Score-P instrumentation and measurement infrastructure

Demo/Hands-on: Instrumentation & initial measurement

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 analysis & report examination

Toolchain and Score-P modules (COSMA)

- Select modules for the GCC + OpenMPI tool chain

```
% module load  gnu_comp/14.1.0 openmpi/5.0.3
```

- Load Score-P and Cube modules
  - Score-P installation is toolchain specific!

```
% module load  scorep/9.4 cube/4.9.1
```

NPB-MZ-MPI / BT instrumentation

```
#--------------------------------------------------------------
# The Fortran compiler used for MPI programs
#--------------------------------------------------------------
#FC = mpif90

# Alternative variants to perform instrumentation
...
FC = scorep --user mpif90

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

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

Prefix the compiler by the Score-P instrumenter command

NPB-MZ-MPI / BT instrumented build

```
% make clean

% make bt-mz CLASS=B
cd BT-MZ; make CLASS=B VERSION=
make: Entering directory 'BT-MZ'
cd ../sys; cc  -o setparams setparams.c -lm
../sys/setparams bt-mz B
scorep --user mpif90 -g -c  -O3 -fopenmp  bt_scorep_user.F90
 [...]
cd ../common;  scorep --user mpif90 -g -c  -O3 -fopenmp  timers.f90
 [...]
scorep --user mpif90 -g -O3 -fopenmp -o ../bin.scorep/bt-mz_B.x \
bt_scorep_user.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.x
make: Leaving directory 'BT-MZ'
```

- Return to root directory and clean-up
- Re-build executable using Score-P compiler wrapper

Measurement configuration: scorep-info

```
% 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 ...]
```

- Score-P measurements are configured via environmental variables

Summary measurement collection

```
% cd bin.scorep
% cp ../jobscript/bluefield1/scorep.sbatch .
% vim scorep.sbatch
...
# set up environment
module purge
module load gnu_comp openmpi
module load scorep

# measurement configuration
export SCOREP_EXPERIMENT_DIRECTORY=scorep_bt-mz_sum
#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 ./bt-mz_B.x

% sbatch scorep.sbatch
```

- Change to the directory containing the new executable before running it with the desired configuration
- Check settings

Leave these lines commented out for the moment

- Submit job

Summary measurement collection

```
% less npb-btmz.o<job_id>

 NAS Parallel Benchmarks (NPB3.4-MZ MPI+OpenMP) - BT-MZ \
>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:    24   (  6.0 threads/process)

 Calculated speedup = 23.94

 Time step    1

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

- Check the output of the application run

BT-MZ summary analysis report examination

```
% ls
bt-mz_B.x slurm-<job_id>.out scorep_bt-mz_sum/

% ls scorep_bt-mz_sum
MANIFEST.md profile.cubex  scorep.cfg




% cube scorep_bt-mz_sum/profile.cubex

      [CUBE GUI showing summary analysis report]
```

- Creates experiment directory including
  - A brief content overview (MANIFEST.md)
  - A record of the measurement configuration (scorep.cfg)
  - The analysis report that was collated after measurement (profile.cubex)
- Interactive exploration with Cube

Hint:
Copy ‘profile.cubex’ to local system (laptop) using ‘scp’ to improve responsiveness of GUI

Reference results available:
`/cosma/apps/do009/shared/Scalasca/experiments`