

## International HPC Summer School 2023: Performance analysis and optimization

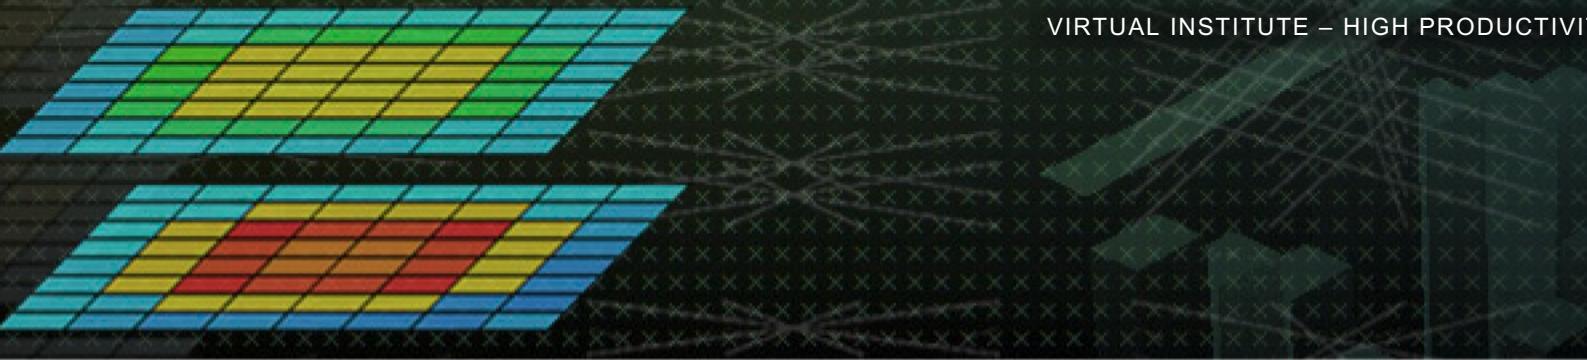
Score-P – A Joint Performance Measurement Run-Time  
Infrastructure for Periscope, Scalasca, TAU, and Vampir

---

VI-HPS Team

Ilya Zhukov – Jülich Supercomputing Centre





## Hands-on: NPB-MZ-MPI / BT



# 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
- 3.0 Trace measurement with filtering
- 3.1 Event trace examination & analysis

## Recap: Local installation

---

- VI-HPS tools not yet installed system-wide
  - Enable and load local tool installations
  - Required for each shell session

```
% module use /jet/home/zhukov/ihpcss23/modules/  
% module load scalasca/2.6-intel_openmpi
```

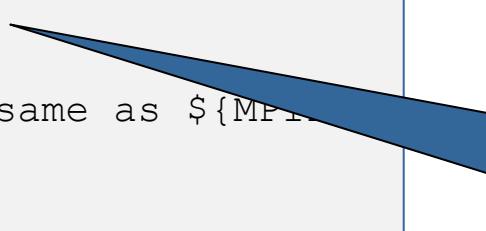
- Copy tutorial sources to your working directory

```
% cd $PROJECT  
% tar zxvf /jet/home/zhukov/ihpcss23/tutorial/NPB3.3-MZ-MPI.tar.gz  
% cd NPB3.3-MZ-MPI
```

# NPB-MZ-MPI / BT instrumentation

```
#-----  
# The Fortran compiler used for MPI programs  
#-----  
#MPIF77 = mpif77  
  
# Alternative variants to perform instrumentation  
...  
MPIF77 = scorep mpif77  
  
# This links MPI Fortran programs; usually the same as ${MPI_F77}  
FLINK    = $(MPIF77)  
...
```

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



Uncomment the Score-P compiler wrapper specification

# NPB-MZ-MPI / BT instrumented build

```
% make clean  
  
% make bt-mz CLASS=C NPROCS=8  
cd BT-MZ; make CLASS=C NPROCS=8 VERSION=  
make: Entering directory 'BT-MZ'  
cd ../sys; cc -o setparams setparams.c -lm  
../sys/setparams bt-mz 8 C  
scorep mpif77 -c -O3 -fopenmp bt.f  
[...]  
cd ..;/common; scorep mpif77 -c -O3 \  
-fopenmp timers.f  
scorep mpif77 -O3 -fopenmp \  
-o ./bin.scorep/bt-mz_C.8 \  
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_C.8  
make: Leaving directory 'BT-MZ'
```

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

# Summary measurement collection

```
% cd bin.scorep
% cp ..../jobscript/bridges2/scorep.sbatch.C.8 .
% less scorep.sbatch.C.8

...
# Score-P measurement configuration
#export SCOREP_EXPERIMENT_DIRECTORY=\
#scorep_bt-mz_${CLASS}_${PROCS}x${OMP_NUM_THREADS}_profile
#export SCOREP_FILTERING_FILE=../config/scorep.filt
#export SCOREP_TOTAL_MEMORY=65M
#export SCOREP_METRIC_PAPI=PAPI_TOT_INS,PAPI_TOT_CYC
#export SCOREP_ENABLE_TRACING=true

scalasca -analyze mpirun \
"--map-by ppr:$TASKS_PER_SOCKET:socket:pe=$SLURM_CPUS_PER_TASK \
-n $SLURM_NTASKS" $EXE

%
% sbatch scorep.sbatch.C.8
```

- 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

# 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

```
% less profile-C.8-<job_id>.out
```

```
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: 8
```

```
Use the default load factors with threads
```

```
Total number of threads: 48 ( 6.0 threads/process)
```

```
Calculated speedup = 47.97
```

```
Time step 1
```

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

- Check the output of the application run
- Compare run time of reference and this run
- What was a run time dilation?

# BT-MZ summary analysis report examination

```
% ls -1  
bt-mz_C.8  
scorep.sbatch.C.8  
profile-C.8-<jobid>.err  
profile-C.8-<jobid>.out  
scorep_bt-mz_C_8x6_profile/  
  
% ls scorep_bt-mz_C_8x6_sum/  
MANIFEST.md  
profile.cubex  
scorep.cfg  
  
% cube profile.cubex
```

[Cube GUI showing summary analysis report]

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

## Hint:

Copy ‘profile.cubex’ to your laptop using ‘scp’ to improve responsiveness of GUI