



MAQAO

Performance Analysis and Optimization Tool

Cédric VALENSI, Emmanuel OSERET

{cedric.valensi, emmanuel.oseret}@uvsq.fr

Performance Evaluation Team, University of Versailles S-Q-Y

<http://www.maqao.org>

VI-HPS 30th TW Barcelona – Spain – 21-25 January 2019



Performance analysis and optimisation

How much can I optimise my application?

- Can it actually be done?
- What would the effort/gain ratio be?

Where can I gain time?

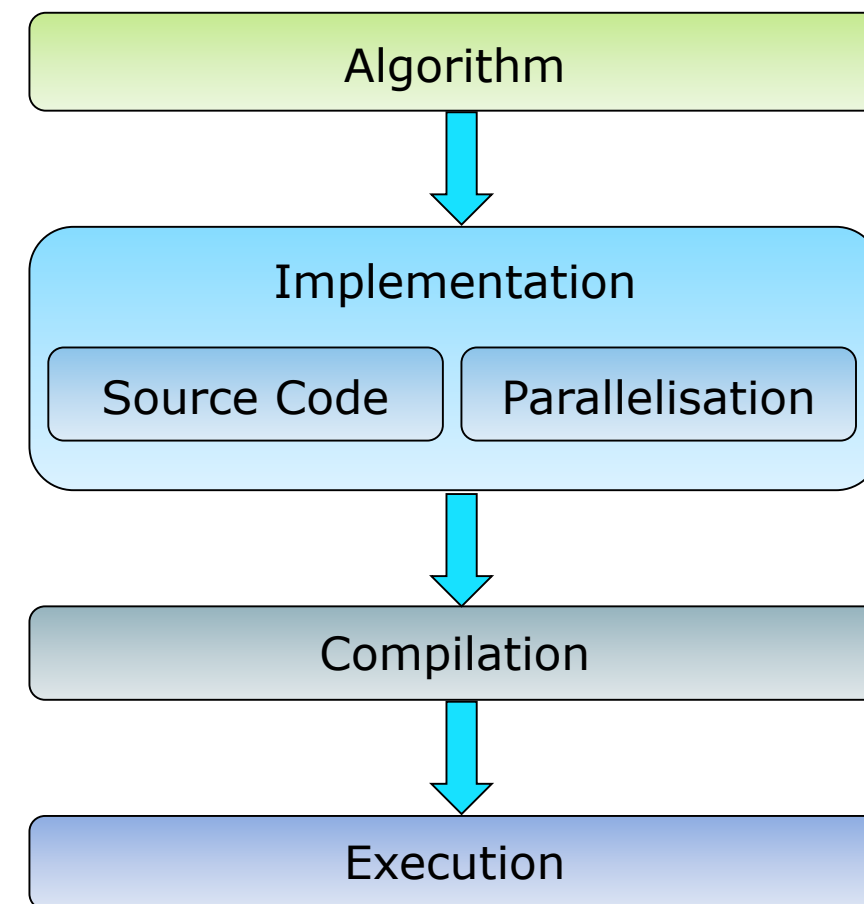
- Where is my application wasting time?

Why is the application spending time there?

- Algorithm, implementation or hardware?
- Data access or computation?

How can I improve the situation?

- In which step(s) of the design process?
- What additional information do I need?



A multifaceted problem

Pinpointing the performance bottlenecks

Identifying the dominant issues

- Algorithms, implementation, parallelisation, ...

Making the **best use** of the machine features

- Complex multicore and manycore CPUs
- Complex memory hierarchy

Finding the **most rewarding** issues to be fixed

- **40%** total time, expected **10%** speedup

- → TOTAL IMPACT: **4%** speedup



- **20%** total time, expected **50%** speedup

- → TOTAL IMPACT: **10%** speedup



→ **Need for dedicated and complementary tools**

Motivating example

Code of a loop representing ~10% walltime

```

do j = ni + nvalue1, nato
  nj1 = ndim3d*j + nc ; nj2 = nj1 + nvalue1 ; nj3 = nj2 + nvalue1
  u1 = x11 - x(nj1) ; u2 = x12 - x(nj2) ; u3 = x13 - x(nj3)
  rtest2 = u1*u1 + u2*u2 + u3*u3 ; cnij = eci*qEold(j)
  rij = demi*(rvwi + rvwalc1(j))
  drtest2 = cnij/(rtest2 + rij) ; drtest = sqrt(drtest2)
  Eq = qq1*qq(j)*drtest
  ntj = nti + ntype(j)
  Ed = ceps(ntj)*drtest2*drtest2*drtest2
  Eqc = Eqc + Eq ; Ephob = Ephob + Ed
  gE = (c6*Ed + Eq)*drtest2 ; virt = virt + gE*rtest2
  u1g = u1*gE ; u2g = u2*gE ; u3g = u3*gE
  g1c = g1c - u1g ; g2c = g2c - u2g ; g3c = g3c - u3g
  gr(nj1, thread_num) = gr(nj1, thread_num) + u1g
  gr(nj2, thread_num) = gr(nj2, thread_num) + u2g
  gr(nj3, thread_num) = gr(nj3, thread_num) + u3g
end do

```

Annotations for the code:

- 1) High number of statements (points to the entire loop body)
- 2) Non-unit stride accesses (points to `ni + nvalue1, nato` and the `gr` array accesses)
- 3) Indirect accesses (points to `ceps(ntj)`)
- 4) DIV/SQRT (points to `drtest = sqrt(drtest2)`)
- 5) Reductions (points to `Eqc = Eqc + Eq` and `Ephob = Ephob + Ed`)
- 6) Variable number of iterations (points to the `do j =` line)

Source code and associated issues:

- 1) High number of statements
- 2) Non-unit stride accesses
- 3) Indirect accesses
- 4) DIV/SQRT
- 5) Reductions
- 6) Variable number of iterations

MAQAO: Modular Assembly Quality Analyzer and Optimizer

Objectives:

- Characterizing performance of HPC applications
- Focusing on performance at the **core level**
- **Guiding users** through optimization process
- Estimating return of investment (**R.O.I.**)

Characteristics:




- **Modular tool** offering complementary views
- Support for **Intel x86-64** and **Xeon Phi**
 - ARM under development
- LGPL3 Open Source software
- Developed at UVSQ since 2004
- Binary release available as **static executable**



Success stories

MAQAO was used for optimizing industrial and academic HPC applications:

- QMC=CHEM (IRSAMC)
 - Quantum chemistry
 - Speedup: **> 3x**
 - Moved invocation of function with identical parameters out of loop body
- Yales2 (CORIA)
 - Computational fluid dynamics
 - Speedup: **up to 2.8x**
 - Removed double structure indirections
- Polaris (CEA)
 - Molecular dynamics
 - Speedup: **1.5x – 1.7x**
 - Enforced loop vectorisation through compiler directives
- AVBP (CERFACS)
 - Computational fluid dynamics
 - Speedup: **1.08x – 1.17x**
 - Replaced division with multiplication by reciprocal
 - Complete unrolling of loops with small number of iterations



Partnerships

MAQAO was funded by UVSQ, Intel and CEA (French department of energy) through Exascale Computing Research (ECR) and the French Ministry of Industry through various FUI/ITEA projects (H4H, COLOC, PerfCloud, ELCI, etc...)



Provides core technology to be integrated with other tools:

- TAU performance tools with MADRAS patcher through MIL (MAQAO Instrumentation Language)
- ATOS bullxprof with MADRAS through MIL
- Intel Advisor
- INRIA Bordeaux HWLOC

PeXL ISV also contributes to MAQAO:

- Commercial performance optimization expertise
- Training and software development
- www.pexl.eu



MAQAO team and collaborators

- Prof. William Jalby
- *Prof. Denis Barthou*
- Prof. David J. Kuck
- Andrés S. Charif-Rubial, Ph D
- *Jean-Thomas Acquaviva, Ph D*
- *Stéphane Zuckerman, Ph D*
- *Julien Jaeger, Ph D*
- *Souad Koliaï, Ph D*
- Cédric Valensi, Ph D
- Eric Petit, Ph D
- *Zakaria Bendifallah, Ph D*
- Emmanuel Oseret, Ph D
- Pablo de Oliveira, Ph D
- *Tipp Moseley, Ph D*
- David C. Wong, Ph D
- *Jean-Christophe Beyler, Ph D*
- Mathieu Tribalat
- Hugo Bolloré
- *Jean-Baptiste Le Reste*
- *Sylvain Henry, Ph D*
- Salah Ibn Amar
- Youenn Lebras
- Othman Bouizi, Ph D
- *José Noudohouenou, Ph D*
- Aleksandre Vardoshvili
- Romain Pillot

Analysis at binary level

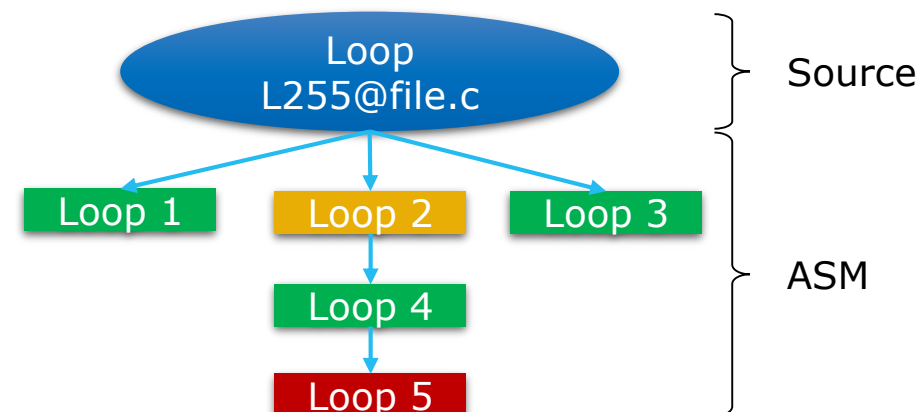
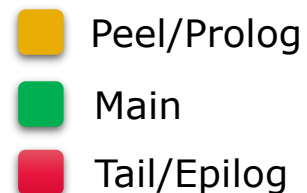
Advantages of binary analysis:

- Compiler optimizations increase the distance between the executed code and the source
- Source code instrumentation may prevent the compiler from applying some transformations

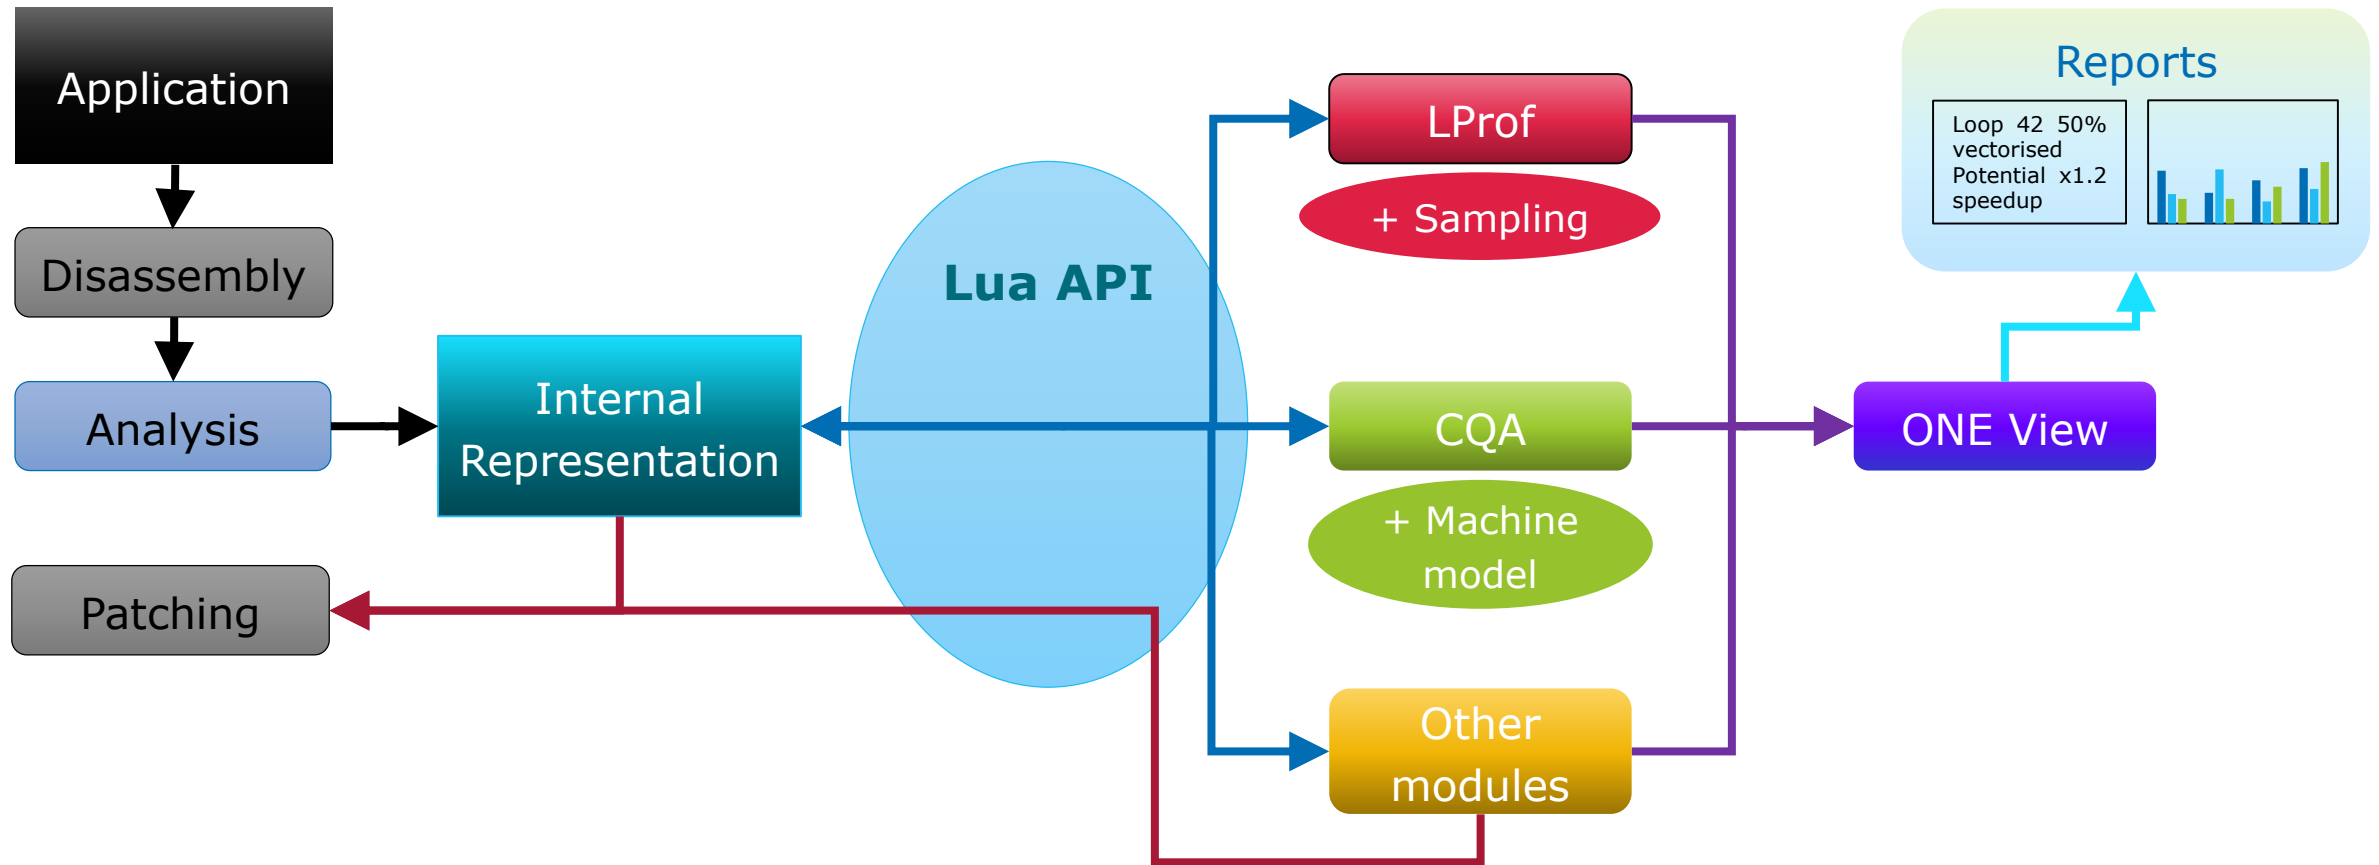
We want to evaluate the “real” executed code: **What You Analyse Is What You Run**

Main steps:

- Reconstruct the program structure
- Relate the analyses to source code
 - A single source loop can be compiled as multiple assembly loops
 - Affecting unique identifiers to loops

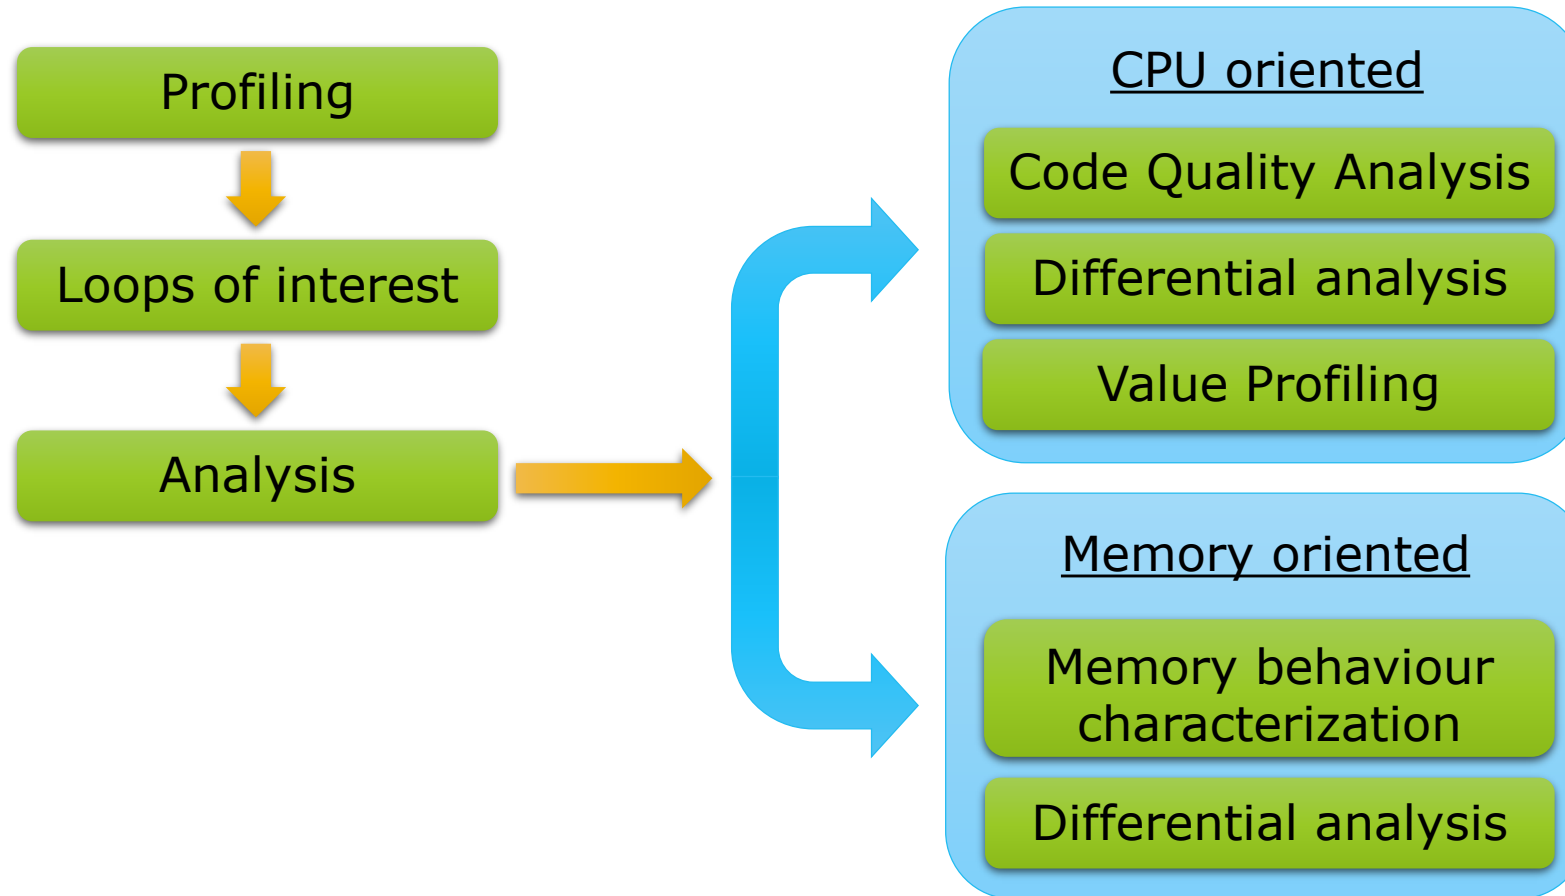


MAQAO Main structure



MAQAO Methodology

Decision tree



MAQAO LProf: Lightweight Profiler

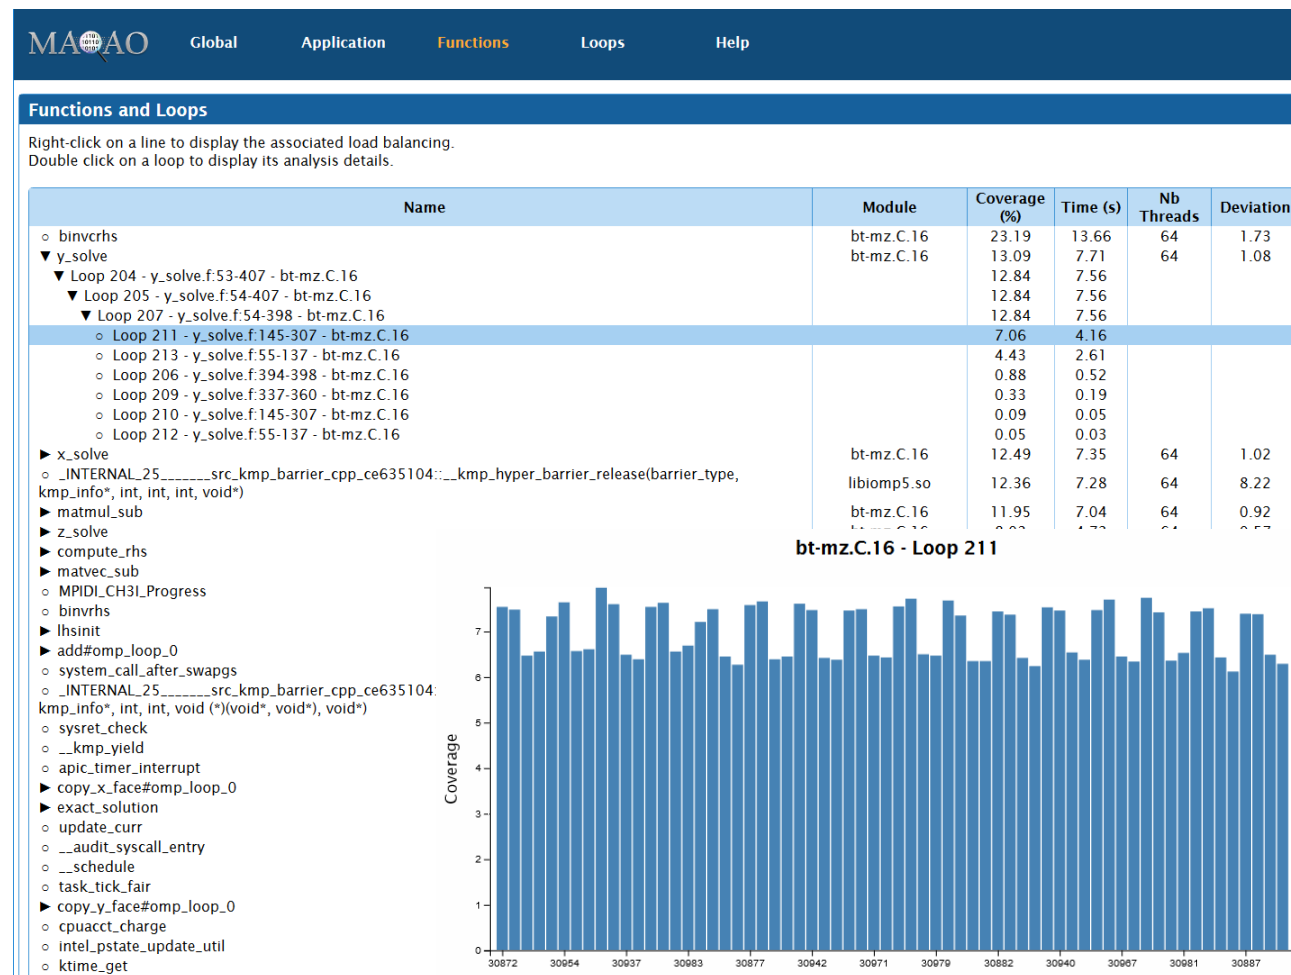
Goal: Lightweight localization of application hotspots

Features:

- **Sampling** based
- Access to hardware counters for additional information
- Results at function and loop granularity

Strengths:

- **Non intrusive:** No recompilation necessary
- **Low overhead**
- Agnostic with regard to parallel runtime



MAQAO CQA: Code Quality Analyzer

Goal: **Assist developers** in improving code performance

Features:

- Evaluates the **quality** of the compiler generated code
- Returns **hints and workarounds** to improve quality
- Focuses on **loops**
 - In HPC most of the time is spent in loops
- Targets **compute-bound** codes

Static analysis:

- Requires **no execution** of the application
- Allows **cross-analysis**

Static Reports

▼ CQA Report

The loop is defined in /tmp/NPB3.3.1-MZ/NPB3.3-MZ-MPI/BT-MZ/z_solve.f:415-423

▼ Path 1

2% of peak computational performance is used (0.77 out of 32.00 FLOP per cycle (GFLOPS @ 1GHz))

gain potential hint expert

Code clean check

Detected a slowdown caused by scalar integer instructions (typically used for address computation). By removing them, you can lower the cost of an iteration from 65.00 to 57.00 cycles (1.14x speedup).

Workaround

- Try to reorganize arrays of structures to structures of arrays
- Consider to permute loops (see vectorization gain report)
- To reference allocatable arrays, use "allocatable" instead of "pointer" pointers or qualify them with the "contiguous" attribute (Fortran 2008)
- For structures, limit to one indirection. For example, use a_b%c instead of a%b%c with a_b set to a%b before this loop

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 65.00 to 8.12 cycles (8.00x speedup).

Workaround

- Try another compiler or update/tune your current one:
 - use the vec-report option to understand why your loop was not vectorized. If "existence of vector dependences", try the IVDEP directive. If, using IVDEP, "vectorization possible but seems inefficient", try the VECTOR ALWAYS directive.
- 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: Fortran storage order is column-major: do i do j a(i,j) = b(i,j) (slow, non stride 1) => do i do j a(j,i) = b(i,j) (fast, stride 1)
 - If your loop streams arrays of structures (AoS), try to use structures of arrays instead (SoA): do i a(i)%x = b(i)%x (slow, non stride 1) => do i a%x(i) = b%x(i) (fast, stride 1)

Execution units bottlenecks

Found no such bottlenecks but see expert reports for more complex bottlenecks.

MAQAO CQA: Main Concepts

Most of the time, applications only exploit at best 5 to 10% of the peak performance.

Main elements of analysis:

- Peak performance
- Execution pipeline
- Resources/Functional units

Key performance levers for core level efficiency:

- Vectorizing
- Avoiding high latency instructions if possible
- Having the compiler generate an efficient code
- Reorganizing memory layout

Same instruction – Same cost



**Process up to
8X data**

MAQAO CQA: Compiler and programmer hints

Compiler can be driven using flags and pragmas:

- Ensuring full use of architecture capabilities (e.g. using flag -xHost on AVX capable machines)
- Forcing optimization (unrolling, vectorization, alignment, ...)
- Bypassing conservative behaviour when possible (e.g. 1/X precision)

Implementation changes:

- Improve data access
 - Loop interchange
 - Changing loop strides
 - Reshaping arrays of structures
- Avoid instructions with high latency

MAQAO ONE View: Performance View Aggregator

Automating the whole analysis process

- Invocation of the required MAQAO modules
- Generation of **aggregated performance views** available as HTML files

Main steps:

- Invokes LProf to **identify hotspots**
- Invokes CQA on **loop hotspots**

Available results:

- **Speedup** predictions
- Global code **quality** metrics
- **Hints** for improving performance
- Parallel **efficiency**



Analysing an application with MAQAO

Execute ONE View

- Provide all parameters necessary for executing the application
 - Parameters can be passed on the command line or into a configuration file
 - Parameters include binary name, MPI commands, dataset directory, ...

```
$ maqao oneview --create-report=one --binary=bt-mz.C.16 --mpi_command="mpirun -n 16"
```

- Analyses can be tweaked if necessary
 - Report one corresponds to profiling and code quality analysis
- ONE View can reuse an existing experiment directory to perform further analyses
- Results available in HTML by default
 - XLS files or console output available

Help and tutorials available on the MAQAO website: www.maqao.org/documentation.html

Analysing an application with MAQAO

MAQAO modules can be invoked separately for advanced analyses

- LProf
 - Profiling

```
$ maqao lprof xp=exp_dir --mpi-command="mpirun -n 16" -- ./bt-mz.C.16
```

- Display

```
$ maqao lprof xp=exp_dir -df
```

- Displaying the results from a ONE View run

```
$ maqao lprof xp=oneview_xp_dir/lprof_npsu -df
```

- CQA

```
$ maqao cqa loop=42 bt-mz.C.16
```

Help and tutorials available on the MAQAO website: www.maqao.org/documentation.html

Global summary

Experiment summary


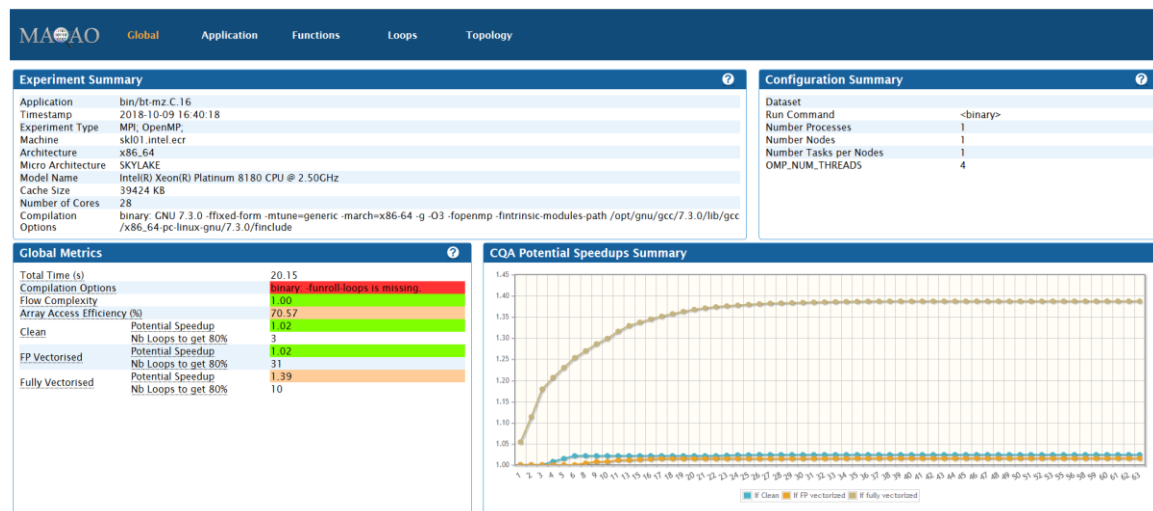
- Characteristics of the machine where the experiment took place

Global metrics

- General quality metrics derived from MAQAO analyses
- Global speedup predictions

CQA potential speedups

- Speedup prediction depending on the number of vectorised loops
- Ordered speedups to identify the loops to optimise in priority



Application Characteristics

Application categorisation

- Time spent in different regions of code

Function based profile

- Functions by coverage ranges

Loop based profile

- Loops by coverage ranges

Detailed loop based profile


- Loop types by coverage ranges



Application Characteristics: Time Categorisation

Identifying at a glance where time is spent

- Application
 - Main executable
- Parallelization
 - Threads
 - OpenMP
 - MPI
- System libraries
 - I/O operations
 - String operations
 - Memory management functions
- External libraries
 - Specialised libraries such as libm / libmkl
 - Application code in external libraries



Functions Profiling

Identifying hotspots

- Exclusive coverage
- Load balancing across threads
- Loops nests by functions

▼ matmul_sub

- Loop 230 - solve_subs.f:71-175 - bt-mz.C.16
- Loop 231 - solve_subs.f:71-175 - bt-mz.C.16

▼ z_solve

- ▼ Loop 232 - z_solve.f:53-423 - bt-mz.C.16
- ▼ Loop 233 - z_solve.f:54-423 - bt-mz.C.16
- ▼ Loop 236 - z_solve.f:54-423 - bt-mz.C.16
- Loop 239 - z_solve.f:146-308 - bt-mz.C.16
- Loop 235 - z_solve.f:55-137 - bt-mz.C.16
- Loop 234 - z_solve.f:415-423 - bt-mz.C.16

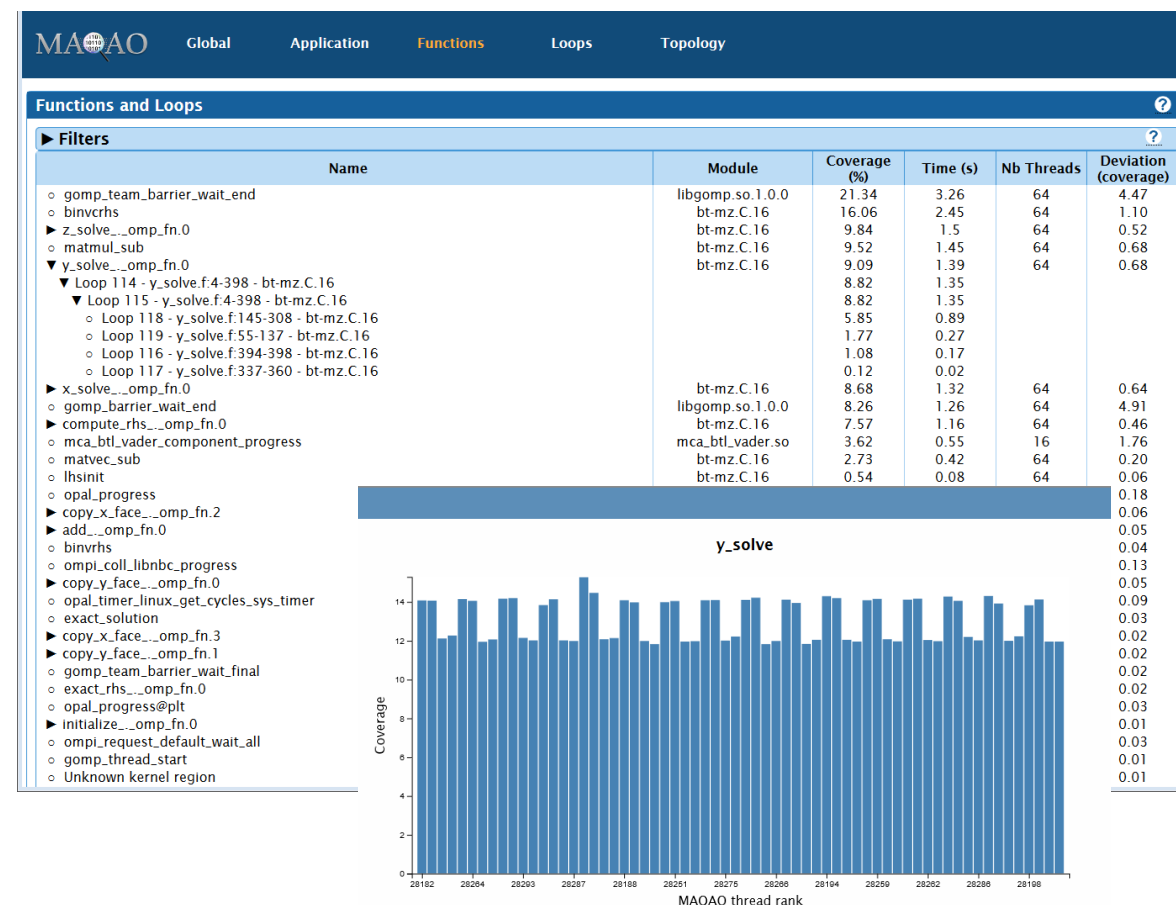
Single

Outermost

Inbetween

Inbetween

Innermost



Loops Profiling Summary

Identifying loop hotspots

- Vectorisation information
- Potential speedups

MAAO Global Application Functions Loops Help								
Loops Index								
Double click on a loop to display its analysis details. Check / uncheck metrics above the table to display / hide them.								
<input checked="" type="checkbox"/> Coverage (%)	<input checked="" type="checkbox"/> Vectorization Ratio (%)	<input checked="" type="checkbox"/> Speedup If Clean	<input checked="" type="checkbox"/> Speedup If FP Vectorized	<input checked="" type="checkbox"/> Speedup If Fully Vectorized				
Loop id	Source Lines	Source File	Source Function	Coverage (%)	Vectorization Ratio (%)	Speedup If Clean	Speedup If FP Vectorized	Speedup If Fully Vectorized
Loop 211	145-307	bt-mz.C.16:y_solve.f	y_solve	7.06	45.13	1	1.22	5.52
Loop 201	146-308	bt-mz.C.16:x_solve.f	x_solve	7.06	45.13	1	1.22	5.52
Loop 230	71-175	bt-mz.C.16:solve_subs.f	matmul_sub	5.57	100	1.02	1.9	4
Loop 213	55-137	bt-mz.C.16:y_solve.f	y_solve	4.43	47	1.05	1.16	5.93
Loop 203	57-139	bt-mz.C.16:x_solve.f	x_solve	3.93	48.36	1.01	1.09	5.83
Loop 239	146-308	bt-mz.C.16:z_solve.f	z_solve	3.06	8.97	1.07	1.8	8
Loop 235	55-137	bt-mz.C.16:z_solve.f	z_solve	2.81	22.08	1.03	1.62	7.49
Loop 234	415-423	bt-mz.C.16:z_solve.f	z_solve	1.54	0	1.14	1.67	8
Loop 122	304-349	bt-mz.C.16:rhs.f	compute_rhs	1.32	71.26	1.33	1.92	5.36
Loop 148	194-238	bt-mz.C.16:rhs.f	compute_rhs	1.25	71.59	1.32	1.93	5.38
Loop 162	83-132	bt-mz.C.16:rhs.f	compute_rhs	1.23	71.59	1.24	1.87	5.26
Loop 231	71-175	bt-mz.C.16:solve_subs.f	matmul_sub	1.11	10.59	1	2.29	8
Loop 227	23-27	bt-mz.C.16:solve_subs.f	matvec_sub	0.97	100	1	1.95	4
Loop 206	394-398	bt-mz.C.16:y_solve.f	y_solve	0.88	0	1.04	1.73	8
Loop 196	395-399	bt-mz.C.16:x_solve.f	x_solve	0.84	0	1.04	2.02	8
Loop 229	23-27	bt-mz.C.16:solve_subs.f	matvec_sub	0.62	100	1	1.95	4
Loop 170	40-50	bt-mz.C.16:rhs.f	compute_rhs	0.4	73.33	1	1.82	4
Loop 105	388-391	bt-mz.C.16:rhs.f	compute_rhs	0.35	100	1.12	1.83	4
Loop 238	313-314	bt-mz.C.16:z_solve.f	z_solve	0.35	0	1	1	8

High level reports

- Reference to the source code
- Bottleneck description
- Hints for improving performance
- Reports categorized by probability that applying hints will yield predicted gain
 - Gain: Good probability
 - Potential gain: Average probability
 - Hints: Lower probability

Loop Id: 224
Module: bt-mz.C.16
Source: solve_subs.f:71-175
Coverage: 4.79%

Source Code

```

71:  cblock(1,1) = cblock(1,1) - ablock(1,1)*bblock(1,1)
72:  >
73:  >
74:  >
75:  >
76:  cblock(2,1) = cblock(2,1) - ablock(2,1)*bblock(2,1)
77:  >
78:  >
79:  >
80:  >
81:  cblock(3,1) = cblock(3,1) - ablock(3,1)*bblock(3,1)
82:  >
83:  >
84:  >
85:  >
86:  cblock(4,1) = cblock(4,1) - ablock(4,1)*bblock(4,1)
87:  >
88:  >
89:  >
90:  >
91:  cblock(5,1) = cblock(5,1) - ablock(5,1)*bblock(5,1)
92:  >
93:  >
94:  >
95:  >
96:  cblock(1,2) = cblock(1,2) - ablock(1,2)*bblock(1,2)
97:  >
98:  >
99:  >
100: >
101: cblock(2,2) = cblock(2,2) - ablock(2,2)*bblock(2,2)
102: >
103: >
104: >

```

CQA

Path 1 / 1

Coverage 4.79 %

Function [matmul_sub](#)

Source file and lines solve_subs.f:71-175

Module bt-mz.C.16

The loop is defined in /ccc/dsku/nfs-server/user/cont001/ocre/valensic/NPB3.3.1-MZ/NPB3.3-MZ-MPI/BT-MZ/solve_subs.f:71-175.

It is main loop of related source loop which is unrolled by 2 (including vectorization).

[gain](#) [potential](#) [hint](#) [expert](#)

Code clean check

Detected a slowdown caused by scalar integer instructions (typically used for address computation). By removing them, you can lower the cost of an iteration from 27.00 to 25.00 cycles (1.08x speedup).

Workaround

[gain](#) [potential](#) [hint](#) [expert](#)

FMA

Presence of both ADD/SUB and MUL operations.

Workaround

[gain](#) [potential](#) [hint](#) [expert](#)

Type of elements and instruction set

195 SSE or AVX instructions are processing arithmetic or math operations on double precision FP elements in scalar mode (one at a time).

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

The binary loop is composed of 195 FP arithmetic operations:

- 70: addition or subtraction
- 125: multiply

The binary loop is loading 1760 bytes (220 double precision FP elements). The binary loop is storing 1632 bytes (204 double precision FP elements).

Arithmetic intensity

Arithmetic intensity is 0.06 FP operations per loaded or stored byte.

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.

Loop Analysis Reports – Expert View

Low level reports for performance experts

- Assembly-level
- Instructions cycles costs
- Instructions dispatch predictions
- Memory access analysis

Gain
Potential gain
Hints
Experts only

ASM code

In the binary file, the address of the loop is: 421409

Instruction	Nb	FU	P0	P1	P2	P3	P4	P5	P6	Latency	Recip. throughput
MOVAPS %XMM13,%XMM5	1	0.50	0.50	0	0	0	0	0	0	2	0.50
INC %RDI	1	0	0	0	0	1.50	0.50	0	1	1	
DIVSD 0x28(%R10,%RDX,1),%XMM5	4	1	0	0.50	0.50	0	0	0	0	40-42	12-32
MOVAPS %XMM5,%XMM15	1	0.50	0.50	0	0	0	0	0	2	0.50	
MULSD %XMM5,%XMM15	1	0.50	0.50	0	0	0	0	0	6	0.50	
MOVSD %XMM5,0x12890(%R14)	1	0	0	0.50	0.50	0	0	1	2	1	
MULSD %XMM15,%XMM5	1	0.50	0.50	0	0	0	0	0	6	0.50	
MOVSD %XMM15,0x12898(%R14)	1	0	0	0.50	0.50	0	0	1	2	1	
MOVSD %XMM5,0x128a0(%R14)	1	0	0	0.50	0.50	0	0	1	2	1	
MOVSD %XMM14,(%R9,%R14,1)	1	0	0	0.50	0.50	0	0	1	2	1	
MOVSD %XMM14,0x28(%R9,%R14,1)	1	0	0	0.50	0.50	0	0	1	2	1	

Loop Id: 224
Module: bt-mz.C.16
Source: solve_subs.f:71-175
Coverage: 4.79%

Assembly Code

Hide groups analysis

Ox424e35 MOVUPS (%RDI,%RAX,8)%XMM4 [3]
Ox424e39 MOVAPS %XMM5,%XMM2
Ox424e3c MULPD %XMM4,%XMM2
Ox424e40 LEA (%RCX,%RAX,8)%RSI
Ox424e44 MOVUPS (%RSI,%XMM15) [2]
Ox424e48 SUBPD %XMM2,%XMM15
Ox424e4d LEA 0x28(%RDI,%RAX,8)%R8
Ox424e52 MOVUPS (%R8,%XMM2) [4]
Ox424e56 MOVUPS 0x1d0(%RSP,%XMM1) [1]
Ox424e5e MOVAPS %XMM12,%XMM14
Ox424e62 MULPD %XMM2,%XMM1
Ox424e66 MOVUPS 0x130(%RSP,%XMM0) [1]
Ox424e6e SUBPD %XMM1,%XMM15
Ox424e73 MOVUPS 0x28(%R8,%XMM1) [4]
Ox424e78 MULPD %XMM1,%XMM0
Ox424e7c SUBPD %XMM0,%XMM15
Ox424e81 MOVUPS 0x50(%R8,%XMM0) [4]
Ox424e86 MOVUPS 0xd0(%RSP,%XMM3) [1]
Ox424e8e MULPD %XMM0,%XMM3
Ox424e92 SUBPD %XMM3,%XMM15
Ox424e97 MOVUPS 0x78(%R8,%XMM3) [4]

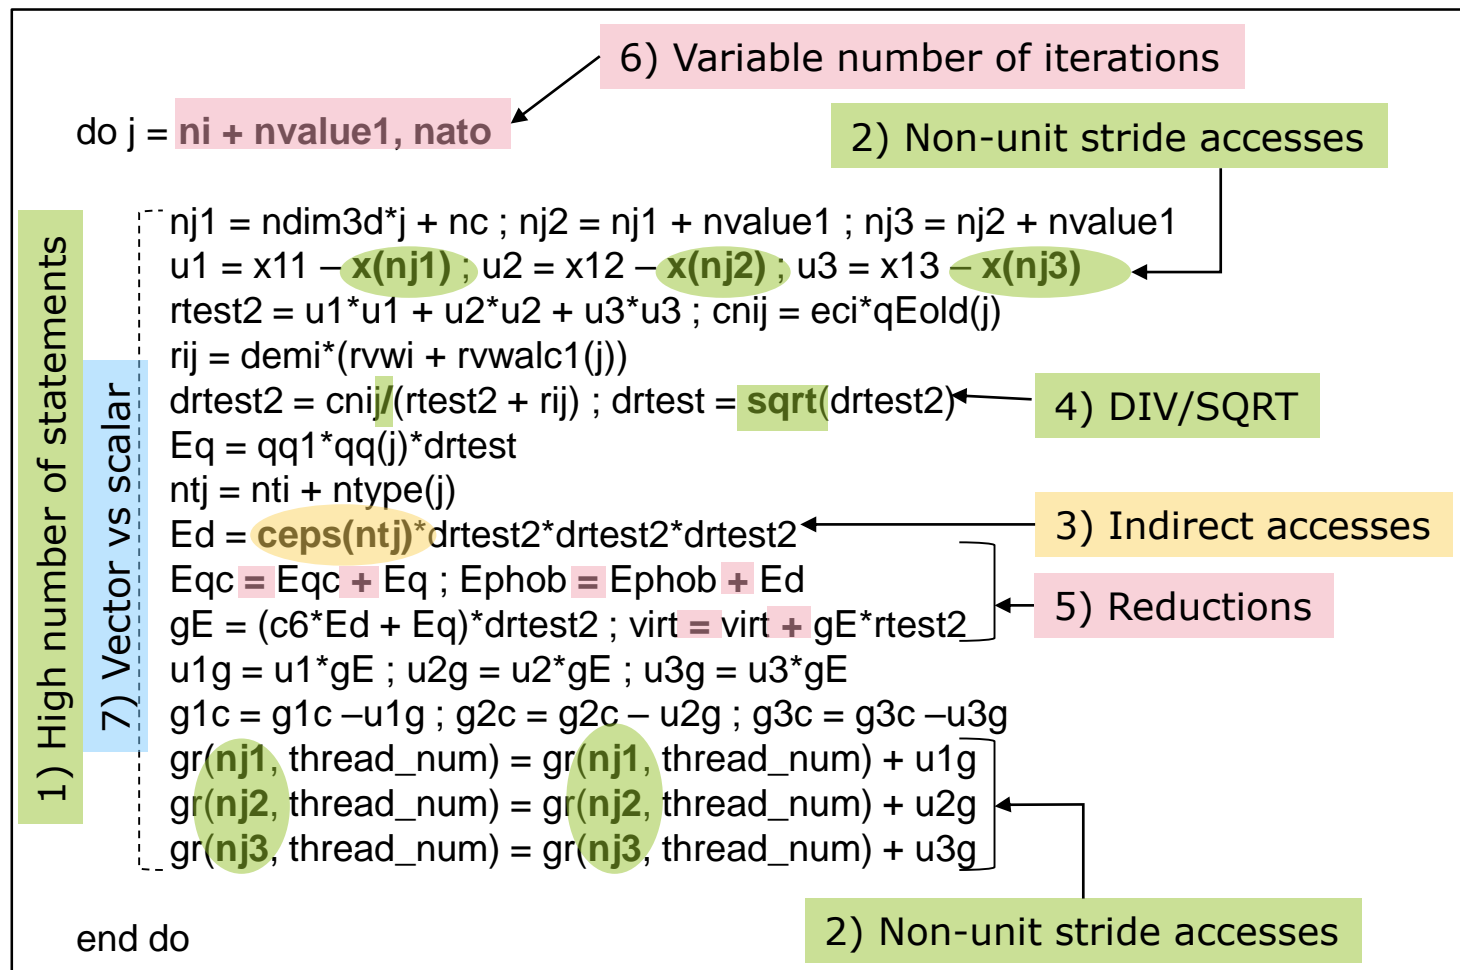
CQA Advanced

Path 1 / 1

Metric	Value
Coverage (% app. time)	4.79
Time (s)	0.23
CQA speedup if clean	1.08
CQA speedup if FP arith vectorized	1.65
CQA speedup if fully vectorized	2.00
CQA speedup if no inter-iteration dependency	NA
CQA speedup if next bottleneck killed	1.08
Source	solve_subs.f:71-175
Source loop unroll info	unrolled by 2
Source loop unroll confidence level	max
Unroll/vectorization loop type	main
Unroll factor	2
CQA cycles	27.00
CQA cycles if clean	25.00
CQA cycles if FP arith vectorized	16.32
CQA cycles if fully vectorized	13.50
Front-end cycles	22.50
P0 cycles	25.00
P1 cycles	27.00
P2 cycles	13.00
P3 cycles	13.00

Application to Motivating Example

Issues identified by CQA



CQA can detect and provide hints to resolve most of the identified issues:

- 1) High number of statements
- 2) Non-unit stride accesses
- 3) Indirect accesses
- 4) DIV/SQRT
- 5) Reductions
- 6) Variable number of iterations
- 7) Vector vs scalar

Application to Motivating Example

Vectorization

Your loop is partially vectorized.
Only 28% of vector register length is used (average across all SSE/AVX instructions).
By fully vectorizing your loop, you can lower the cost of an iteration from 57.00 to 21.50 cycles (2.65x speedup).
51% of SSE/AVX instructions are used in vector version (process two or more data elements in vector registers):

- 24% of SSE/AVX loads are used in vector version.
- 0% of SSE/AVX stores are used in vector version.

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

Proposed solution(s):

- Try another compiler or update/tune your current one:
 - use the vec-report option to understand why your loop was not vectorized. If "existence of vector dependences", try the IVDEP directive. If, using IVDEP, "vectorization possible but seems inefficient", try the VECTOR ALWAYS directive.
- 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:
Fortran storage order is column-major: do i do j a(i,j) = b(i,j) (slow, non stride 1) => do i do j a(j,i) = b(j,i) (fast, stride 1)
 - If your loop streams arrays of structures (AoS), try to use structures of arrays instead (SoA):
do i a(i)%x = b(i)%x (slow, non stride 1) => do i a%x(i) = b%x(i) (fast, stride 1)

Execution units bottlenecks

Performance is limited by:

- execution of divide and square root operations (the divide/square root unit is a bottleneck)
- execution of INT/FP operations in vector registers (the VPU is a bottleneck)

By removing all these bottlenecks, you can lower the cost of an iteration from 57.00 to 48.00 cycles (1.19x speedup).

Proposed solution(s):

- 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.
Check whether you really need double precision. If not, switch to single precision to speedup execution.
- Reduce arithmetical operations on array elements

FMA

Detected 48 FMA (fused multiply-add) operations.
Presence of both ADD/SUB and MUL operations.

Proposed solution(s):

Try to change order in which elements are evaluated (using parentheses) in arithmetic expressions containing both ADD/SUB and MUL operations to enable your compiler to generate FMA instructions wherever possible.
For instance $a + b * c$ is a valid FMA (MUL then ADD).
However $(a+b) * c$ cannot be translated into FMA.

Slow data structures access

Detected data structures (typically arrays) that cannot be efficiently read/written:

- Constant non-unit stride: 1 occurrence(s)
- Irregular (variable stride) or indirect: 1 occurrence(s)

7) Motivating factors:

- 1) High number of statements
- 2) Non-unit stride accesses
- 3) Indirect accesses
- 4) DIV/SQRT
- 5) Reductions
- 6) Variable number of iterations
- 7) Vector vs scalar

MAQAO ONE View Thread/Process View

Software Topology

- Processes by node
- Thread by process

View by thread

- Function profile at the thread or process level


MAQAO			Global	Application	Functions	Loops	Topology	Help
Software Topology								
							ID	Time(s)
▼ Node skylake01								11.22
▼ Process 359337								11.22
○ Thread 359337								11.22
▶ Process 359338								11.16
▶ Process 359352								11.22
▶ Process 359351								11.2
▶ Process 359353								11.22
▶ Process 359354								11.22
▶ Process 359355								11.18
▶ Process 359356								11.18
▶ Process 359357								11.22
▶ Process 359358								11.18
▶ Process 359359								11.18
▶ Process 359360								11.16
▶ Process 359361								11.18
▶ Process 359362								11.08
▶ Process 359364								11.22
▶ Process 359366								11.19
○ AVERAGE								11.22

MAQAO					Global	Application	Functions	Loops	Topology	Help
Profiling node skylake01 - process 359337 - thread 359337										
			Name	Module	Coverage (%)	Time (s)				
○ MPIDI_CH3I_Progress				libmpi.so.12.0	20.62	2.31				
▶ calc_data_gradient				3D_cylinder	4.95	0.56				
▶ ics_advance_velocity_tfv4a_4th				3D_cylinder	3.75	0.42				
▶ calc_data_tridiag_op_product				3D_cylinder	3.58	0.4				
○ MPIR_Allreduce_group				libmpi.so.12.0	3.22	0.36				
▶ filter_real_data				3D_cylinder	2.43	0.27				
▶ update_int_comm				3D_cylinder	2.42	0.27				
○ system_call_after_swaps				SYSTEM CALL	1.66	0.19				
▶ adv_scalar_w_u_tfv4a_4th				3D_cylinder	1.59	0.18				
▶ solve_linear_system_deflated_pcg				3D_cylinder	1.45	0.16				

MAQAO ONE View Scalability Reports

Goal: Provide a view of the application scalability

- Profiles with different numbers of threads/processes
- Displays efficiency metrics for application



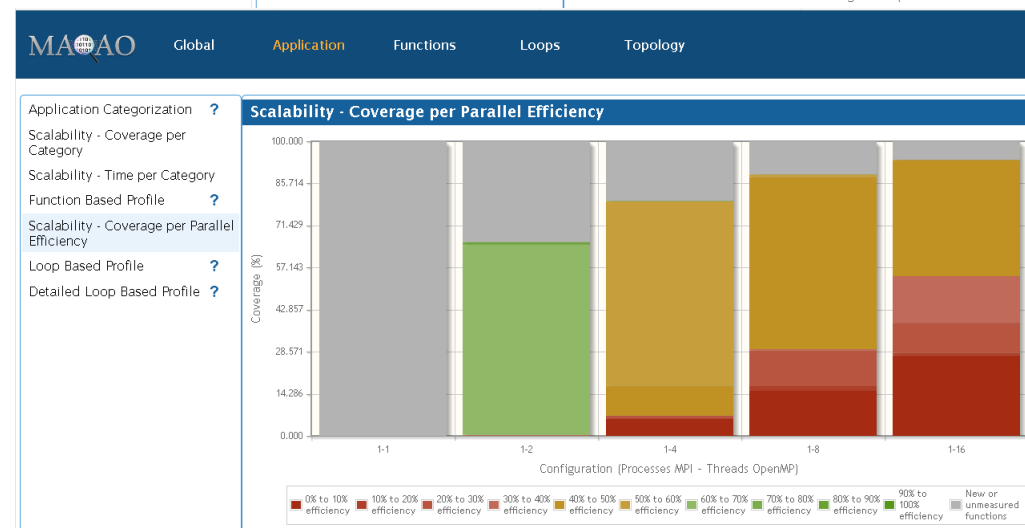
MAQAO ONE View Scalability Reports – Application View

Coverage per category

- Comparison of categories for each run

Coverage per parallel efficiency

- Efficiency = $T_{\text{sequential}} / (T_{\text{parallel}} * N_{\text{threads}})$
 - Distinguishing functions only represented in parallel or sequential
- Displays efficiency by coverage



Displays metrics for each function/loop

- Efficiency
- Potential speedup if efficiency=1

[illegible]

Thank you for your attention !

Questions ?