

# Score-P – A Joint Performance Measurement Run-Time Infrastructure for Periscope, Scalasca, TAU, and Vampir

---

VI-HPS Team



# Congratulations!?

---

- If you made it this far, you successfully used Score-P to
  - instrument the application
  - analyze its execution with a summary measurement, and
  - examine it with one of the interactive analysis report explorer GUIs
- ... revealing the call-path profile annotated with
  - the “Time” metric
  - Visit counts
  - MPI message statistics (bytes sent/received)
- ... but how **good** was the measurement?
  - The measured execution produced the desired valid result
  - however, the execution took rather longer than expected!
    - even when ignoring measurement start-up/completion, therefore
    - it was probably dilated by instrumentation/measurement overhead

# Performance analysis steps

---

- 0.0 Reference preparation for validation
  
- 1.0 Program instrumentation
  - 1.1 Summary measurement collection
  - 1.2 Summary analysis report examination
  
- 2.0 Summary experiment scoring
  - 2.1 Summary measurement collection with filtering
  - 2.2 Filtered summary analysis report examination
  
- 3.0 Event trace collection
  - 3.1 Event trace examination & analysis

## BT-MZ summary analysis result scoring

```
% scorep-score scorep_bt-mz_sum/profile.cubex
```

Estimated aggregate size of event trace:

Estimated requirements for largest trace buffer (max\_buf):

Estimated memory requirements (SCOREP\_TOTAL\_MEMORY):

(warning: The memory requirements cannot be satisfied by Score-P to avoid intermediate flushes when tracing. Set SCOREP\_TOTAL\_MEMORY=4G to get the maximum supported memory or reduce requirements using USR regions filters.)

flt	type	max_buf[B]	visits	time[s]	time[%]	time/visit[us]	region
	ALL	21,518,477,721	6,591,910,449	2608.77	100.0	0.40	ALL
	USR	21,431,996,118	6,574,793,529	1096.94	42.0	0.17	USR
	OMP	83,841,856	16,359,424	1454.29	55.7	88.90	OMP
	COM	2,351,570	723,560	3.28	0.1	4.54	COM
	MPI	288,136	33,928	54.26	2.1	1599.27	MPI
	SCOREP	41	8	0.00	0.0	216.83	SCOREP

160GB


21GB

21GB

- Report scoring as textual output

160 GB total memory  
21 GB per rank!

- Region/callpath classification
  - MPI** pure MPI functions
  - OMP** pure OpenMP regions
  - USR** user-level computation
  - COM** "combined" USR+OpenMP/MPI
  - ALL** aggregate of all region types




# BT-MZ summary analysis report breakdown

```
% scorep-score -r scorep_bt-mz_sum/profile.cubex
```

```
[...]
[...]
```

flt	type	max_buf[B]	visits	time[s]	time[%]	time/visit[us]	region
	ALL	21,518,477,721	6,591,910,449	2608.77	100.0	0.40	ALL
	USR	21,431,996,118	6,574,793,529	1096.94	42.0	0.17	USR
	OMP	83,841,856	16,359,424	1454.29	55.7	88.90	OMP
	COM	2,351,570	723,560	3.28	0.1	4.54	COM
	MPI	288,136	33,928	54.26	2.1	1599.27	MPI
	SCOREP	41	8	0.00	0.0	216.83	SCOREP

USR	6,883,222,086	2,110,313,472	370.34	14.2	0.18	matmul_sub_
USR	6,883,222,086	2,110,313,472	415.67	15.9	0.20	binvrhs_
USR	6,883,222,086	2,110,313,472	276.56	10.6	0.13	matvec_sub_
USR	293,617,584	87,475,200	11.97	0.5	0.14	binvrhs_
USR	293,617,584	87,475,200	13.76	0.5	0.16	lhsinit_
USR	224,028,792	68,892,672	8.63	0.3	0.13	exact_solution



More than  
20 GB just for these  
6 regions

## BT-MZ summary analysis score

---

- Summary measurement analysis score reveals
  - Total size of event trace would be ~160 GB
  - Maximum trace buffer size would be ~21 GB per rank
    - smaller buffer would require flushes to disk during measurement resulting in substantial perturbation
  - 99.5% of the trace requirements are for USR regions
    - purely computational routines never found on COM call-paths common to communication routines or OpenMP parallel regions
  - These USR regions contribute around 39% of total time
    - however, much of that is very likely to be measurement overhead for frequently-executed small routines
- Advisable to tune measurement configuration
  - Specify an adequate trace buffer size
  - Specify a filter file listing (USR) regions not to be measured

## BT-MZ summary analysis report filtering

```
% cat ../config/scorep.filt
SCOREP_REGION_NAMES_BEGIN
EXCLUDE
  binvcrhs*
  matmul_sub*
  matvec_sub*
  exact_solution*
  binvrhs*
  lhs*init*
  timer_*
SCOREP_REGION_NAMES_END

% scorep-score -f ../config/scorep.filt -c 2 \
  scorep_bt-mz_sum/profile.cubex
```

```
Estimated aggregate size of event trace:
Estimated requirements for largest trace buffer (max_buf):
Estimated memory requirements (SCOREP_TOTAL_MEMORY):
(hint: When tracing set SCOREP_TOTAL_MEMORY=215MB to avoid
intermediate flushes or reduce requirements using
USR regions filters.)
```

1624MB  
203MB  
215MB

- Report scoring with prospective filter listing 7 USR regions

1.6 GB of memory in total,  
203 MB per rank!  
(Including 2 metric values)

## BT-MZ summary analysis report filtering

```
% scorep-score -r -f ../config/scorep.filt \
  scorep_bt-mz_sum/profile.cubex
flt      type      max_buf[B]      visits time[s] time[%] time/
          region
          visit[us]
-        ALL      21,518,477,721  6,591,910,449  2608.77  100.0  0.40  ALL
-        USR      21,431,996,118  6,574,793,529  1096.94  42.0   0.17  USR
-        OMP      83,841,856     16,359,424    1454.29  55.7   88.90  OMP
-        COM      2,351,570     723,560      3.28    0.1    4.54  COM
-        MPI      288,136      33,928       54.26   2.1    1599.27  MPI
-        SCOREP   41           8            0.00    0.0    216.83  SCOREP

*        ALL      86,513,609    17,126,761   1511.84  58.0   88.27  ALL-FLT
+        FLT      21,431,964,112  6,574,783,688  1096.94  42.0   0.17  FLT
-        OMP      83,841,856    16,359,424   1454.29  55.7   88.90  OMP-FLT
*        COM      2,351,570     723,560      3.28    0.1    4.54  COM-FLT
-        MPI      288,136      33,928       54.26   2.1    1599.27  MPI-FLT
*        USR      32,006       9,841        0.00    0.0    0.35  USR-FLT
-        SCOREP   41           8            0.00    0.0    216.83  SCOREP-FLT

+        USR      6,883,222,086  2,110,313,472  370.34  14.2   0.18  matmul_sub_
+        USR      6,883,222,086  2,110,313,472  415.67  15.9   0.20  binvrhs_
+        USR      6,883,222,086  2,110,313,472  276.56  10.6   0.13  matvec_sub_
+        USR      293,617,584   87,475,200    11.97   0.5    0.14  binvrhs_
+        USR      293,617,584   87,475,200    13.76   0.5    0.16  lhsinit_
+        USR      224,028,792   68,892,672    8.63    0.3    0.13  exact_solution
```

- Score report breakdown by region (w/o additional metrics)

Filtered routines marked with '+'



## BT-MZ filtered summary measurement

```
% cd bin.scorep
% cp ../jobscript/dine/scorep.sbatch .
% vi scorep.sbatch

# Score-P measurement configuration
export SCOREP_EXPERIMENT_DIRECTORY=scorep_bt-mz_sum_filter
export SCOREP_FILTERING_FILE=../config/scorep.filt
#export SCOREP_TOTAL_MEMORY=100M
#export SCOREP_METRIC_PAPI=PAPI_TOT_INS,PAPI_TOT_CYC,...
#export SCOREP_ENABLE_TRACING=true

set -x
export OMP_NUM_THREADS=6
time -p mpiexec -np 8 ./bt-mz_C.8

% sbatch scorep.sbatch
```

- Set new experiment directory and re-run measurement with new filter configuration
- Submit job

# Score-P filtering

```
% cat ../config/scorep.filt
SCOREP_REGION_NAMES_BEGIN
EXCLUDE
  binvcrhs*
  matmul_sub*
  matvec_sub*
  exact_solution*
  binvrhs*
  lhs*init*
  timer_*
SCOREP_REGION_NAMES_END

% export SCOREP_FILTERING_FILE=\
../config/scorep.filt
```

Region name  
filter block  
using wildcards

Apply filter

- Filtering by source file name
  - All regions in files that are excluded by the filter are ignored
- Filtering by region name
  - All regions that are excluded by the filter are ignored
  - Overruled by source file filter for excluded files
- Apply filter by
  - exporting `SCOREP_FILTERING_FILE` environment variable
- Apply filter at
  - Run-time
  - Compile-time (GCC-plugin only, Intel in 7.0 release)
    - Add cmd-line option `--instrument-filter`
    - No overhead for filtered regions but recompilation

# Source file name filter block

---

- Keywords
  - Case-sensitive
  - SCOREP\_FILE\_NAMES\_BEGIN, SCOREP\_FILE\_NAMES\_END
    - Define the source file name filter block
    - Block contains EXCLUDE, INCLUDE rules
  - EXCLUDE, INCLUDE rules
    - Followed by one or multiple white-space separated source file names
    - Names can contain bash-like wildcards \*, ?, []
    - Unlike bash, \* may match a string that contains slashes
- EXCLUDE, INCLUDE rules are applied in sequential order
- Regions in source files that are excluded after all rules are evaluated, get filtered

```
# This is a comment
SCOREP_FILE_NAMES_BEGIN
  # by default, everything is included
  EXCLUDE */foo/bar*
  INCLUDE */filter_test.c
SCOREP_FILE_NAMES_END
```

# Region name filter block

---

- Keywords
  - Case-sensitive
  - SCOREP\_REGION\_NAMES\_BEGIN,  
SCOREP\_REGION\_NAMES\_END
    - Define the region name filter block
    - Block contains EXCLUDE, INCLUDE rules
  - EXCLUDE, INCLUDE rules
    - Followed by one or multiple white-space separated region names
    - Names can contain bash-like wildcards \*, ?, []
- EXCLUDE, INCLUDE rules are applied in sequential order
- Regions that are excluded after all rules are evaluated, get filtered

```
# This is a comment
SCOREP_REGION_NAMES_BEGIN
# by default, everything is included
EXCLUDE *
INCLUDE bar foo
        baz
        main
SCOREP_REGION_NAMES_END
```

## Region name filter block, mangling


- Name mangling
  - Filtering based on names seen by the measurement system
    - Dependent on compiler
    - Actual name may be mangled
- `scorep-score` names as starting point  
(e.g. `matvec_sub_`)
  - Use `*` for Fortran trailing underscore(s) for portability
  - Use `?` and `*` as needed for full signatures or overloading
  - Use `\` to escape special characters

```
void bar(int* a) {
    *a++;
}
int main() {
    int i = 42;
    bar(&i);
    return 0;
}
```

```
# filter bar:
# for gcc-plugin, scorep-score
# displays 'void bar(int*)',
# other compilers may differ

SCOREP_REGION_NAMES_BEGIN
    EXCLUDE void?bar(int?)
SCOREP_REGION_NAMES_END
```

# Mastering build systems



- Hooking up the Score-P instrumenter `scorep` into complex build environments like *Autotools* or *CMake* was always challenging
- Score-P provides convenience wrapper scripts to simplify this (since Score-P 2.0)
- *Autotools* and *CMake* need the used compiler already in the *configure step*, but instrumentation should not happen in this step, only in the *build step*

```
% SCOREP_WRAPPER=off \  
> cmake .. \  
> -DCMAKE_C_COMPILER=scorep-icc \  
> -DCMAKE_CXX_COMPILER=scorep-icpc
```


Disable instrumentation in the *configure step*

Specify the wrapper scripts as the compiler to use

- Allows to pass addition options to the Score-P instrumenter and the compiler via environment variables without modifying the *Makefiles*:  
`SCOREP_WRAPPER_INSTRUMENTER_FLAGS`, `SCOREP_WRAPPER_COMPILER_FLAGS`
- Run `scorep-wrapper --help` for a detailed description and the available wrapper scripts of the Score-P installation

# Score-P user instrumentation API

---



- Can be used to partition application into coarse grain phases
  - E.g., initialization, solver, & finalization
- Can be used to further subdivide functions
  - E.g., multiple loops inside a function
- Enabled with `--user` flag to Score-P instrumenter
- Available for Fortran / C / C++

# Score-P user instrumentation API (Fortran)



```
#include "scorep/SCOREP_User.inc"

subroutine foo(...)
  ! Declarations
  SCOREP_USER_REGION_DEFINE( solve )


  ! Some code...
  SCOREP_USER_REGION_BEGIN( solve, "<solver>", \
                           SCOREP_USER_REGION_TYPE_LOOP )

  do i=1,100
    [...]
  end do
  SCOREP_USER_REGION_END( solve )
  ! Some more code...
end subroutine
```

- Requires processing by the C preprocessor
  - For most compilers, this can be automatically achieved by having an uppercase file extension, e.g., `main.F` or `main.F90`



# Score-P user instrumentation API (C/C++)

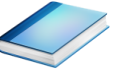


```
#include "scorep/SCOREP_User.h"

void foo()
{
    /* Declarations */
    SCOREP_USER_REGION_DEFINE( solve )

    /* Some code... */
    SCOREP_USER_REGION_BEGIN( solve, "<solver>",
                             SCOREP_USER_REGION_TYPE_LOOP )
    for (i = 0; i < 100; i++)
    {
        [...]
    }
    SCOREP_USER_REGION_END( solve )
    /* Some more code... */
}
```

# Score-P user instrumentation API (C++)



```
#include "scorep/SCOREP_User.h"

void foo()
{
    // Declarations

    // Some code...
    {
        SCOREP_USER_REGION( "<solver>",
                           SCOREP_USER_REGION_TYPE_LOOP )
        for (i = 0; i < 100; i++)
        {
            [...]
        }
    }
    // Some more code...
}
```

# Score-P measurement control API



- Can be used to temporarily disable measurement for certain intervals
  - Annotation macros ignored by default
  - Enabled with `--user` flag

```
#include "scorep/SCOREP_User.inc"

subroutine foo(...)
  ! Some code...
  SCOREP_RECORDING_OFF()
  ! Loop will not be measured
  do i=1,100
    [...]
  end do
  SCOREP_RECORDING_ON()
  ! Some more code...
end subroutine
```

Fortran (requires C preprocessor)

```
#include "scorep/SCOREP_User.h"

void foo(...) {
  /* Some code... */
  SCOREP_RECORDING_OFF()
  /* Loop will not be measured */
  for (i = 0; i < 100; i++) {
    [...]
  }
  SCOREP_RECORDING_ON()
  /* Some more code... */
}
```

C / C++

## Further information

---

- Community instrumentation & measurement infrastructure
  - Instrumentation (various methods)
  - Basic and advanced profile generation
  - Event trace recording
  - Online access to profiling data
- Available under 3-clause BSD open-source license
- Documentation & Sources:
  - <http://www.score-p.org>
- User guide also part of installation:
  - `<prefix>/share/doc/scorep/{pdf,html}/`
- Support and feedback: [support@score-p.org](mailto:support@score-p.org)
- Subscribe to [news@score-p.org](mailto:news@score-p.org), to be up to date