

MAQAO Hands-on exercises

Profiling bt-mz
Optimising a code

Setup

Login to the cluster with X11 forwarding

```
[LOCAL] ssh -Y <login>@lxlogin[1-4].lrz.de
```

Load VIHPS (incl. MAQAO) environment

```
> export TW45=/lrz/sys/courses/vihps/2024
> export WORK=$SCRATCH # assumed fastest filesystem
Hint: copy in ~/.profile
```

Copy handson material to your WORK directory

```
> cd $WORK
> tar xf $TW45/material/maqao/MAQAO_HANDSON.tgz
> tar xf $TW45/material/maqao/NPB3.4-MZ-MPI.tgz
```

(If not already done) Load MAQAO, compiler + MPI

```
> module use $TW45/modulefiles
> module load maqao/2.20.1
```

Setup (bt-mz compilation with debug symbols)

Ensure that the NAS are compiled with debug information (make.def)

```
> cd $WORK/NPB3.4-MZ-MPI  
> cp $WORK/MAQAO_HANDSON/bt/make.def config
```

```
FFLAGS = -O3 -qopenmp -g -fno-omit-frame-pointer
```

Compile bt-mz with debug information

```
> make bt-mz CLASS=C
```

Running bt-mz

```
> cp $WORK/MAQAO_HANDSON/bt/bt.slurm bin  
> cd bin  
> sbatch bt.slurm
```

Profiling bt-mz with MAQAO

Cédric VALENSI
Emmanuel OSERET

Setup ONE View for batch execution

The ONE View configuration file must contain all variables for executing the application.

Retrieve the configuration file prepared for bt-mz in batch mode from the MAQAO HANDSON directory

```
> cd $WORK/NPB3.4-MZ-MPI/bin
> cp $WORK/MAQAO_HANDSON/bt/config_bt_oneview_sbatch.json .
> less config_bt_oneview_sbatch.json
```

```
"executable": "bt-mz.C.x"
...
"batch_script": "maqao_bt.slurm"
...
"batch_command": "sbatch <batch_script>"
...
"number_processes": 4
...
"number_nodes": 2
...
"mpi_command": "mpirun -n <number_processes>"
...
"envv_OMP_NUM_THREADS": 14
```


Review jobscript for use with ONE View

All variables in the jobscript defined in the configuration file must be replaced with their name from it.

Retrieve jobscript modified for ONE View from the MAQAO_HANDSON directory.

```
> cd $WORK/NPB3.4-MZ-MPI/bin
> cp $WORK/MAQAO_HANDSON/bt/maqao_bt.slurm .
> less maqao_bt.slurm
```

```
...
#SBATCH -N 2 <number_nodes>
#SBATCH -n 4 <number_processes>
#SBATCH -c 14 <OMP_NUM_THREADS>
...
export OMP_NUM_THREADS=14<OMP_NUM_THREADS>
...
srrun ./bt-mz-C.x
<mpi_command> <run_command>
...
```

Launch MAQAO ONE View on bt-mz (batch)

Launch ONE View

```
> cd $WORK/NPB3.4-MZ-MPI/bin
> maqao oneview --create-report=one \
-config=config_bt_oneview_sbatch.json -xp=ov_sbatch
```

The `-xp` parameter allows to set the path to the experiment directory, where ONE View stores the analysis results and where the reports will be generated.

If `-xp` is omitted, the experiment directory will be named `maqao_<timestamp>`.

WARNING:

- If the directory specified with `-xp` already exists, ONE View will reuse its content but not overwrite it.

Setup ONE View for scalability mode

Parameters for scalability mode are defined in `multirun_params`.

```
> less config_bt_oneview_sbatch.json
```

```
"executable": "bt-mz.C.x"  
"batch_script": "maqao_bt.slurm"  
"batch_command": "sbatch <batch_script>"  
"number_processes": 4  
"number_nodes": 2  
"mpi_command": "mpirun -n <number_processes>"  
"envv_OMP_NUM_THREADS": 14  
...  
"multiruns_params": [  
  { "name": "2P_1N", "number_nodes": 1, "number_processes": 2,  
    "number_processes_per_node": 2, "envv_OMP_NUM_THREADS": 14 },  
  { "name": "2P_2N", "number_nodes": 2, "number_processes": 2,  
    "number_processes_per_node": 1, "envv_OMP_NUM_THREADS": 14 },  
],  
"scalability_reference": "lowest-threads"
```


Launch MAQAO ONE View on bt-mz in scalability mode

Launch ONE View

```
> cd $WORK/NPB3.4-MZ-MPI/bin  
> maqao oneview --create-report=one --with-scalability=strong \  
-config=config_bt_oneview_sbatch.json -xp=ov_sbatch_scal
```

The `-xp` parameter allows to set the path to the experiment directory, where ONE View stores the analysis results and where the reports will be generated.

If `-xp` is omitted, the experiment directory will be named `maqao_<timestamp>`.

WARNING:

- If the directory specified with `-xp` already exists, ONE View will reuse its content but not overwrite it.

Display MAQAO ONE View results

The HTML files are located in `<exp-dir>/RESULTS/<binary>_one_html`, where `<exp-dir>` is the path of the experiment directory (set with `-xp`) and `<binary>` the name of the executable.

```
> firefox <exp-dir>/RESULTS/bt-mz.C.x_one_html/index.html
```

It is also possible to compress and download the results to display them:

```
> tar -zcf $HOME/ov_html.tgz <exp-dir>/RESULTS/bt-mz.C.x_one_html
```

```
[LOCAL] scp <login>@:lxlogin[1-4].lrz.de:ov_html.tgz .
```

```
[LOCAL] tar xf ov_html.tgz
```

```
[LOCAL] firefox <exp-dir>/RESULTS/bt-mz.C.x_one_html/index.html
```

Results can also be viewed directly on the console:

```
> maqao oneview -R1 -xp=<exp-dir> --output-format=text | less
```

The archive `MAQAO_HANDSON/bt/offline.tgz` contains a sample result directory.

Display MAQAO ONE View results using sshfs

- To install sshfs on Debian-based Linux distributions (like Ubuntu)

```
[LOCAL] sudo apt install sshfs
```

- Recommended to close a sshfs directory after use

```
[LOCAL] fusermount -u /path/to/sshfs/directory
```

Mount \$SCRATCH locally:

```
[LOCAL] mkdir coolmuc2_work  
[LOCAL] sshfs <login>@lxlogin[1-4].lrz.de:/gpfs/scratch/<project>/<user> \  
coolmuc2_work  
[LOCAL] firefox coolmuc2_work/NPB3.4-MZ-MPI/bin/ov_sbatch/RESULTS/bt-  
mz.C.x_one_html/index.html
```

Optimising a code with MAQAO

Emmanuel OSERET

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

Setup environment

Load MAQAO environment (if needed)

```
> module use $TW45/modulefiles  
> module load maqao/2.20.1
```

Load latest GCC compiler (if not already loaded)

```
> module load gcc/11.2.0
```

Analysing matrix multiply with MAQAO

Compile naïve implementation of matrix multiply

```
> cd $WORK/MAQAO_HANDSON/matmul
> make matmul_orig
```

Execution

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \
  ./matmul_orig/matmul 400 300
cycles per FMA: 3.40
```

Analyse matrix multiply with ONE View

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \
  maqao OV -R1 xp=ov_orig -- ./matmul_orig/matmul 400 300
```

OR

```
> maqao OV -R1 c=ov_orig.json xp=ov_orig
```

Viewing results (HTML)

```
> tar -zcf $HOME/ov_orig.tgz ov_orig/RESULTS/matmul_orig_one_html
```

```
[LOCAL] scp <login>@lxlogin[1-4].lrz.de:ov_orig.tgz .
```

```
[LOCAL] tar xf ov_orig.tgz
```

```
[LOCAL] firefox ov_orig/RESULTS/matmul_orig_one_html/index.html
```

⌘

Global Metrics ?

| | |
|--------------------------------------|-------|
| Total Time (s) | 25.30 |
| Profiled Time (s) | 25.30 |
| Time in analyzed loops (%) | 100 |
| Time in analyzed innermost loops (%) | 99.9 |
| Time in user code (%) | 100 |
| Compilation Options Score (%) | 50.0 |
| Array Access Efficiency (%) | 83.3 |

Potential Speedups

| | | |
|--|---------------------|------|
| Perfect Flow Complexity | Potential Speedup | 1.00 |
| Perfect OpenMP + MPI + Pthread | Potential Speedup | 1.00 |
| Perfect OpenMP + MPI + Pthread + Perfect Load Distribution | Potential Speedup | 1.00 |
| No Scalar Integer | Potential Speedup | 1.00 |
| | Nb Loops to get 80% | 1 |
| FP Vectorised | Potential Speedup | 2.51 |
| | Nb Loops to get 80% | 1 |
| Fully Vectorised | Potential Speedup | 8.00 |
| | Nb Loops to get 80% | 1 |
| FP Arithmetic Only | Potential Speedup | 1.00 |
| | Nb Loops to get 80% | 1 |

Viewing results (text)

```
> maqao OV -R1 -xp=ov_orig \  
  --output-format=text --text-global | less
```

```
+-----+  
+                               Global Metrics                               +  
+-----+  
  
Total Time:                      25.30 s  
Profiled Time:                    25.30 s  
Time spent in analyzed loops:     100 %  
Time spent in analyzed innermost loops: 99.9 %  
Time spent in user code:          100 %  
Compilation Options Score:        50  
Array Access Efficiency:          83.3 %  
  
Potential Speedups  
-----  
Perfect Flow Complexity:          1.00  
Perfect OpenMP + MPI + Pthread:   1.00  
Perfect OpenMP + MPI + Pthread + Load Distribution: 1.00
```

Viewing results (text)

```
> maqao oneview -R1 -xp=ov_orig \  
  --output-format=text --text-loops | less
```

```
+-----+  
+                1.1 - Top 10 Loops                +  
+-----+  
  
  Loop Id | Module | Source Location | Coverage (%) |  
-----+-----+-----+-----+  
  1      | matm... | kernel_orig.c:9-10 | 99.64 |  
  2      | matm... | kernel_orig.c:7-10 | 0.35 |  
  3      | matm... | kernel_orig.c:6-10 | 0.02 |
```



Loop ID

Viewing CQA output (text)

```
> maqao oneview -R1 -xp=ov_orig \  
  --output-format=text --text-cqa=1
```

Vectorization

Your loop is not vectorized.

8 data elements could be processed at once in vector registers.

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

Workaround

- Try another compiler or update/tune your current one:

- * recompile with `fassociative-math` (included in `Ofast` or `ffast-math`) to extend loop vectorization to FP reductions.

- 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):

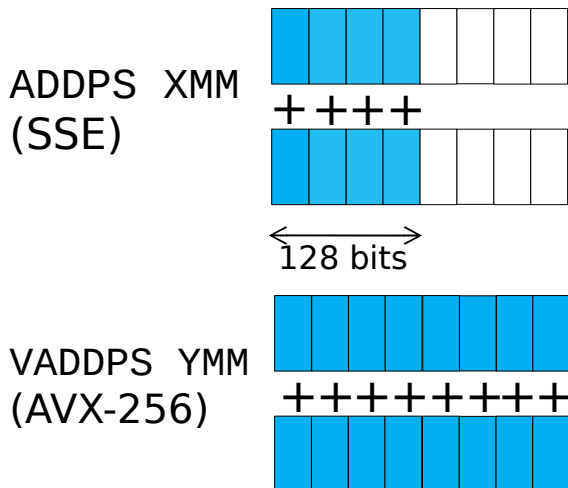
`for(i) a[i].x = b[i].x;` (slow, non stride 1) => `for(i) a.x[i] = b.x[i];` (fast, stride 1)

Loop ID

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 , <C>, %XMM0

VADDPS <A>, %XMM0, <A>

can be replaced with something like:

VFMADD312PS , <C>, <A>

Analyse matrix multiply with architecture specialisation

Compile architecture specialisation version of matrix multiply

```
> cd $WORK/MAQAO_HANDSON/matmul
> make matmul_opt
```

Execution

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \
./matmul_opt/matmul 400 300
cycles per FMA: 4.95
```

Analyse matrix multiply with ONE View

```
> maqao oneview -R1 c=ov_opt.json xp=ov_opt
```

CQA output for the arch-specialized kernel

| Global Metrics | | ? |
|--|---------------------|-------|
| Total Time (s) | | 36.73 |
| Profiled Time (s) | | 36.73 |
| Time in analyzed loops (%) | | 100.0 |
| Time in analyzed innermost loops (%) | | 99.9 |
| Time in user code (%) | | 100.0 |
| Compilation Options Score (%) | | 100 |
| Array Access Efficiency (%) | | 83.3 |
| Potential Speedups | | |
| Perfect Flow Complexity | | 1.00 |
| Perfect OpenMP + MPI + Pthread | | 1.00 |
| Perfect OpenMP + MPI + Pthread + Perfect Load Distribution | | 1.00 |
| No Scalar Integer | Potential Speedup | 1.00 |
| | Nb Loops to get 80% | 1 |
| FP Vectorised | Potential Speedup | 1.60 |
| | Nb Loops to get 80% | 1 |
| Fully Vectorised | Potential Speedup | 7.99 |
| | Nb Loops to get 80% | 1 |
| FP Arithmetic Only | Potential Speedup | 1.00 |
| | Nb Loops to get 80% | 1 |

SLOWDOWN (from 26.3s)

CQA output for the arch-specialized kernel

Vectorization

Your loop is not vectorized. 8 data elements could be processed at once in vector registers.



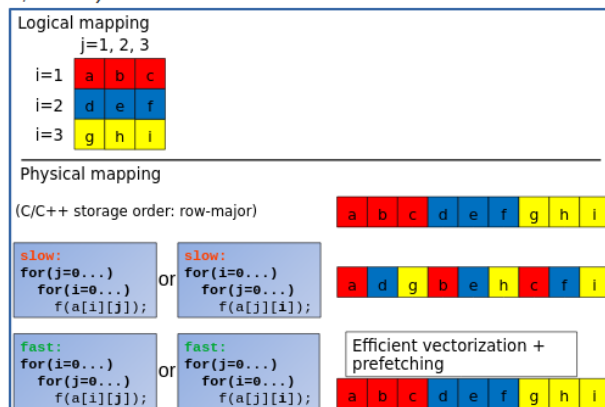
By vectorizing your loop, you can lower the cost of an iteration from 40.00 to 5.00 cycles (8.00x speedup).

Details

All SSE/AVX instructions are used in scalar version (process only one data element in vector registers). Since your execution units are vector units, only a vectorized loop can use their full power.

Workaround

- Try another compiler or update/tune your current one:
 - recompile with fassociative-math (included in Ofast or ffast-math) to extend loop vectorization to FP reductions.
- 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)

Impact of loop permutation on data access

Logical mapping

$j=0,1\dots$

| | | | | | | | | |
|-------|---|---|---|---|---|---|---|---|
| $i=0$ | 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];  
    }  
}
```

Analyse matrix multiply with permuted loops

Compile permuted loops version of matrix multiply

```
> cd $WORK/MAQAO_HANDSON/matmul
> make matmul_perm_opt
```

Execution

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \
./matmul_perm_opt/matmul 400 300
cycles per FMA: 0.46
```

Analyse matrix multiply with ONE View

```
> maqao oneview -R1 c=ov_perm_opt.json xp=ov_perm_opt
```

Loop permutation results

| Global Metrics | | ? |
|--|---------------------|------|
| Total Time (s) | | 3.48 |
| Profiled Time (s) | | 3.46 |
| Time in analyzed loops (%) | | 99.6 |
| Time in analyzed innermost loops (%) | | 90.5 |
| Time in user code (%) | | 99.6 |
| Compilation Options Score (%) | | 100 |
| Array Access Efficiency (%) | | 83.3 |
| Potential Speedups | | |
| Perfect Flow Complexity | | 1.00 |
| Perfect OpenMP + MPI + Pthread | | 1.00 |
| Perfect OpenMP + MPI + Pthread + Perfect Load Distribution | | 1.00 |
| No Scalar Integer | Potential Speedup | 1.04 |
| | Nb Loops to get 80% | 1 |
| FP Vectorised | Potential Speedup | 1.00 |
| | Nb Loops to get 80% | 1 |
| Fully Vectorised | Potential Speedup | 1.03 |
| | Nb Loops to get 80% | 1 |
| FP Arithmetic Only | Potential Speedup | 2.18 |
| | Nb Loops to get 80% | 2 |

Much faster (orig was 26.3s)

Much better (orig was close to 8)

CQA output after loop permutation

The screenshot displays a web interface for CQA (Code Quality Analyzer) output. At the top, there are four tabs: 'gain', 'potential', 'hint', and 'expert'. The 'gain' tab is selected. Below the tabs, there are three main sections:

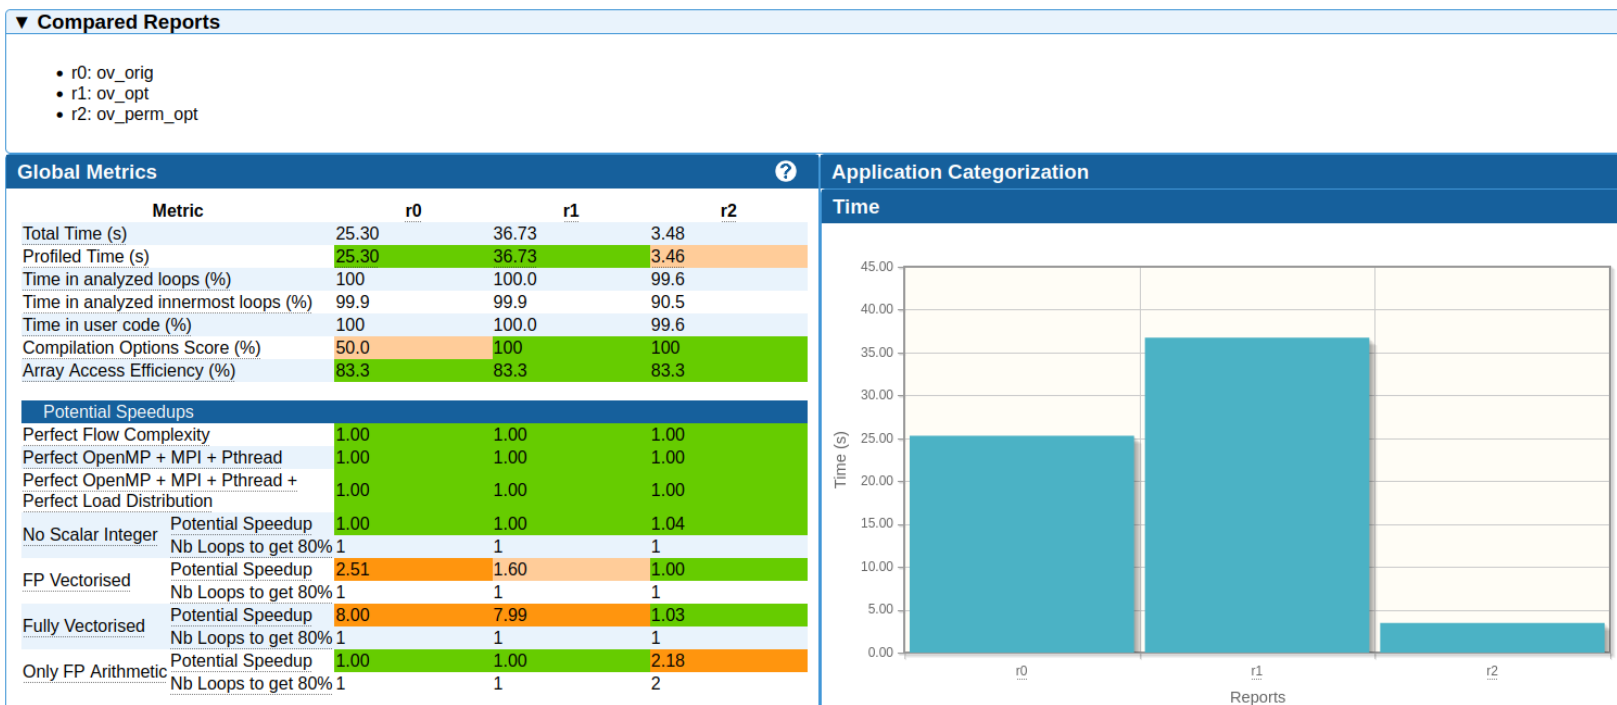
- Vectorization**: A blue header section containing the text: "Your loop is fully vectorized, using full register length." Below this is a sub-section titled **Details** with the text: "All SSE/AVX instructions are used in vector version (process two or more data elements in vector registers)."
- Execution units bottlenecks**: A blue header section containing the text: "Found no such bottlenecks but see expert reports for more complex bottlenecks."

Using comparison mode

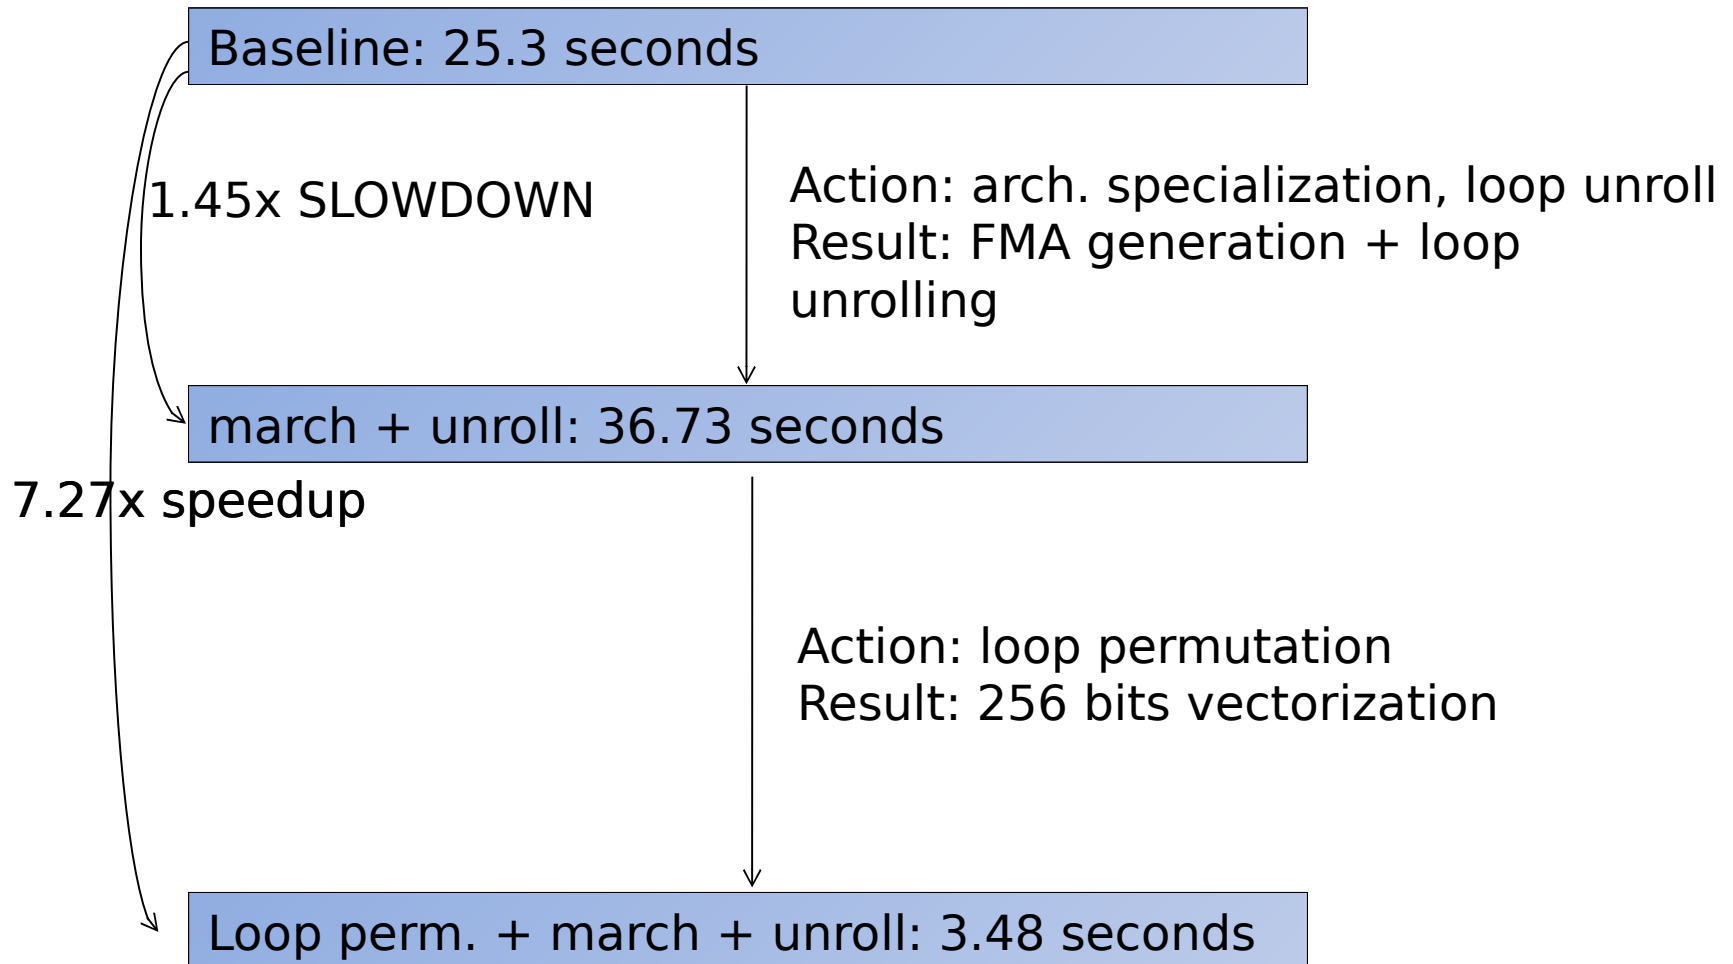
Generating a comparison report from experiment directories

```
> maqao oneview --compare-reports -xp=ov_matmul_cmp \
-inputs=ov_orig,ov_opt,ov_perm_opt
```

Open `ov_matmul_cmp/RESULTS/ov_matmul_cmp/index.html`



Summary of optimizations and gains



Hydro example

Switch to the hydro handson folder

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

Load MAQAO (if not already loaded)

```
> module use $TW45/modulefiles  
> module load maqao/2.20.1
```

Compile

```
> make
```


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



Running and analyzing original kernel (icx -O3 -xHost)

Execution

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \  
  ./hydro_orig 300 200  
Cycles per element for solvers: 1882.88
```

Analysing with MAQAO

```
> maqao OV -R1 xp=ov_orig c=ov_orig.json
```

OR

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \  
  maqao OV -R1 xp=ov_orig -- ./hydro_orig 300 200
```

```
> maqao OV -R1 xp=ov_orig \  
  --output-format=text --text-global --text-loops | less  
> maqao onview -R1 xp=ov_orig \  
  --output-format=text --text-global --text-cqa=15  
> ...  
> Total time: 13.22s
```

CQA output for original kernel

Workaround

- Try another compiler or update/tune your current one:
 - recompile with O2 or higher to enable loop vectorization and with fast-math (included in Ofast) to extend vectorization to FP reductions.
- 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)

Logical mapping

j=1, 2, 3
 i=1

| | | |
|---|---|---|
| a | b | c |
|---|---|---|

 i=2

| | | |
|---|---|---|
| d | e | f |
|---|---|---|

 i=3

| | | |
|---|---|---|
| g | h | i |
|---|---|---|

Physical mapping

(C/C++ storage order: row-major)

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| a | b | c | d | e | f | g | h | i |
|---|---|---|---|---|---|---|---|---|

slow:

```
for(j=0...)
  for(i=0...)
    f(a[i][j]);
```

OR

slow:

```
for(i=0...)
  for(j=0...)
    f(a[j][i]);
```

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| a | d | g | b | e | h | c | f | i |
|---|---|---|---|---|---|---|---|---|

fast:

```
for(i=0...)
  for(j=0...)
    f(a[i][j]);
```

OR

fast:

```
for(j=0...)
  for(i=0...)
    f(a[j][i]);
```

Efficient vectorization +
prefetching

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| a | b | c | d | e | f | g | h | i |
|---|---|---|---|---|---|---|---|---|

As for matmul, loops should be permuted.
CF build_index

Unroll opportunity

Loop is data access bound.

Workaround

Unroll your loop if trip count is significantly higher than target unroll factor and if some data references are common to consecutive iterations. This can be done manually. Or by recompiling with `-funroll-loops` and/or `-floop-unroll-and-jam`.


Consider loop unrolling

Kernel with loop permutation

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \  
./hydro_perm 300 200  
Cycles per element for solvers: 1579.47
```


```
> maqao oneview -R1 xp=ov_perm c=ov_perm.json  
> ...  
> Total time: 11.08s
```

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



Memory references reuse : 4x4 unroll footprint on loads



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

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

Kernel with manual 4x4 unroll and jam

```
> srun --cluster=cm2_tiny --reservation=hhps1s24 \  
./hydro_unroll 300 200  
Cycles per element for solvers: 764.50
```

```
> maqao oneview -R1 xp=ov_unroll c=ov_unroll.json  
> ...  
> Total time: 5.41s
```


CQA output for unrolled kernel

Matching between your loop (in the source code) and the binary loop

The binary loop is composed of 96 FP arithmetical operations:

- 64: addition or subtraction (16 inside FMA instructions)
- 32: multiply (16 inside FMA instructions)

The binary loop is loading 260 bytes (65 single precision FP elements). The binary loop is storing 64 bytes (16 single precision FP elements).

4x4 Unrolling were applied

Lower than 80: 64 (from x) + 16 (from x0)

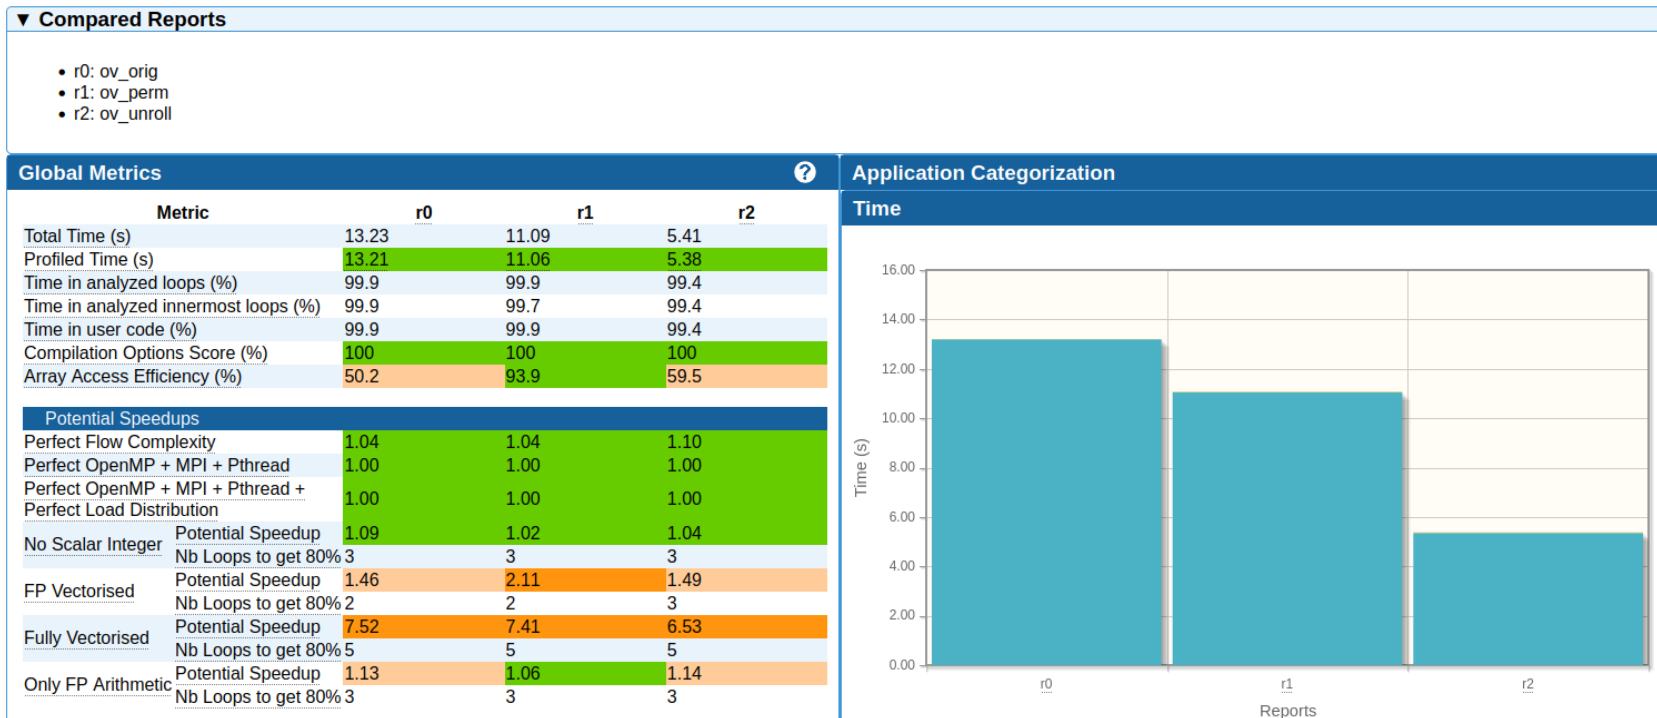
Now divides appears:
compiler failed to remove
common divide across
macros

Using comparison mode

Generating a comparison report from experiment directories

```
> maqao oneview --compare-reports -xp=ov_hydro_cmp \
-inputs=ov_orig,ov_perm,ov_unroll
```

Open `ov_hydro_cmp/RESULTS/ov_hydro_cmp/index.html`



Summary of optimizations and gains

kernel-orig: 13.23s

2.45x speedup

Actions: loop perm, 4x4 unroll
Result: big loop body with mem reuse

kernel-unroll: 5.41s

More sample codes

More codes to study with MAQAO in

```
$WORK/MAQAO_HANDSON/loop_optim_tutorial.tgz
```