



# 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 33<sup>rd</sup> TW Jülich – Germany – 24-28 June 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?

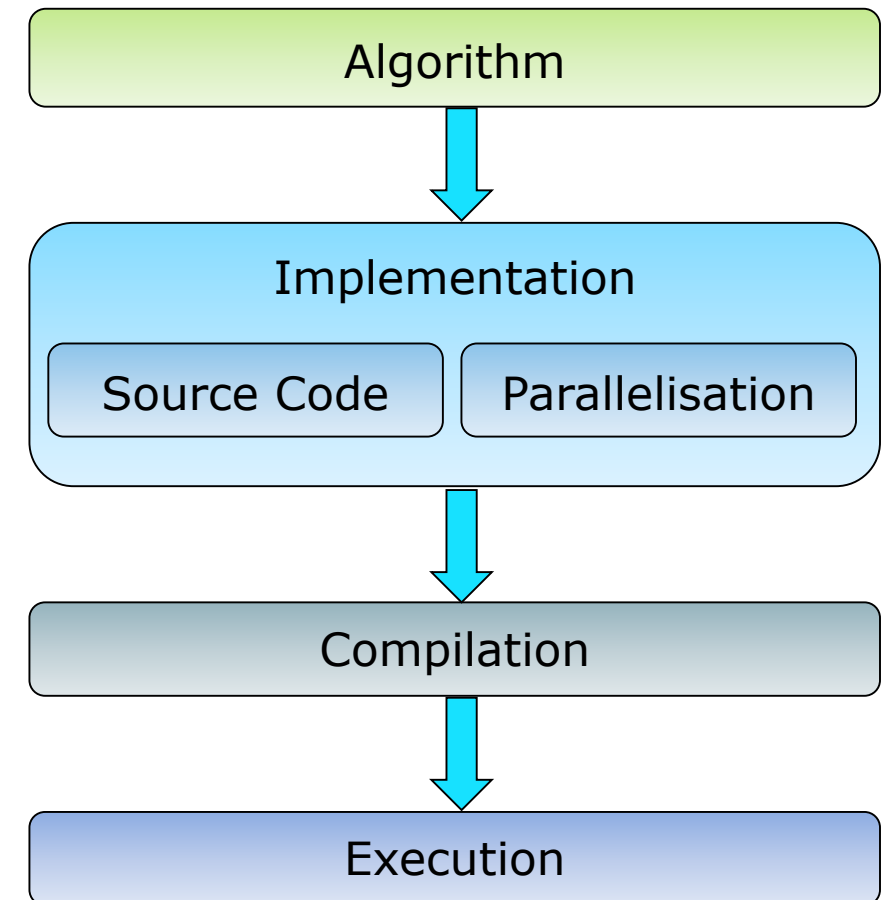
- Where is my application wasting time?

**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 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` loop header)

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:

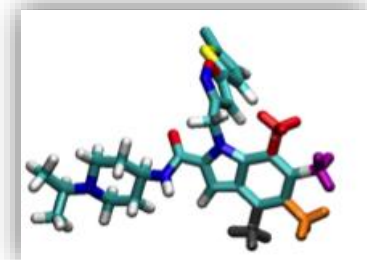
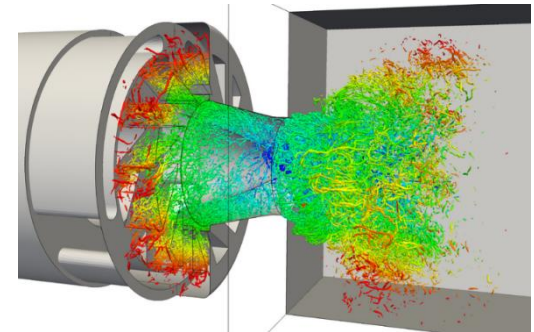
- **Modular tool** offering complementary views
- Support for **Intel x86-64** and **Xeon Phi**
  - ARM under development
- LGPL3 Open Source software
- Developed at UVSQ since 2004
- Binary release available as **static executable**



## 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, ELICI, 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](http://www.pexl.eu)



## MAQAO team and collaborators

---

### MAQAO Team

- Prof. William Jalby
- Cédric Valensi, Ph D
- Emmanuel Oseret, Ph D
- Mathieu Tribalat
- Salah Ibn Amar
- Youenn Lebras
- Kévin Camus

### Collaborators

- Prof. David J. Kuck
- Andrés S. Charif-Rubial, Ph D
- Eric Petit, Ph D
- Pablo de Oliveira, Ph D
- David C. Wong, Ph D
- Othman Bouizi, Ph D

### Past Collaborators or Team Members

- Prof. Denis Barthou
- Jean-Thomas Acquaviva, Ph D
- Stéphane Zuckerman, Ph D
- Julien Jaeger, Ph D
- Souad Koliaï, Ph D
- Zakaria Bendifallah, Ph D
- Tipp Moseley, Ph D
- Jean-Christophe Beyler, Ph D
- Hugo Bolloré
- Jean-Baptiste Le Reste
- Sylvain Henry, Ph D
- José Noudouhouenou, Ph D
- Aleksandre Vardoshvili
- Romain Pillot



## Analysis at binary level

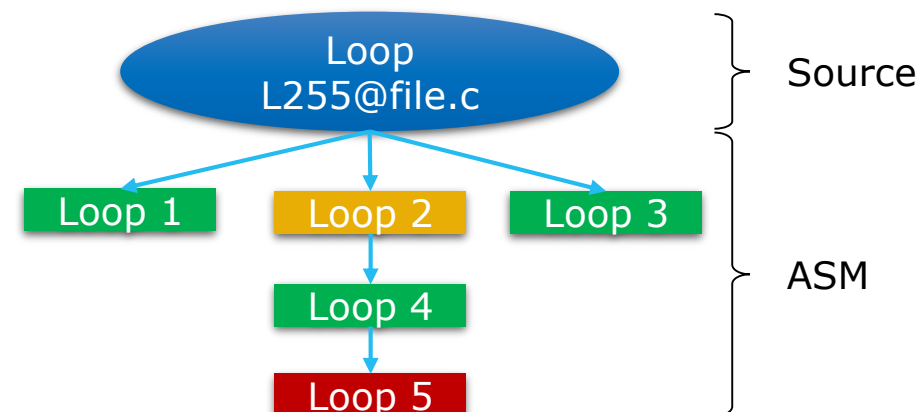
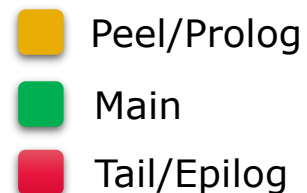
Advantages of binary analysis:

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

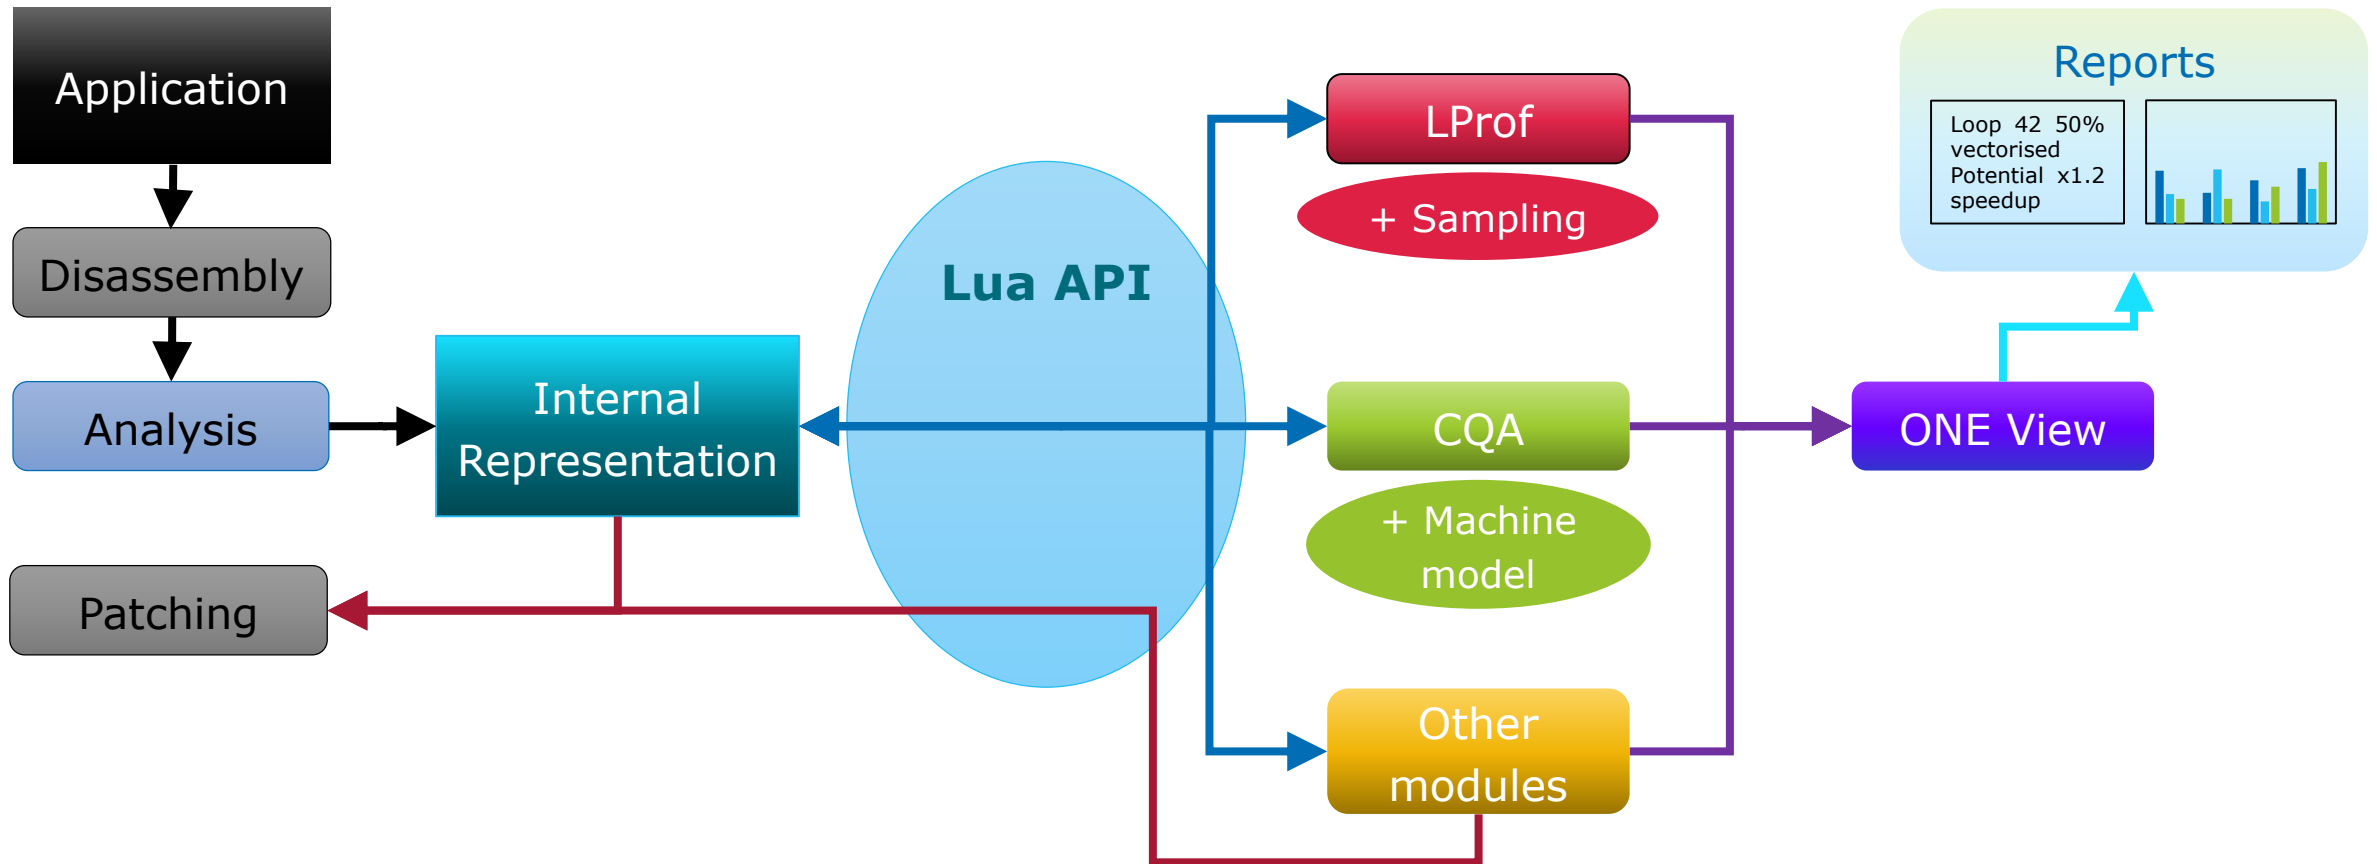
We want to evaluate the “real” executed code: **What You Analyse Is What You Run**

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

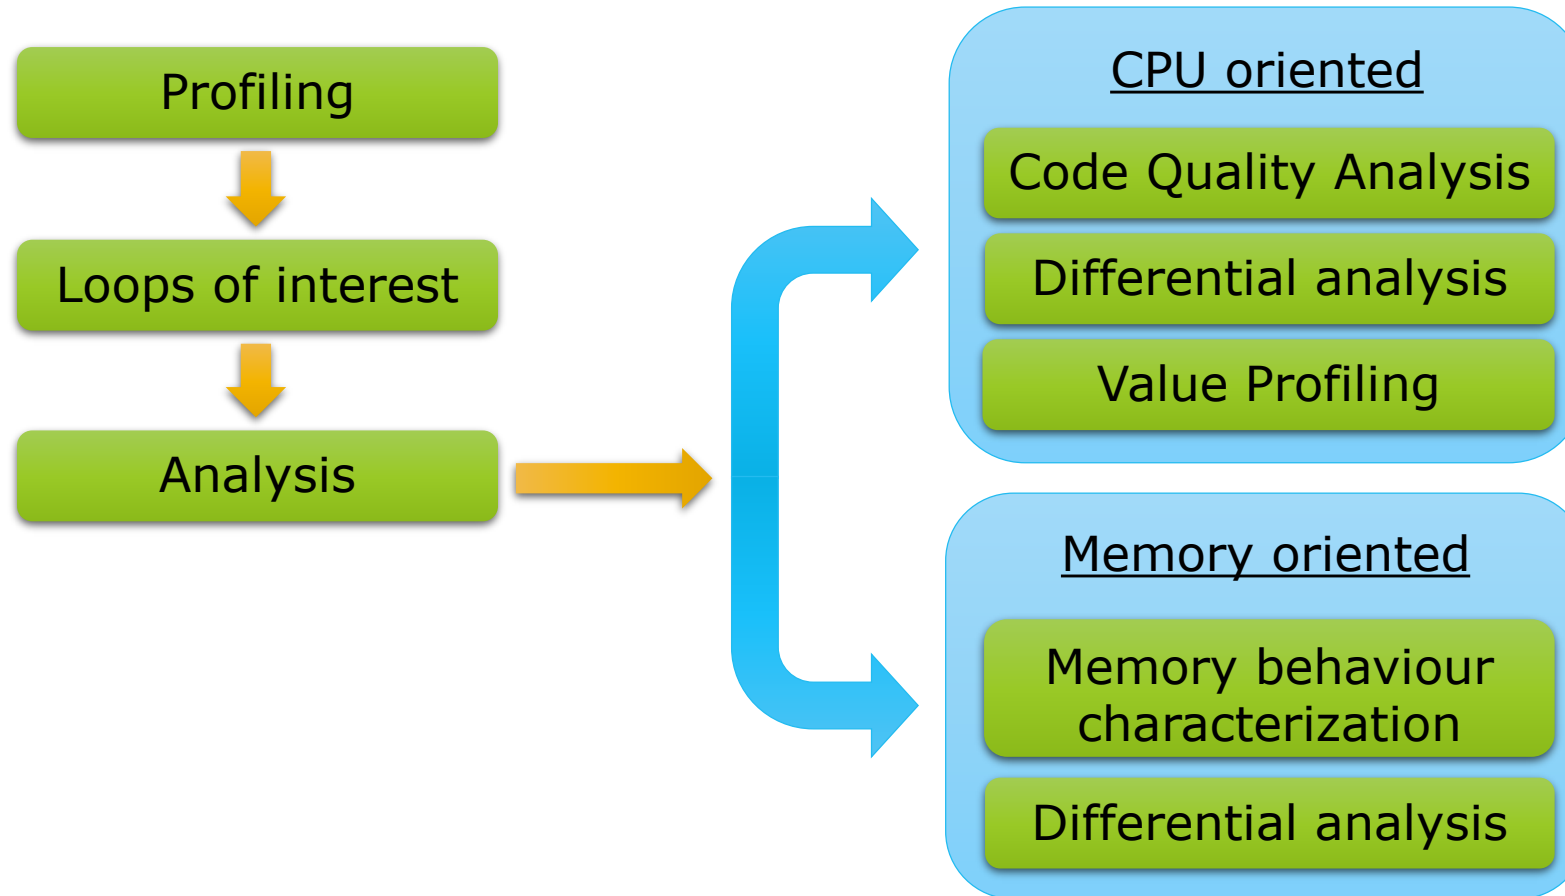


## MAQAO Main structure



## MAQAO Methodology

### Decision tree



# MAQAO LProf: Lightweight Profiler

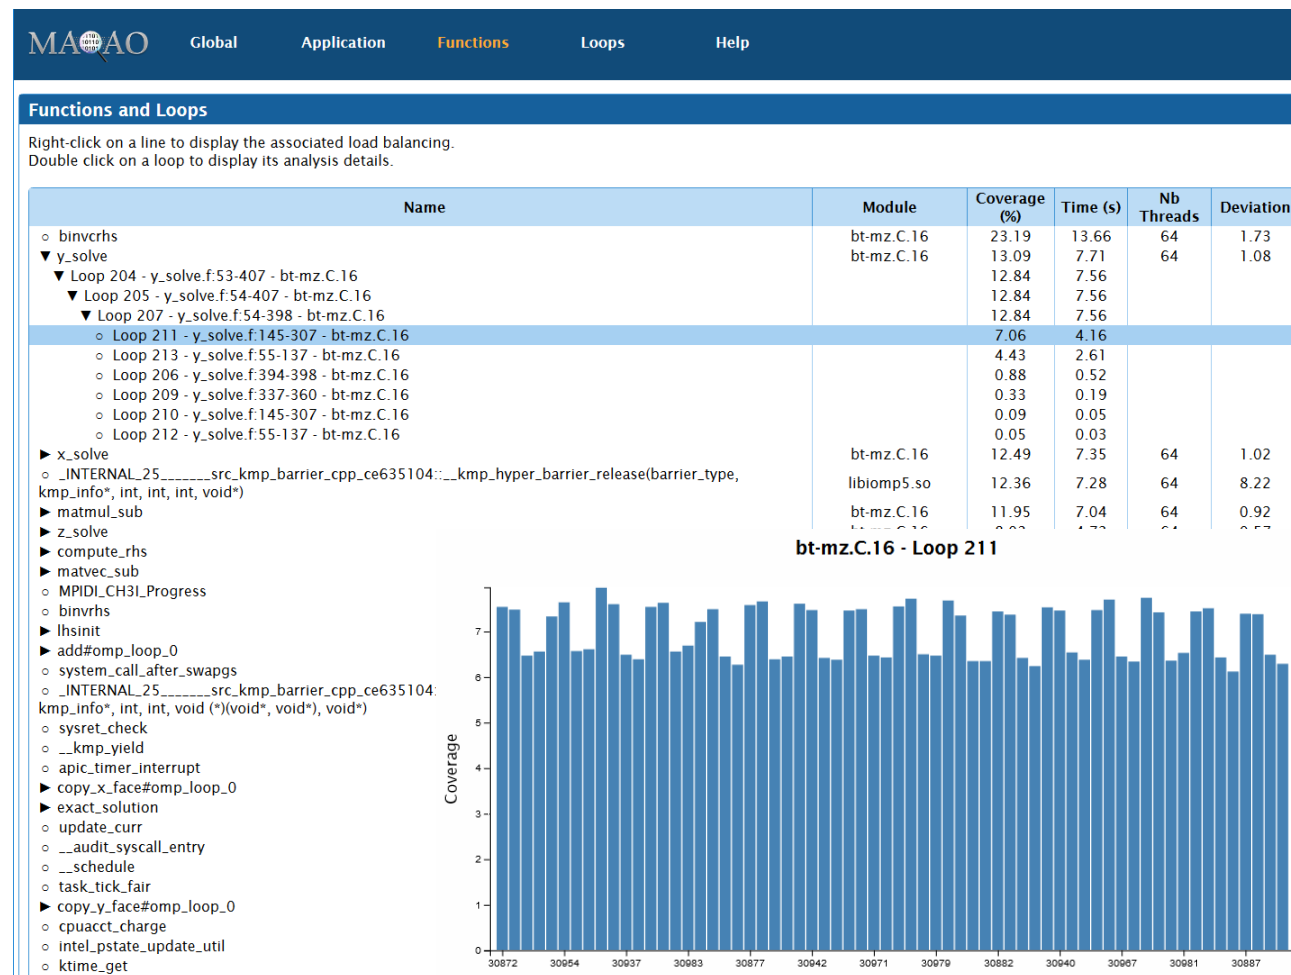
**Goal:** Lightweight localization of application hotspots

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"
```

```
$ maqao oneview --create-report=one --config=my_config.lua"
```

- 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

Online help available:

```
$ maqao oneview --help
```

## 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
```

Online help available:

```
$ maqao lprof --help
```

```
$ maqao cqa --help
```



# 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
  - 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

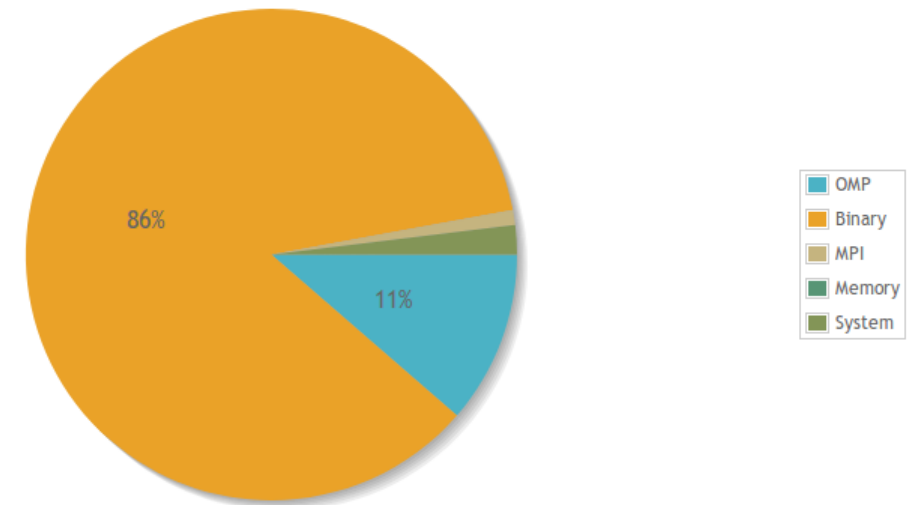


## Application Characteristics: Time Categorisation

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

Application Categorization



## 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

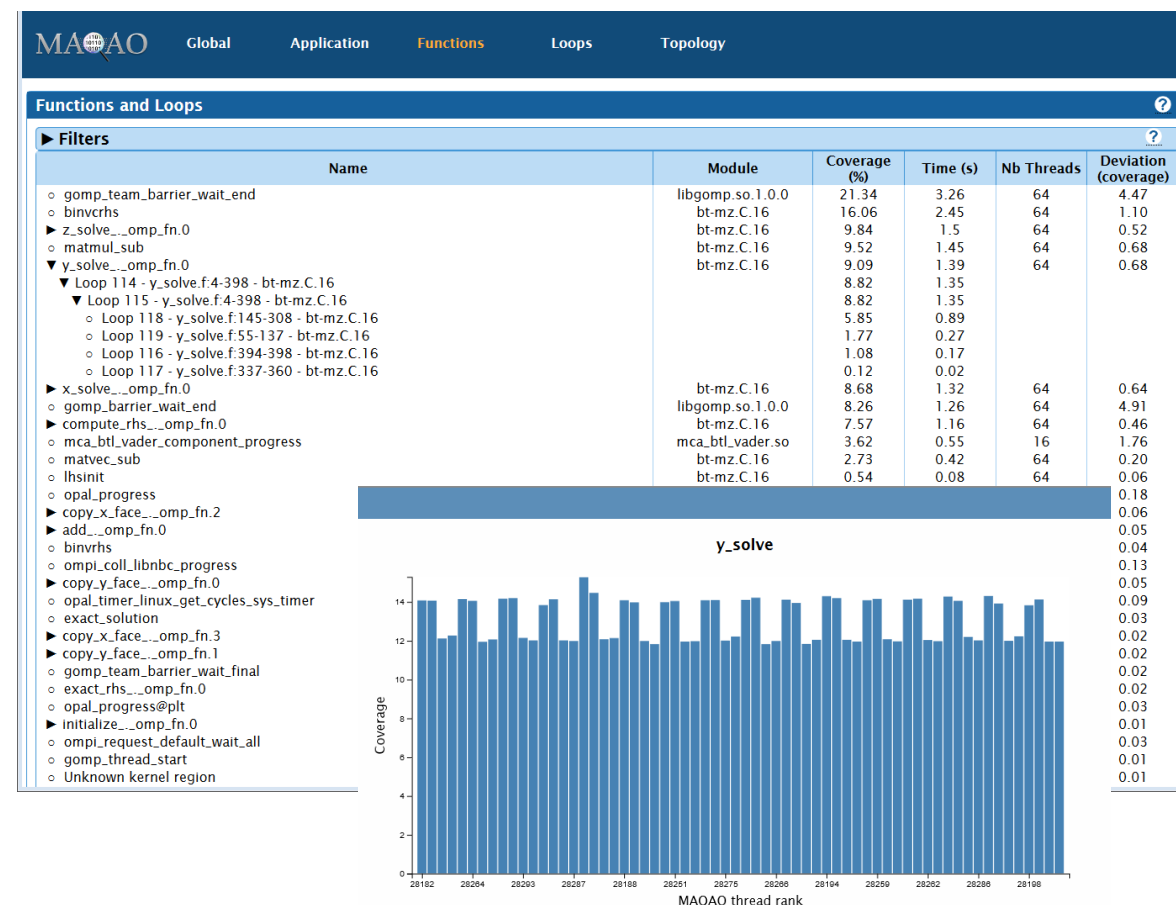
Single

Outermost

Inbetween

Inbetween

Innermost



## Loops Profiling Summary

### Identifying loop hotspots

- Vectorisation information
- Potential speedup by optimisation
  - Clean: Removing address computations
  - FP Vectorised: Vectorising floating-point computations
  - Fully Vectorised: Vectorising floating-point computations and memory accesses

MAAO Global Application Functions <b>Loops</b> Topology										
Loops Index ?										
<input checked="" type="checkbox"/> Coverage (%)	<input checked="" type="checkbox"/> Level	<input checked="" type="checkbox"/> Time (s)	<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	<input type="checkbox"/> Select none			
Loop id	Source Lines	▲ Source File	Source Function	Coverage (%)	Level	Time (s)	Vectorization Ratio (%)	Speedup If Clean	Speedup If FP Vectorized	Speedup If Fully Vectorized
Loop 236	146-308	bt-mz_B.16:z_solve.f	z_solve	4.2	Innermost	0.23	8.97	1.06	1.39	4
Loop 235	313-314	bt-mz_B.16:z_solve.f	z_solve	0.21	Innermost	0.01	0	1	1	4
Loop 234	351-373	bt-mz_B.16:z_solve.f	z_solve	0.15	Innermost	0.01	0	1	1	8
Loop 231	415-423	bt-mz_B.16:z_solve.f	z_solve	1.32	Innermost	0.07	0	1.14	1.07	4
Loop 233	54-423	bt-mz_B.16:z_solve.f	z_solve	0.14	InBetween	0.01	NA	NA	NA	NA
Loop 232	55-137	bt-mz_B.16:z_solve.f	z_solve	2.43	Innermost	0.13	22.08	1	1	4
Loop 208	145-307	bt-mz_B.16:y_solve.f	y_solve	5.14	Innermost	0.28	45.13	1	1	3.09
Loop 207	145-307	bt-mz_B.16:y_solve.f	y_solve	0.12	Innermost	0.01	8.97	1.06	1.38	4
Loop 206	337-360	bt-mz_B.16:y_solve.f	y_solve	0.38	Innermost	0.02	0	1	1	7.52
Loop 203	394-398	bt-mz_B.16:y_solve.f	y_solve	0.87	Innermost	0.05	0	1.01	1.17	4
Loop 201	53-407	bt-mz_B.16:y_solve.f	y_solve	0.01	Outermost	0	NA	NA	NA	NA
Loop 204	54-398	bt-mz_B.16:y_solve.f	y_solve	0.27	InBetween	0.01	NA	NA	NA	NA
Loop 210	55-137	bt-mz_B.16:y_solve.f	y_solve	2.67	Innermost	0.14	47	1	1	4
Loop 209	55-137	bt-mz_B.16:y_solve.f	y_solve	0.06	Innermost	0	22.08	1	1	4



# Loop Analysis Reports

## 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

The screenshot displays the MAOAO (Matrix Analysis and Optimization) interface for analyzing a loop. The top navigation bar includes tabs for Global, Application, Functions, Loops, and Topology. The main window shows the source code of a loop (Loop Id: 224, Module: bt-mz.C.16) with line numbers 711 to 1041. The analysis results on the right indicate a coverage of 4.79% and a function named `matmul_sub`. The source file and lines are `solve_subs.f:71-175`, and the module is `bt-mz.C.16`. The loop is defined in `/ccc/dsku/nfs-server/user/cont001/occe/valensic/NPB3.3.1-MZ-MPI/BT-MZ/solve_subs.f:71-175`. It is the main loop of related source loop which is unrolled by 2 (including vectorization).

The analysis results are categorized by probability of gain:

- gain**: Potential gain (Average probability)
- potential**: Potential gain (Average probability)
- hint**: Hint (Lower probability)
- expert**: Expert (Lower probability)

The **Code clean check** section indicates 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).

The **Workaround** section provides several hints for improvement:

- FMA**: Presence of both ADD/SUB and MUL operations. Workaround: Use parentheses in arithmetic to enable your compiler to use a valid FMA (MUL then ADD then MUL).
- AVX2**: AVX2 processors. By fully utilizing AVX2 registers, you can improve performance. Since your execution is data access bound, you can try to use the "contiguous" attribute to set to `contiguous` before this loop.
- Unit-stride**: Unit-stride elements are accessed regularly. Fortran storage order `1) => do i do j` structures of arrays `do i a%(i) = b%(i)`.
- 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`.

The **Type of elements and instruction set** section states: 195 SSE or AVX instructions are processing arithmetic or math operations on double precision FP elements in scalar mode (one at a time).

The **Matching between your loop (in the source code) and the binary loop** section states: The binary loop is composed of 195 FP arithmetical 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).

The **Arithmetic intensity** section states: Arithmetic intensity is 0.06 FP operations per loaded or stored byte.

The **Unroll opportunity** section states: 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

Assembly code

- Highlights groups of instructions accessing the same memory addresses

CQA internal metrics

**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	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	0	2	0.50
MULSD %XMM5,%XMM15	1	0.50	0.50	0	0	0	0	0	0	6	0.50
MOVSD %XMM5,0x12890(%R14)	1	0	0	0.50	0.50	0	0	0	1	2	1
MULSD %XMM15,%XMM5	1	0.50	0.50	0	0	0	0	0	0	6	0.50

Loop Id: 224 Module: bt-mz.C.16

Assembly Code

Hide groups analysis

Source: solve\_subs.f:71-175 Coverage: 4.79%

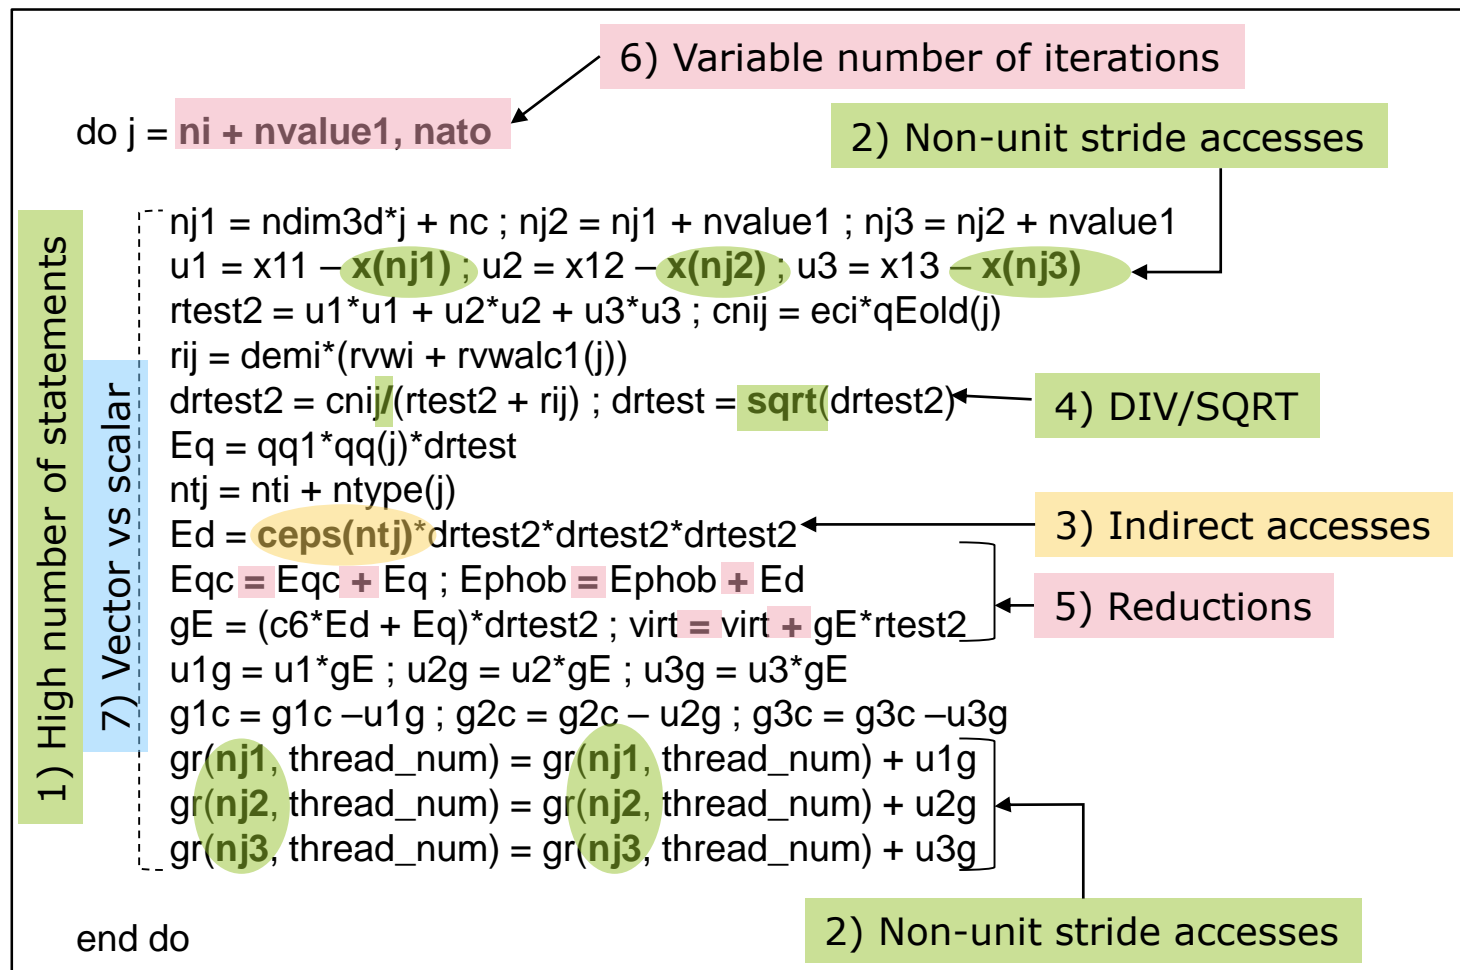
CQA Advanced

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

**Gain Potential gain Hints Experts only**

### 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

**Gain Potential gain Hints Experts only**

### 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.

**Gain Potential gain Hints Experts only**

### 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)

- 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

- List of nodes
- Processes by node
- Thread by process

### View by thread

- Function profile at the thread or process level

MAQAO
Global
Application
Functions
Loops
Topology

Software Topology
?

ID	Processes	Threads	Time(s)
▼ Node c251-109.wrangler.tacc.utexas.edu	8	32	5.34
▼ Process 145897		4	5.34
○ Thread 145897			5.34
○ Thread 145933			5.32
○ Thread 145952			5.32
○ Thread 145969			5.3
▶ Process 145899		4	5.34
▶ Process 145901		4	5.34
▶ Process 145903		4	5.34
▶ Process 145898		4	5.34
▶ Process 145900		4	5.34
▶ Process 145895		4	5.34
▶ Process 145896		4	5.34
▶ Node c251-110.wrangler.tacc.utexas.edu	8	32	5.36
○ AVERAGE			5.36

MAQAO
Global

Profiling node c251-109.wrangler.tacc.utexas.edu - process 145897 - thread 145897

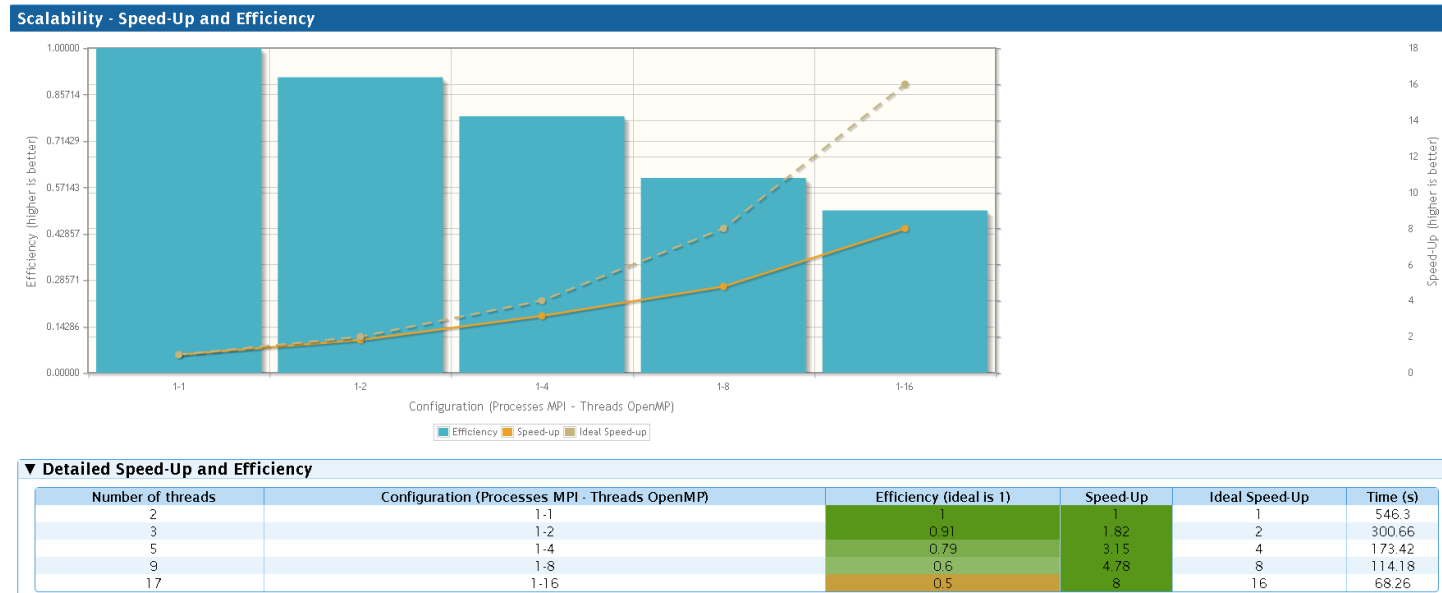
Name	Module	Coverage (%)	Time (s)
○ binvcrhs	bt-mz_B.16	24.34	1.3
○ _INTERNAL_25____src_kmp_barrier_cpp_fa608613::__kmp_hy per_barrier_gather(barrier_type, kmp_info*, int, int, void (*)(void*, v oid*), void*)	libiomp5.so	17.6	0.94
▶ matmul_sub	bt-mz_B.16	12.73	0.68
▶ y_solve	bt-mz_B.16	7.87	0.42
▶ compute_rhs	bt-mz_B.16	7.49	0.4
▶ x_solve	bt-mz_B.16	7.12	0.38
▶ z_solve	bt-mz_B.16	6.74	0.36



## 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



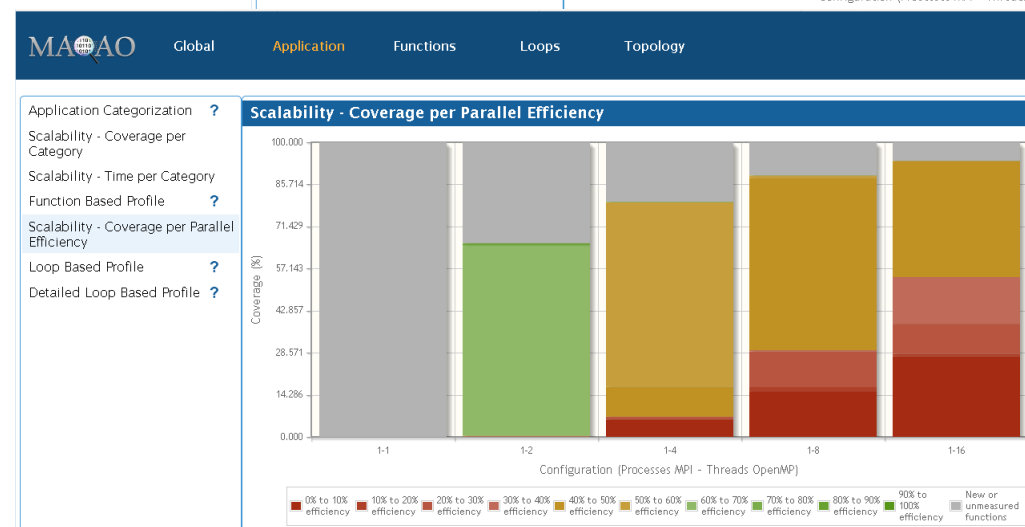
## MAQAO ONE View Scalability Reports – Application View

### Coverage per category

- Comparison of categories for each run

### Coverage per parallel efficiency

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





## More on MAQAO

---

MAQAO website: [www.maqao.org](http://www.maqao.org)

- Documentation: [www.maqao.org/documentation.html](http://www.maqao.org/documentation.html)
  - Tutorials for ONE View, LProf and CQA
  - Lua API documentation
- Latest release: <http://www.maqao.org/downloads.html>
  - Binary releases (2-3 per year)
  - Core sources
- Publications around MAQAO: <http://www.maqao.org/publications.html>

# Thank you for your attention !

## Questions ?