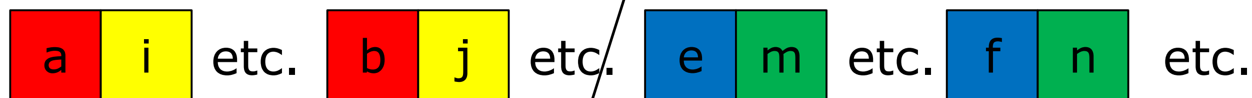


```
for (i=0; i<n; i++)
  for (j=0; j<n; j++)
    f(a[i][j]);
```

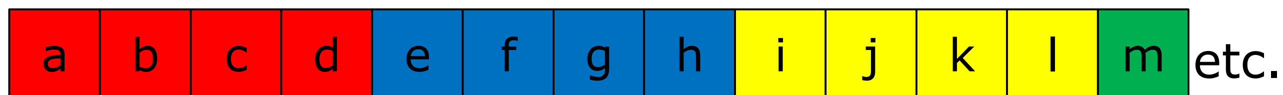

Removing inter-iteration dependences and getting stride 1 by permuting loops on j and k

```
void kernel1 (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

```
> make clean
> make OPTFLAGS=-O3 KERNEL=1
> ./matmul 100 1000
Cycles per FMA: 0.65
> maqao cqa matmul fct-loops=kernel1
```

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 (hornet)

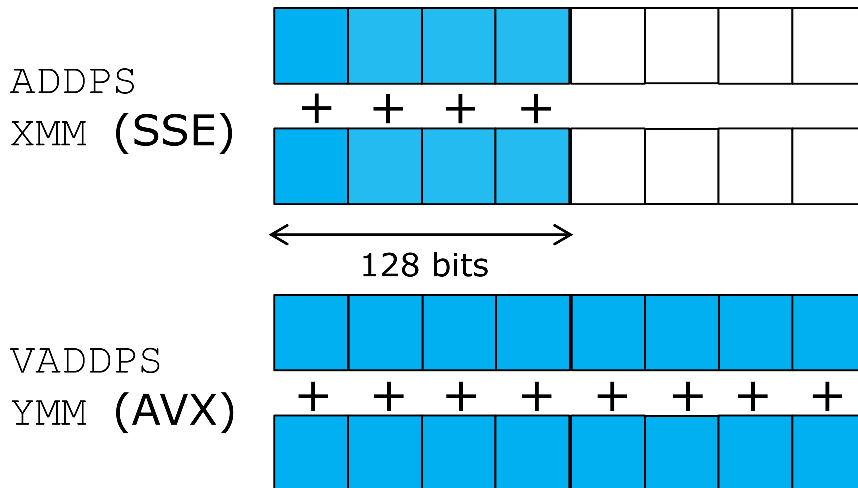


```
# A = A + BC

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

Kernel1 + -march=native

```
> make clean
> make OPTFLAGS="-O3 -march=native" KERNEL=1
> ./matmul 100 1000
Cycles per FMA: 0.47
> maqao cqa matmul fct-loops=kernel1 --confidence-levels=gain,hint
```

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: 0.35
```

Summary of optimizations and gains

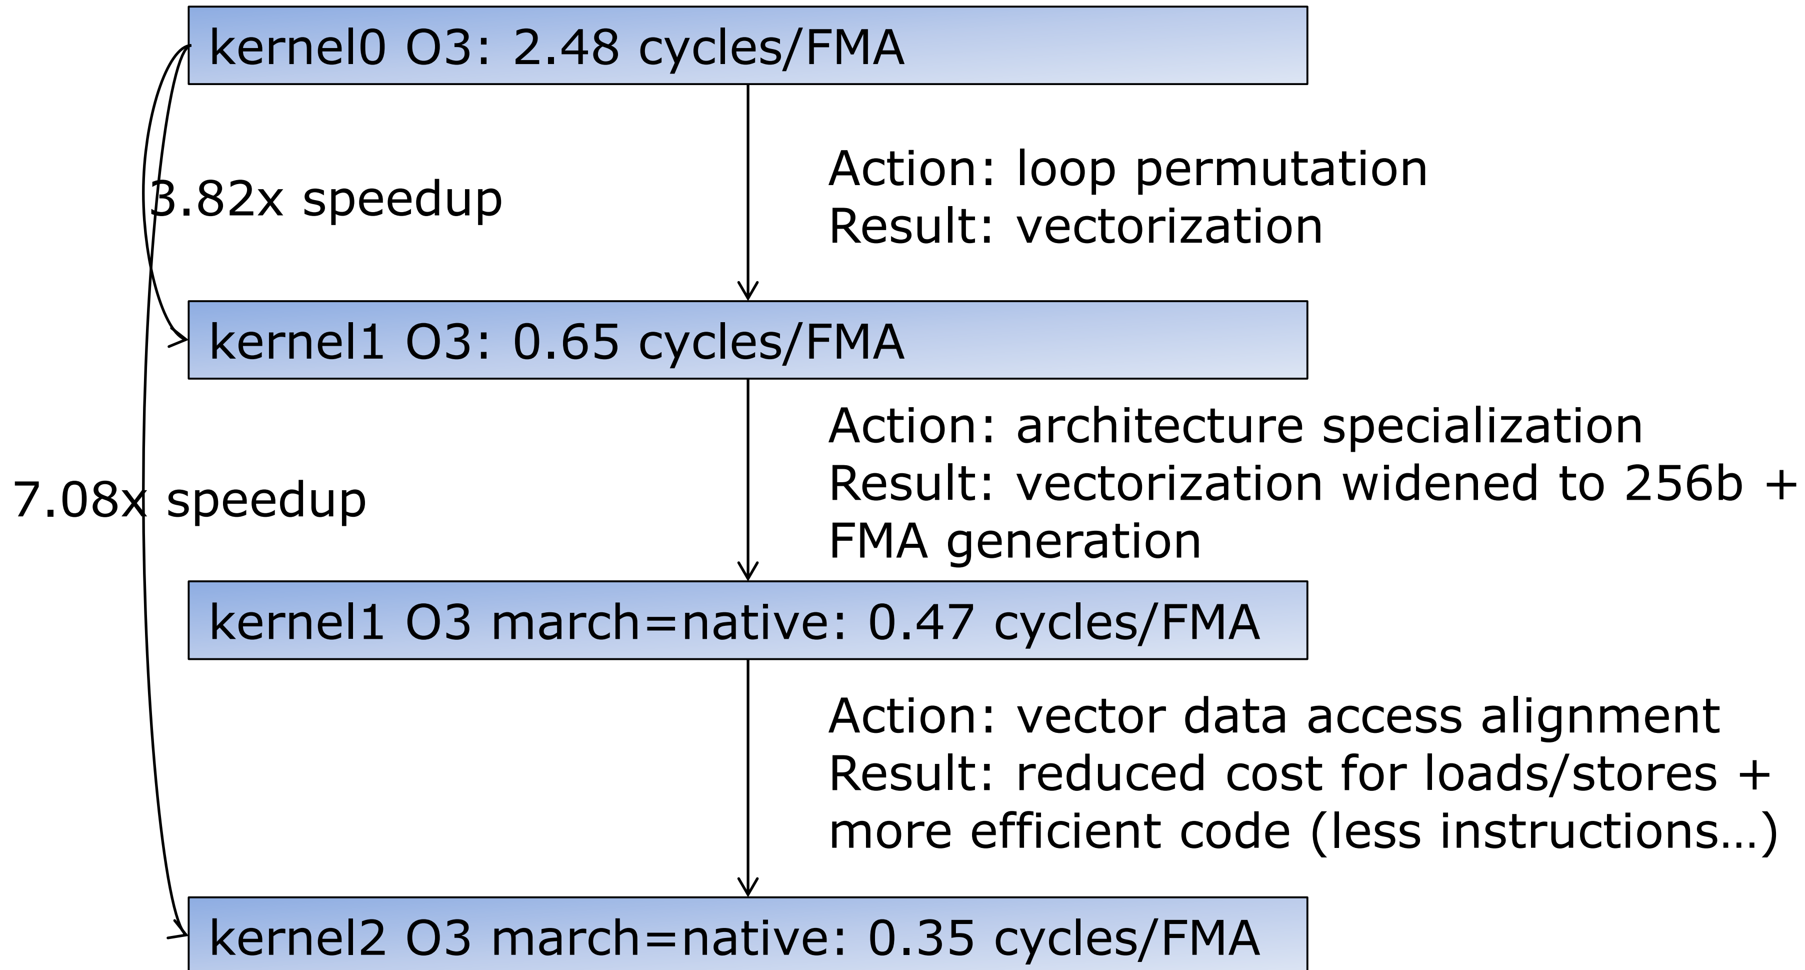

Hydro example

Switch to the other CQA handson folder

```
> cd $HOME/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

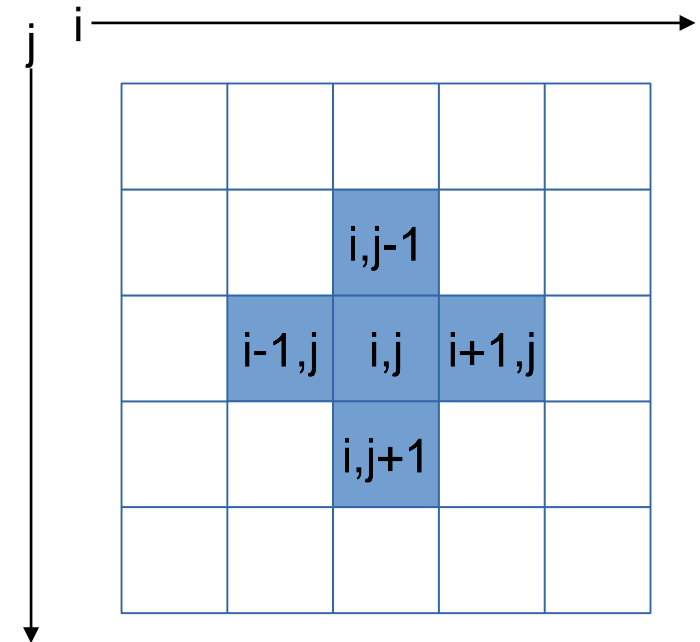

Compiling, running and analyzing kernel0 (icc -O3 -xHost)

```
> make KERNEL=0
> ./hydro 250 10 # 1st param: grid_size and 2nd param: repet nb
Cycles per element for solvers: 2064.14
> maqao lprof xp=sx -- ./hydro 250 10
> maqao lprof xp=sx -dl | head
################################################################################
#  Loop ID   |    Function Name     |     Source Info     |    Level     |  Time (%)
################################################################################
#  142       |  project             |  103,105@kernel.c   |  Innermost   |  28.29
#  54        |  c_densitySolver     |  103,105@kernel.c   |  Innermost   |  23.03
#  94        |  c_velocitySolver    |  103,105@kernel.c   |  Innermost   |  21.71
#  87        |  c_velocitySolver    |  103,105@kernel.c   |  Innermost   |  19.08
#  140       |  project             |  371,374@kernel.c   |  Innermost   |  1.32

> maqao cqa hydro loop=142
```

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 need analysis.

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

```
Bottlenecks
-----------
The divide/square root unit is a
bottleneck.
By removing all these bottlenecks, you
can lower the cost of an iteration
from 7.00 to 5.00 cycles (1.40x
speedup).
Try to reduce the number of division
or square root instructions.
If denominator is constant over
iterations, use reciprocal (replace
x/y with x*(1/y)). Check precision
impact. This will be done by your
compiler with no-prec-div or Ofast.
```

Cost of a division (about 10-50 cycles) is much higher than for a addition or a multiplication (typically 1 cycle per instruction), especially on double precision elements and on older processors

Removing (hoisting) division

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

void linearSolver0 (...) {
  int i, j, k;
  const float inv_c = 1.0f / c;

  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)]
  ) * inv_c;
}
```

Dividing by c is equivalent to multiplying by (1/c). Since c is constant in the loop, the divide-by-c operation was hoisted out of the loop and replaced by a multiply inside it

Remark : in some cases applying this optimization can change numerical results (due to rounding). That is why it is not applied by default by compilers.
On this example, no difference when comparing 6 first digits of the decimal part

kernel1: division removal

```
> make clean
> make KERNEL=1
> ./hydro 250 10
Cycles per element for solvers: 1824.16
> maqao lprof xp=sx2 -- ./hydro 250 50
> maqao lprof xp=sx2 -dl | head
(output similar to KERNEL=0, with same loop Ids)
> maqao cqa loop=142
```

CQA output for kernel1

```
Composition and unrolling
-------------------------
It is composed of the loop 142
and is not unrolled or unrolled with
no peel/tail loop.
The analysis will be displayed for
the requested loops: 142

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…
```

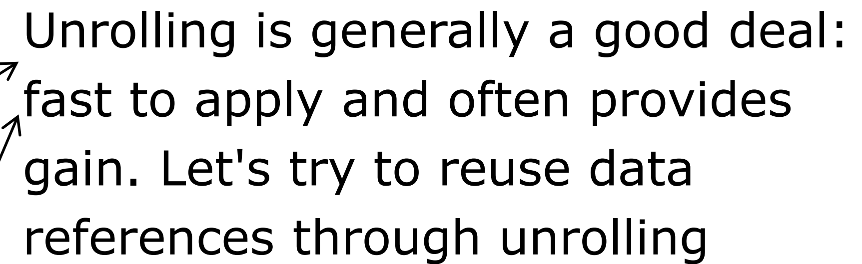

Memory references reuse : 4x4 unroll footprint on loads

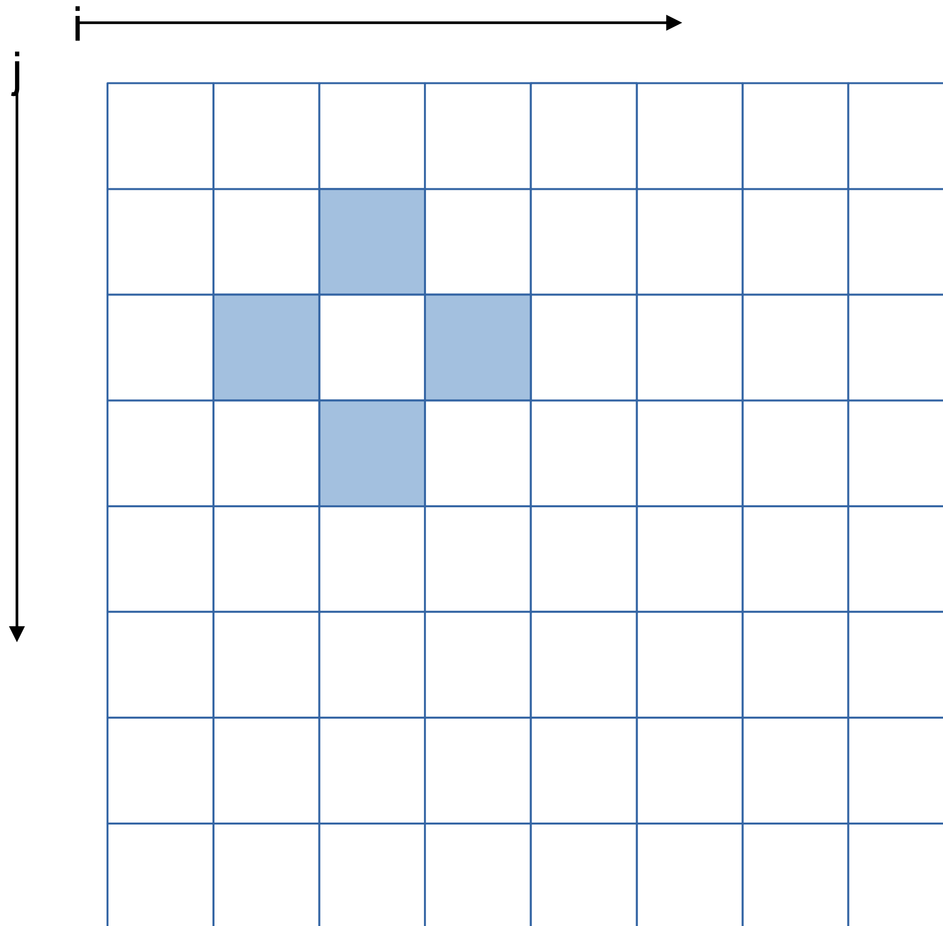

LINEAR_SOLVER(i+0,j+0)

Memory references reuse : 4x4 unroll footprint on loads

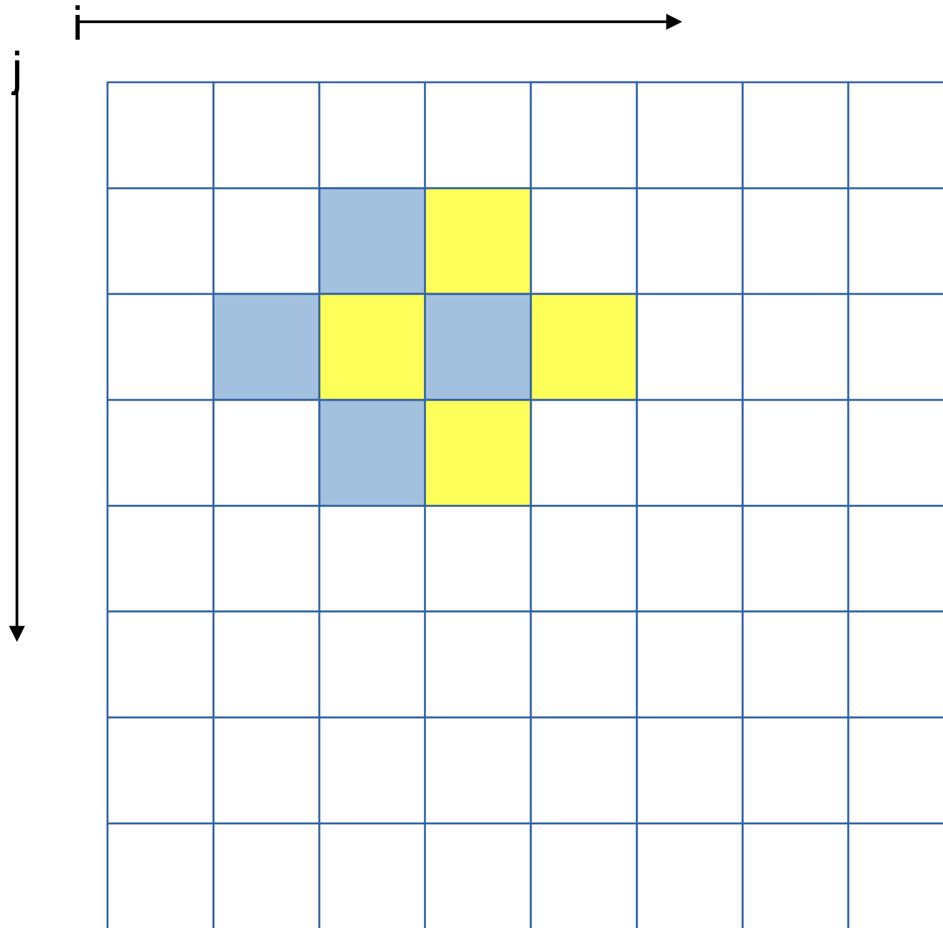

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)

Memory references reuse : 4x4 unroll footprint on loads

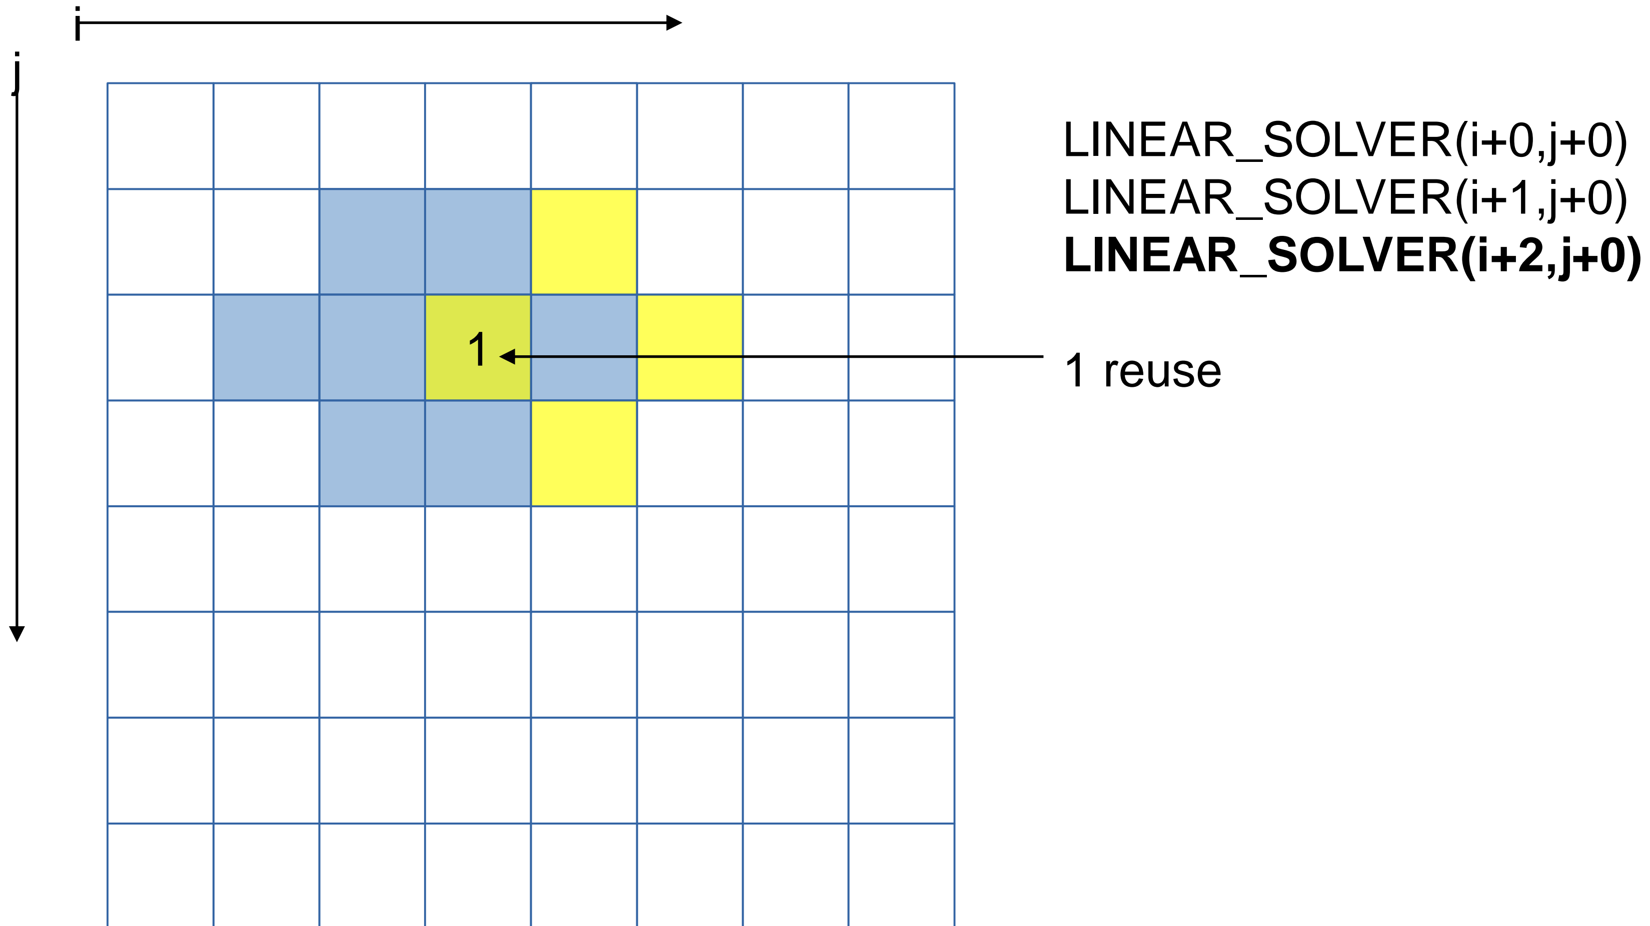

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)

Memory references reuse : 4x4 unroll footprint on loads

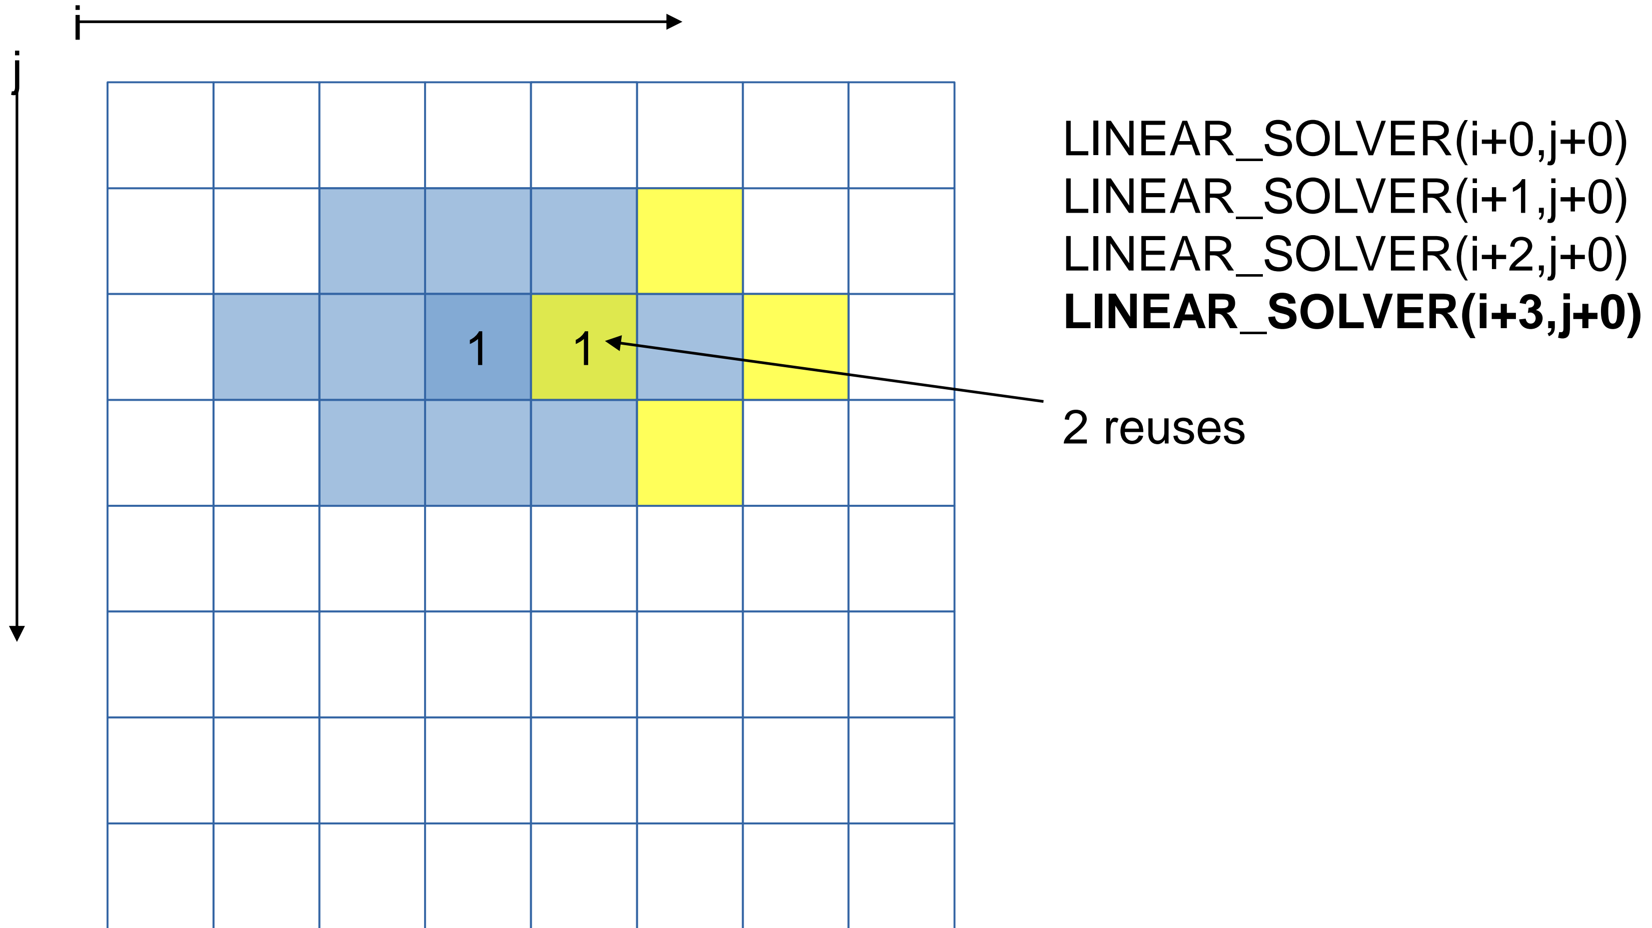

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

2 reuses

Memory references reuse : 4x4 unroll footprint on loads

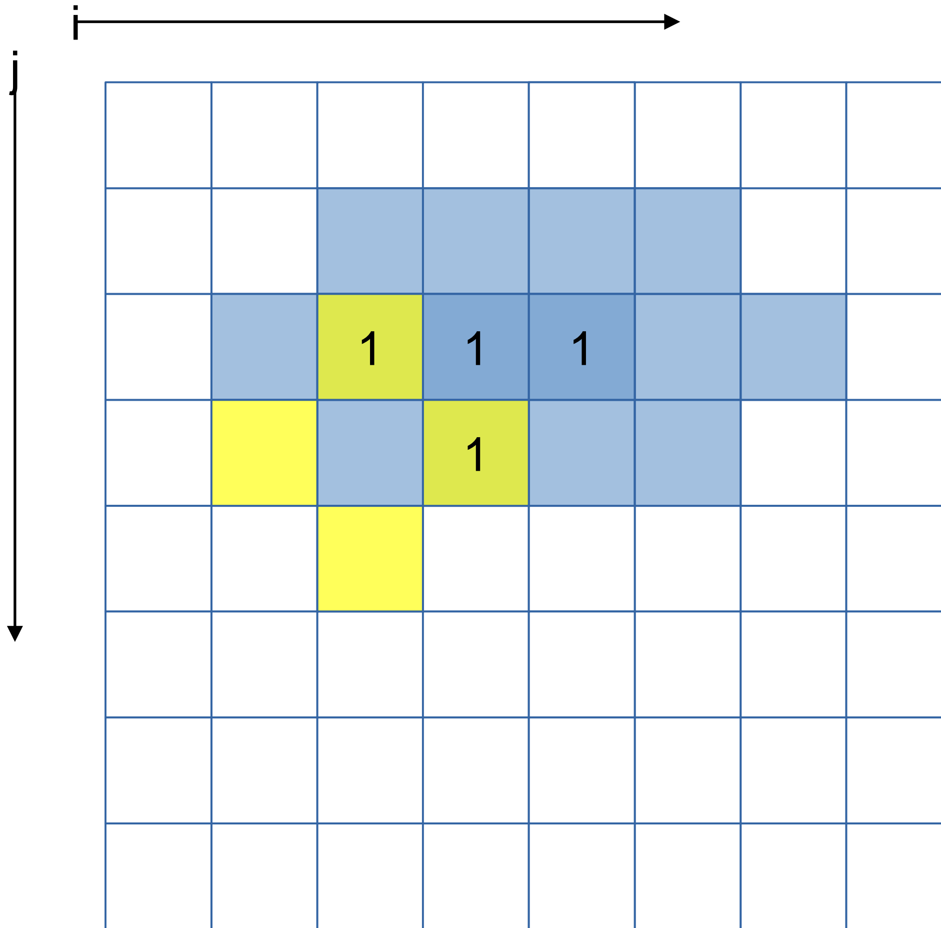

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

LINEAR_SOLVER(i+0,j+1)

4 reuses

Memory references reuse : 4x4 unroll footprint on loads

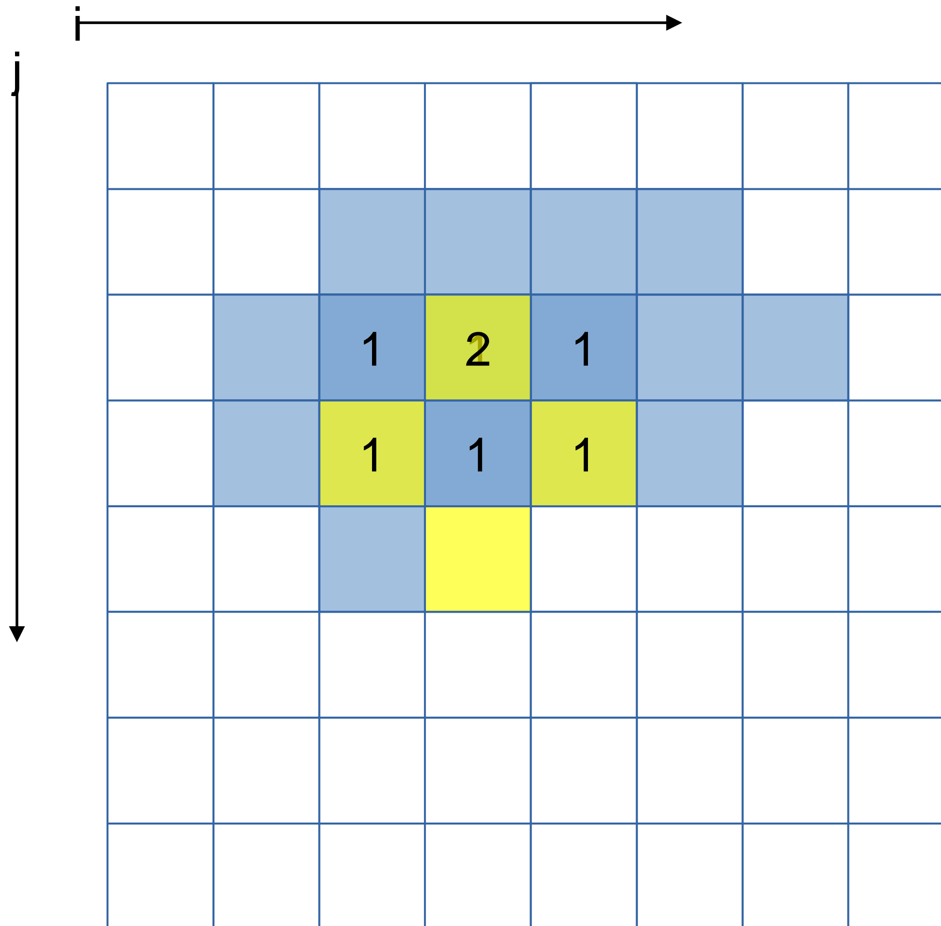

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

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

7 reuses

Memory references reuse : 4x4 unroll footprint on loads

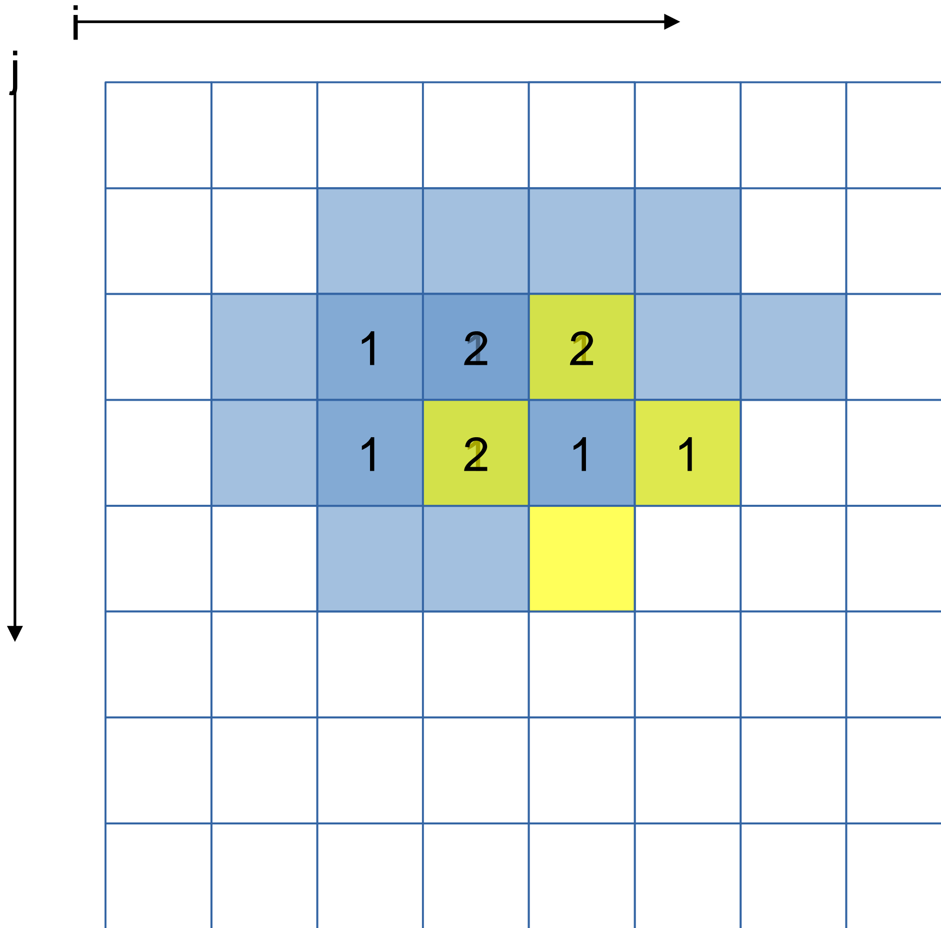

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

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

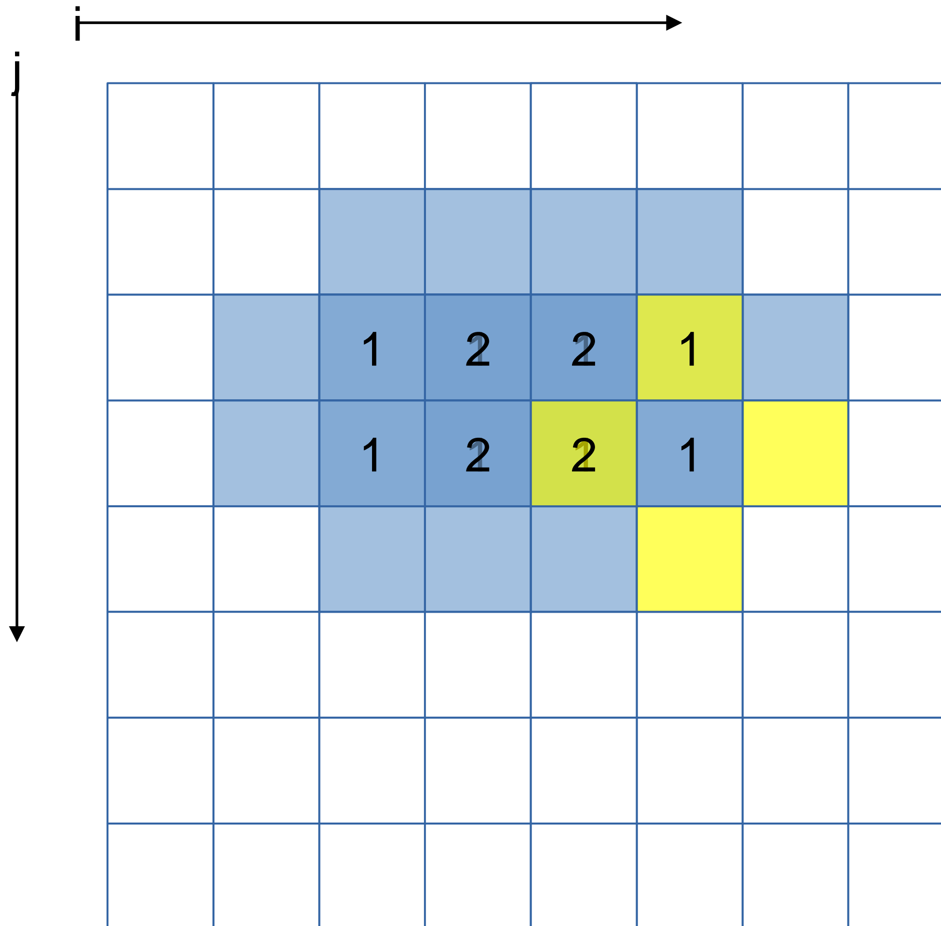

LINEAR_SOLVER(i+0,j+0)
LINEAR_SOLVER(i+1,j+0)
LINEAR_SOLVER(i+2,j+0)
LINEAR_SOLVER(i+3,j+0)

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

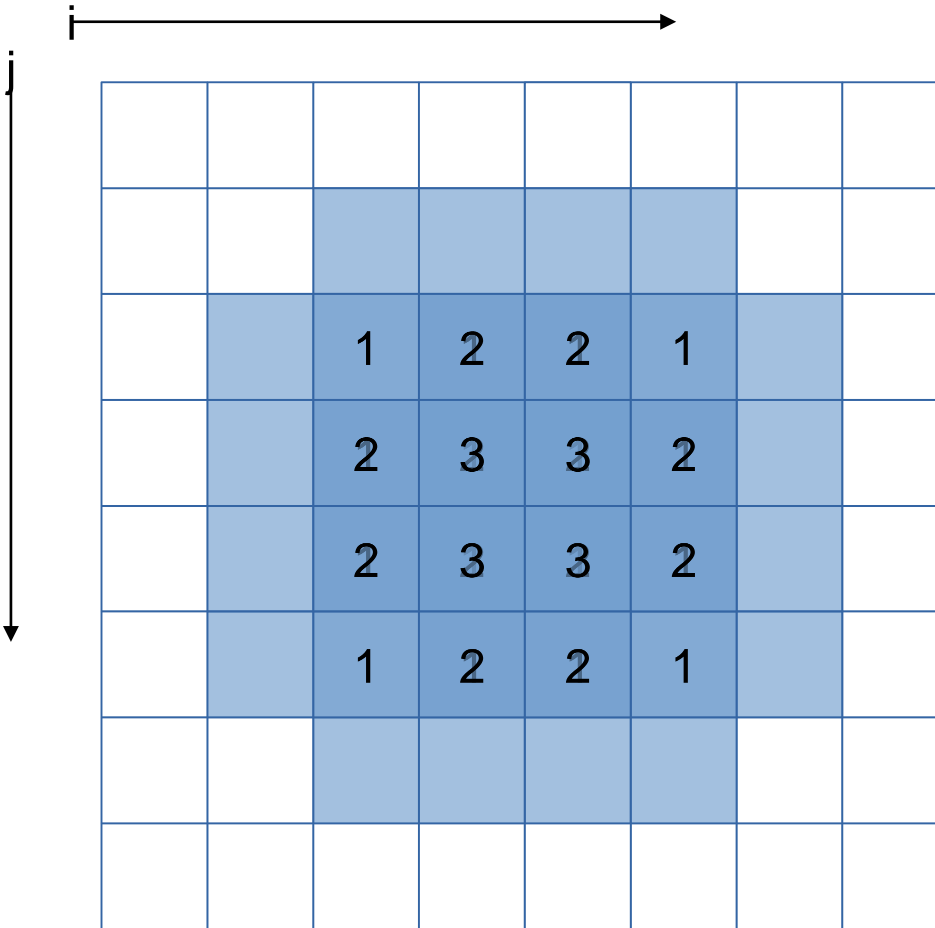

LINEAR_SOLVER(i+0-3,j+0)

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 4x4x4 = 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

Kernel2

```
> make clean
> make KERNEL=2
> ./hydro 250 10
Cycles per element for solvers: 735.97
> maqao lprof xp=sx3 -- ./hydro 250 50
> maqao lprof xp=sx3 -dl | head
################################################################################
#  Loop ID  |    Function Name    |     Source Info     |    Level     |  Time (%)
################################################################################
#  143      |  linearSolver2      |  14,167@kernel.c    |  Innermost   |  57.14
#  55       |  c_densitySolver    |  14,167@kernel.c    |  Innermost   |  19.64
#  76       |  c_velocitySolver   |  14,283@kernel.c    |  Innermost   |  3.57

> maqao cqa hydro loop=143
```

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

CQA output for kernel2

```
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 were applied

Expected 48... But still better than 80

Summary of optimizations and gains