



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

Ronny Tschüter<sup>1)</sup>, Frank Winkler<sup>1)</sup>, Brian Wylie<sup>2)</sup>

With contributions from  
Andreas Knüpfer<sup>1)</sup> and Christian Rössel<sup>2)</sup>

<sup>1)</sup>ZIH TU Dresden , <sup>2)</sup>FZ Jülich

- Several performance tools co-exist
- Separate measurement systems and output formats
- Complementary features and overlapping functionality
- Redundant effort for development and maintenance
- Limited or expensive interoperability
- Complications for user experience, support, training



- Start a community effort for a common infrastructure
  - Score-P instrumentation and measurement system
  - Common data formats OTF2 and CUBE4
- Developer perspective:
  - Save manpower by sharing development resources
  - Invest in new analysis functionality and scalability
  - Save efforts for maintenance, testing, porting, support, training
- User perspective:
  - Single learning curve
  - Single installation, fewer version updates
  - Interoperability and data exchange
- SILC project funded by BMBF
- Close collaboration PRIMA project funded by DOE



GEFÖRDERT VOM

Bundesministerium  
für Bildung  
und Forschung



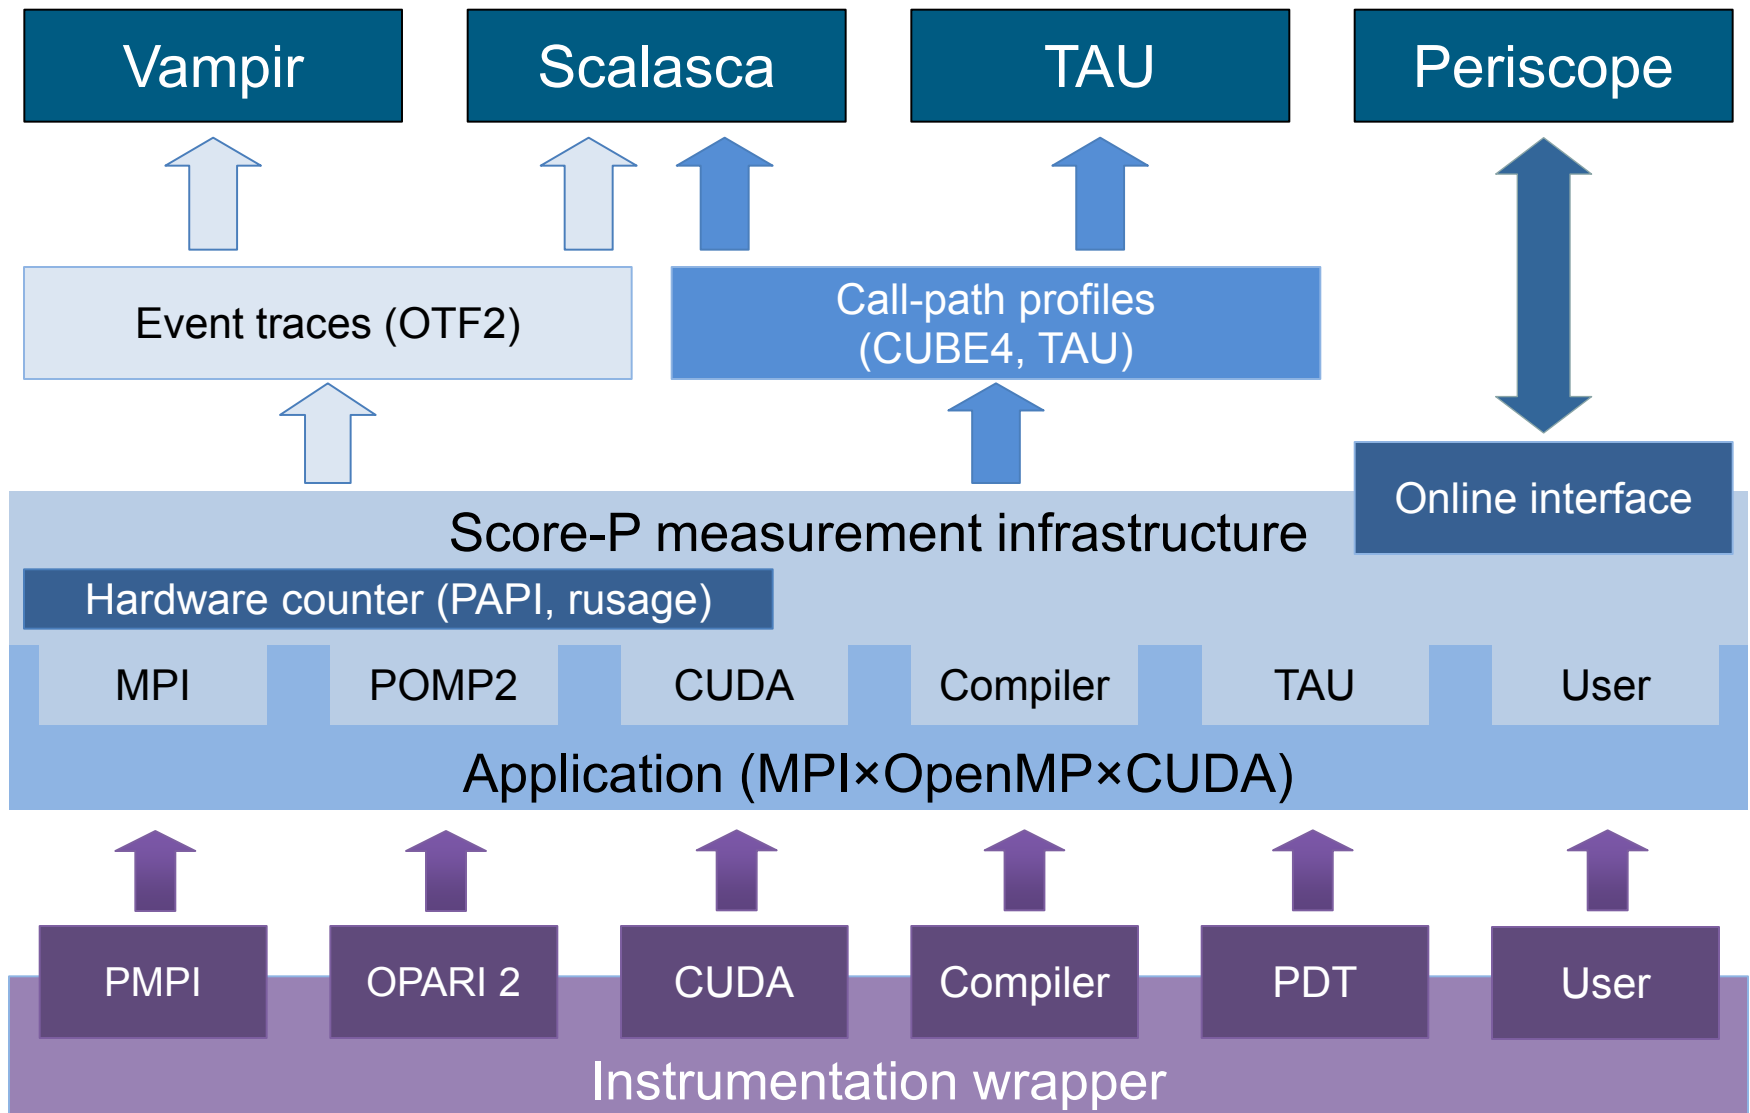
- Forschungszentrum Jülich, Germany
- German Research School for Simulation Sciences, Aachen, Germany
- Gesellschaft für numerische Simulation mbH Braunschweig, Germany
- RWTH Aachen, Germany
- Technische Universität Dresden, Germany
- Technische Universität München, Germany
- University of Oregon, Eugene, USA



UNIVERSITY OF OREGON

- Provide typical functionality for HPC performance tools
- Support all fundamental concepts of partner's tools
- Instrumentation (various methods)
- Flexible measurement without re-compilation:
  - Basic and advanced profile generation
  - Event trace recording
  - Online access to profiling data
- MPI, OpenMP, and hybrid parallelism (and serial)
- Enhanced functionality (OpenMP 3.0, CUDA, highly scalable I/O)

- Functional requirements
  - Generation of call-path profiles and event traces
  - Using direct instrumentation, later also sampling
  - Recording time, visits, communication data, hardware counters
  - Access and reconfiguration also at runtime
  - Support for MPI, OpenMP, basic CUDA, and all combinations
    - Later also OpenCL/HMPP/PTHREAD/...
- Non-functional requirements
  - Portability: all major HPC platforms
  - Scalability: petascale
  - Low measurement overhead
  - Easy and uniform installation through UNITE framework
  - Robustness
  - Open Source: New BSD License



- Scalability to maximum available CPU core count
- Support for OpenCL, HMPP, PTHREAD
- Support for sampling, binary instrumentation
- Support for new programming models, e.g., PGAS
- Support for new architectures
- Ensure a single official release version at all times which will always work with the tools
- Allow experimental versions for new features or research
- Commitment to joint long-term cooperation



# VI-HPS

SOFTWARE



0.00 <<time step loop>>  
0.00 updatedt  
6.62 updatex  
372.85 updateien  
0.00 gene  
0.00 <<iteration loop>>  
293.65 genbc



PRODUCTIVITY

FAST SOLUTIONS

- ☒ PAPI\_L1\_DCM
- ☒ PAPI\_L1\_ICM
- ☐ PAPI\_L2\_DCM
- ☒ PAPI\_L2\_ICM
- ☒ PAPI\_L2\_TCM
- ☐ PAPI\_L2\_TCM

## Score-P hands-on: NPB-MZ-MPI / BT

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

- Load modules

```
% module use /gpfslocal/pub/vihps/UNITE/local  
  
% module load VI-HPS-TW  
VI-HPS-TW loaded  
(load) intel version 13.0.1  
(load) intelmpi version 4.0.3  
papi/5.1.0 loaded  
scorep/1.1.1-intel13-intelmpi-papi-static loaded  
cube4/4.1.6-gnu loaded  
scalasca/2.0b3-intel13-intelmpi loaded  
tau/2.22.2-intelmpi-intel13 loaded  
vampir/8.1-dev loaded
```

- Change to directory containing NPB BT-MZ sources

- Edit `config/make.def` to adjust build configuration
  - Modify specification of compiler/linker: `MPIF77`

```
#           SITE- AND/OR PLATFORM-SPECIFIC DEFINITIONS
#-----
# Items in this file may need to be changed for each platform.
#-----
...
#-----
# The Fortran compiler used for MPI programs
#-----
#MPIF77 = mpiifort -fpp

# Alternative variants to perform instrumentation
...
MPIF77 = scorep mpiifort

# This links MPI Fortran programs; usually the same as ${MPIF77}
FLINK    = $(MPIF77)
...
```

Uncomment the  
Score-P compiler  
wrapper specification

- Return to root directory and clean-up

```
% make clean
```

- Re-build executable using Score-P instrumenter

```
% make bt-mz CLASS=B NPROCS=4
cd BT-MZ; make CLASS=B NPROCS=4 VERSION=
make: Entering directory 'BT-MZ'
cd ../sys; cc -o setparams setparams.c -lm
../sys/setparams bt-mz 4 B
scorep mpiifort -c -O3 -openmp bt.f
[...]
cd ../common; scorep mpiifort -c -O3 -fopenmp timers.f
scorep mpiifort -O3 -openmp -o ../bin.scorep/bt-mz_B.4 \
bt.o initialize.o exact_solution.o exact_rhs.o set_constants.o \
adi.o rhs.o zone_setup.o x_solve.o y_solve.o exch_qbc.o \
solve_subs.o z_solve.o add.o error.o verify.o mpi_setup.o \
../common/print_results.o ../common/timers.o
Built executable ../bin.scorep/bt-mz_B.4
make: Leaving directory 'BT-MZ'
```

- Score-P measurements are configured via environment variables:

```
% scorep-info config-vars --full
SCOREP_ENABLE_PROFILING
  Description: Enable profiling
  [...]
SCOREP_ENABLE_TRACING
  Description: Enable tracing
  [...]
SCOREP_TOTAL_MEMORY
  Description: Total memory in bytes for the measurement system
  [...]
SCOREP_EXPERIMENT_DIRECTORY
  Description: Name of the experiment directory
  [...]
SCOREP_FILTERING_FILE
  Description: A file name which contain the filter rules
  [...]
SCOREP_METRIC_PAPI
  Description: PAPI metric names to measure
  [...]
SCOREP_METRIC_RUSAGE
  Description: Resource usage metric names to measure
  [...] More configuration variables ...
```

- Change to the directory containing the new executable before running it and adjust configuration

```
% cd bin.scorep
% cp ../jobscript/mds/run_scorep.ll .
% vim run_scorep.ll
```

```
...
#
# Score-P configuration
#
export SCOREP_EXPERIMENT_DIRECTORY=scorep_bt-mz_B_4x4_sum
...
```

- Run executable

```
% llsubmit run_scorep.ll
% cat NPB_B.<id>
NAS Parallel Benchmarks (NPB3.3-MZ-MPI) - BT-MZ MPI+OpenMP Benchmark

Number of zones:      8 x      8
Iterations: 200      dt:    0.000300
Number of active processes:      4

Use the default load factors with threads
Total number of threads:      16  (  4.0 threads/process)

Calculated speedup =      15.96

Time step      1

[... More application output ...]
```



- Creates experiment directory `./scorep_bt-mz_B_4x4_sum` containing
  - a record of the measurement configuration (`scorep.cfg`)
  - the analysis report that was collated after measurement (`profile.cubex`)

```
% ls
... scorep_bt-mz_B_4x4_sum
% ls scorep_bt-mz_B_4x4_sum
profile.cubex scorep.cfg
```

- Interactive exploration with CUBE / ParaProf

```
% cube scorep_bt-mz_B_4x4_sum/profile.cubex

[CUBE GUI showing summary analysis report]

% paraprof scorep_bt-mz_B_4x4_sum/profile.cubex

[TAU ParaProf GUI showing summary analysis report]
```

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

- 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 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

- Report scoring as textual output

```
% scorep-score scorep_bt-mz_B_4x4_sum/profile.cubex
Estimated aggregate size of event trace (total_tbc):
Estimated requirements for largest trace buffer (max_tbc):
(hint: When tracing set SCOREP_TOTAL_MEMORY > max_tbc to avoid intermediate flushes
or reduce requirements using file listing names of USR regions to be filtered.)
```


flt type	max_tbc	time	% region
ALL	9043348074	933.55	100.0 ALL
USR	9025830154	450.52	48.3 USR
OMP	16431872	480.67	51.5 OMP
COM	997150	0.67	0.1 COM
MPI	88898	1.69	0.2 MPI

35955109198 bytes  
9043348074 bytes

33.5 GB total memory  
8.4 GB per rank!

- Region/callpath classification

- MPI (pure MPI library functions)
- OMP (pure OpenMP functions/regions)
- USR (user-level source local computation)
- COM (“combined” USR + OpenMP/MPI)
- ANY/ALL (aggregate of all region types)




- Score report breakdown by region

```
% scorep-score -r scorep_bt-mz_B_4x4_sum/profile.cubex
[...]
```

flt type	max_tbc	time	% region
ALL	9043348074	933.55	100.0 ALL
USR	9025830154	450.52	48.3 USR
OMP	16431872	480.67	51.5 OMP
	997150	0.67	0.1 COM
	88898	1.69	0.2 MPI
USR	2894950740	137.99	14.8 matmul_sub_
USR	2894950740	119.71	12.8 matvec_sub_
USR	2894950740	175.59	18.8 binvcrhs_
USR	127716204	6.08	0.7 binvrhs_
USR	127716204	7.38	0.8 lhsinit_
USR	94933520	3.76	0.4 exact_solution_
OMP	771840	0.05	0.0 !\$omp parallel @exch_...
OMP	771840	0.04	0.0 !\$omp parallel @exch_...
OMP	771840	0.05	0.0 !\$omp parallel @exch_...
[...]			

More than  
8 GB just for  
these 6 regions



- Summary measurement analysis score reveals
  - Total size of event trace would be ~34 GB
  - Maximum trace buffer size would be ~8.5 GB per rank
    - smaller buffer would require flushes to disk during measurement resulting in substantial perturbation
  - 99.8% 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 32% 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

- Report scoring with prospective filter listing 6 USR regions

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

% scorep-score -f ../config/scorep.filt scorep_bt-mz_B_4x4_sum/profile.cubex
Estimated aggregate size of event trace (total_tbc): 70086838 bytes
Estimated requirements for largest trace buffer (max_tbc): 17521726 bytes
(hint: When tracing set SCOREP_TOTAL_MEMORY > max_tbc to avoid intermediate flushes
or reduce requirements using file listing names of USR regions to be filtered.)
```

67 MB of memory in total,  
17 MB per rank!

- Score report breakdown by region

```
% scorep-score -r -f ../config/scorep.filt \
> scorep_bt-mz_B_4x4_sum/profile.cubex
```

flt	type	max_tbc	time	%	region
*	ALL	17521726	483.03	51.7	ALL-FLT
+	FLT	9025826370	450.51	48.3	FLT
-	OMP	16431872	480.67	51.5	OMP-FLT
*	COM	997150	0.67	0.1	COM-FLT
-	MPI	88898	1.69	0.2	MPI-FLT
*	USR	3806	0.00	0.0	USR-FLT
+	USR	2894950740	137.99	14.8	matmul_sub_
+	USR	2894950740	119.71	12.8	matvec_sub_
+	USR	2894950740	175.59	18.8	binvcrhs_
+	USR	127716204	6.08	0.7	binvrhs_
+	USR	127716204	7.38	0.8	lhsinit_
+	USR	94933520	3.76	0.4	exact_solution_
-	OMP	771840	0.05	0.0	!\$omp parallel @exch_...
-	OMP	771840	0.04	0.0	!\$omp parallel @exch_...
-	OMP	771840	0.05	0.0	!\$omp parallel @exch_...
[...]					

Filtered  
routines  
marked  
with '+'



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

```
% vim run_scorep.ll
```

- Adjust configuration

```
...  
export SCOREP_EXPERIMENT_DIRECTORY=scorep_bt-mz_B_4x4_sum_with_filter  
export SCOREP_FILTERING_FILE=../config/scorep.filt  
...
```

- Submit job

```
% llsubmit run_scorep.ll
```

- Scoring of new analysis report as textual output

```
% scorep-score scorep_bt-mz_B_4x4_sum_with_filter/profile.cubex
Estimated aggregate size of event trace (total_tbc):      70086838 bytes
Estimated requirements for largest trace buffer (max_tbc): 17521726 bytes
(hint: When tracing set SCOREP_TOTAL_MEMORY > max_tbc to avoid intermediate flushes
or reduce requirements using file listing names of USR regions to be filtered.)
```

flt type	max_tbc	time	% region
ALL	17521726	215.07	100.0 ALL
OMP	16431872	212.86	99.0 OMP
COM	997150	0.68	0.3 COM
MPI	88898	1.54	0.7 MPI
USR	3806	0.00	0.0 USR

- Significant reduction in runtime (measurement overhead)
  - Not only reduced time for USR regions, but MPI/OMP reduced too!
- Further measurement tuning (filtering) may be appropriate
  - e.g., use “timer\_\*” to filter timer\_start\_, timer\_read\_, etc.

- Recording hardware counters via PAPI

```
export SCOREP_METRIC_PAPI=PAPI_L2_TCM,PAPI_FP_OPS
```

- Also possible to record them only per rank

```
export SCOREP_METRIC_PAPI_PER_PROCESS=PAPI_L3_TCM
```

- Recording operating system resource usage

```
export SCOREP_METRIC_RUSAGE_PER_PROCESS=ru_maxrss,ru_stime
```

Note: Additional memory is needed to store metric values.  
Therefore, you may have to adjust SCOREP\_TOTAL\_MEMORY,  
for example as reported using “scorep-score -c”

- Available PAPI metrics
  - Preset events: common set of events deemed relevant and useful for application performance tuning
    - Abstraction from specific hardware performance counters, mapping onto available events done by PAPI internally

```
% papi_avail
```

- Native events: set of all events that are available on the CPU (**platform dependent**)

```
% papi_native_avail
```

Note:

Due to hardware restrictions

- number of concurrently measured events is limited
- there may be unsupported combinations of concurrent events

- Available resource usage metrics

```
% man getrusage
```

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

```
struct rusage {  
    struct timeval ru_utime; /* user CPU time used */  
    struct timeval ru_stime; /* system CPU time used */  
    long    ru_maxrss;      /* maximum resident set size */  
    long    ru_ixrss;       /* integral shared memory size */  
    long    ru_idrss;       /* integral unshared data size */  
    long    ru_isrss;       /* integral unshared stack size */  
    long    ru_minflt;      /* page reclaims (soft page faults) */  
    long    ru_majflt;      /* page faults (hard page faults) */  
    long    ru_nswap;       /* swaps */  
    long    ru_inblock;     /* block input operations */  
    long    ru_oublock;     /* block output operations */  
    long    ru_msgsnd;      /* IPC messages sent */  
    long    ru_msgrcv;      /* IPC messages received */  
    long    ru_nsignals;    /* signals received */  
    long    ru_nvcsw;       /* voluntary context switches */  
    long    ru_nivcsw;      /* involuntary context switches */  
};
```

```
[... More output ...]
```

Note:

- (1) Not all fields are maintained on each platform.
- (2) Check scope of metrics (per process vs. per thread)

- Record only for subset of the MPI functions events

```
export SCOREP_MPI_ENABLE_GROUPS=cg,coll,p2p,xnonblock
```

- All possible sub-groups

- cg Communicator and group management
- coll Collective functions
- env Environmental management
- err MPI Error handling
- ext External interface functions
- io MPI file I/O
- misc Miscellaneous
- perf PControl
- p2p Peer-to-peer communication
- rma One sided communication
- spawn Process management
- topo Topology
- type MPI datatype functions
- xnonblock Extended non-blocking events
- xreqtest Test events for uncompleted requests

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**

- Traces can become extremely large and unwieldy
  - Size is proportional to number of processes/threads (width), duration (length) and detail (depth) of measurement
- Traces containing intermediate flushes are of little value
  - Uncoordinated flushes result in cascades of distortion
  - Reduce size of trace
  - Increase available buffer space
- Traces should be written to a parallel file system
  - /work or /scratch are typically provided for this purpose
- Moving large traces between file systems is often impractical
  - However, systems with more memory can analyze larger traces
  - Alternatively, run trace analyzers with undersubscribed nodes



- Edit `config/make.def` to adjust build configuration
  - Modify specification of compiler/linker: `MPIF77`

```
#                SITE- AND/OR PLATFORM-SPECIFIC DEFINITIONS
#-----
# Items in this file may need to be changed for each platform.
#-----
...
#-----
# The Fortran compiler used for MPI programs
#-----
#MPIF77 = mpiifort -fpp

# Alternative variants to perform instrumentation
...
#MPIF77 = scorep mpiifort
MPIF77 = scorep --user mpiifort

# This links MPI Fortran programs; usually the same as ${MPIF77}
FLINK    = $(MPIF77)
...

```

Uncomment the  
Score-P compiler  
wrapper specification

- Return to root directory and clean-up

```
% make clean
```

- Re-build executable using Score-P compiler wrapper

```
% make bt-mz CLASS=B NPROCS=4
cd BT-MZ; make CLASS=B NPROCS=4 VERSION=
make: Entering directory 'BT-MZ'
cd ../sys; cc -o setparams setparams.c -lm
../sys/setparams bt-mz 4 B
scorep --user mpiifort -c -O3 -openmp bt.f
[...]
cd ../common; scorep --user mpiifort -c -O3 -fopenmp timers.f
scorep --user mpiifort -O3 -openmp -o ../bin.scorep/bt-mz_B.4 \
bt.o initialize.o exact_solution.o exact_rhs.o set_constants.o \
adi.o rhs.o zone_setup.o x_solve.o y_solve.o exch_qbc.o \
solve_subs.o z_solve.o add.o error.o verify.o mpi_setup.o \
../common/print_results.o ../common/timers.o
Built executable ../bin.scorep/bt-mz_B.4
make: Leaving directory 'BT-MZ'
```

- Score-P measurements are configured via environmental variables:

```
% scorep-info config-vars --full
SCOREP_ENABLE_PROFILING
  Description: Enable profiling
  [...]
SCOREP_ENABLE_TRACING
  Description: Enable tracing
  [...]
SCOREP_TOTAL_MEMORY
  Description: Total memory in bytes for the measurement system
  [...]
SCOREP_EXPERIMENT_DIRECTORY
  Description: Name of the experiment directory
  [...]
SCOREP_FILTERING_FILE
  Description: A file name which contain the filter rules
  [...]
SCOREP_METRIC_PAPI
  Description: PAPI metric names to measure
  [...]
SCOREP_METRIC_RUSAGE
  Description: Resource usage metric names to measure
  [...] More configuration variables ...
```

- Re-run the application using the tracing mode of Score-P
  - Edit `run_scorep.ll` to adjust configuration

```
export SCOREP_EXPERIMENT_DIRECTORY=scorep_bt-mz_B_4x4_trace
export SCOREP_FILTERING_FILE=../config/scorep.filt
export SCOREP_ENABLE_TRACING=true
export SCOREP_ENABLE_PROFILING=false
export SCOREP_METRIC_PAPI=PAPI_L2_TCM,PAPI_FP_OPS
export SCOREP_TOTAL_MEMORY=50M
```

- Submit job

```
% llsubmit run_scorep.ll
```

- Separate trace file per thread written straight into new experiment directory `./scorep_bt-mz_B_4x4_trace`
- Interactive trace exploration with Vampir

```
% vampir scorep_bt-mz_B_4x4_trace/traces.otf2
```

- Can be used to mark initialization, solver & other phases
  - Annotation macros ignored by default
  - Enabled with [**--user**] flag
- Appear as additional regions in analyses
  - Distinguishes performance of important phase from rest
- Can be of various type
  - E.g., function, loop, phase
  - See user manual for details
- Available for Fortran / C / C++

```
#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

```
#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... */
}
```

```
#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...
}
```



- 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++

### Score-P

- Community instrumentation & measurement infrastructure
  - Instrumentation (various methods)
  - Basic and advanced profile generation
  - Event trace recording
  - Online access to profiling data
- Available under New BSD open-source license
- Documentation & Sources:
  - <http://www.score-p.org>
- User guide also part of installation:
  - `<prefix>/share/doc/scorep/{pdf,html}/`
- Contact: [info@score-p.org](mailto:info@score-p.org)
- Bugs: [scorep-bugs@groups.tu-dresden.de](mailto:scorep-bugs@groups.tu-dresden.de)