

The Caliper Profiling Library

41st VI-HPS Tuning Workshop

Feb 7, 2022



David Boehme



LLNL-PRES-821032

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore
National Laboratory

Caliper: A Performance Instrumentation and Profiling Library

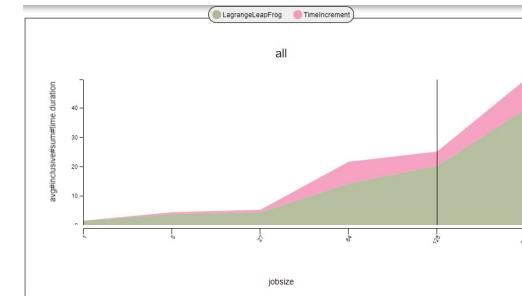
- Integrates a performance profiler into your program
 - Profiling is always available
 - Simplifies performance profiling for application end users
- Common instrumentation interface
 - Provides program context information for other tools
- Advanced profiling features
 - MPI, CUDA, ROCm, Kokkos support; call-stack sampling; hardware counters; memory profiling

Caliper Use Cases

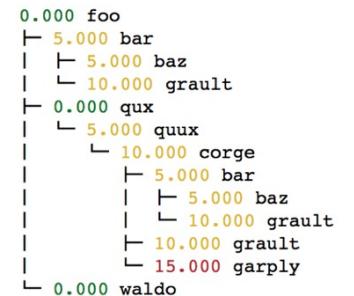
- Lightweight always-on profiling
 - Performance summary report for each run
- Performance debugging
- Performance introspection
- Comparison studies across runs
 - Performance regression testing
 - Configuration and scaling studies
- Automated workflows

Performance reports

Path	Min time/rank	Max time/rank	Avg time/rank	Time %
main	0.000119	0.000119	0.000119	7.079120
mainloop	0.000067	0.000067	0.000067	3.985723
foo	0.000646	0.000646	0.000646	38.429506
init	0.000017	0.000017	0.000017	1.011303

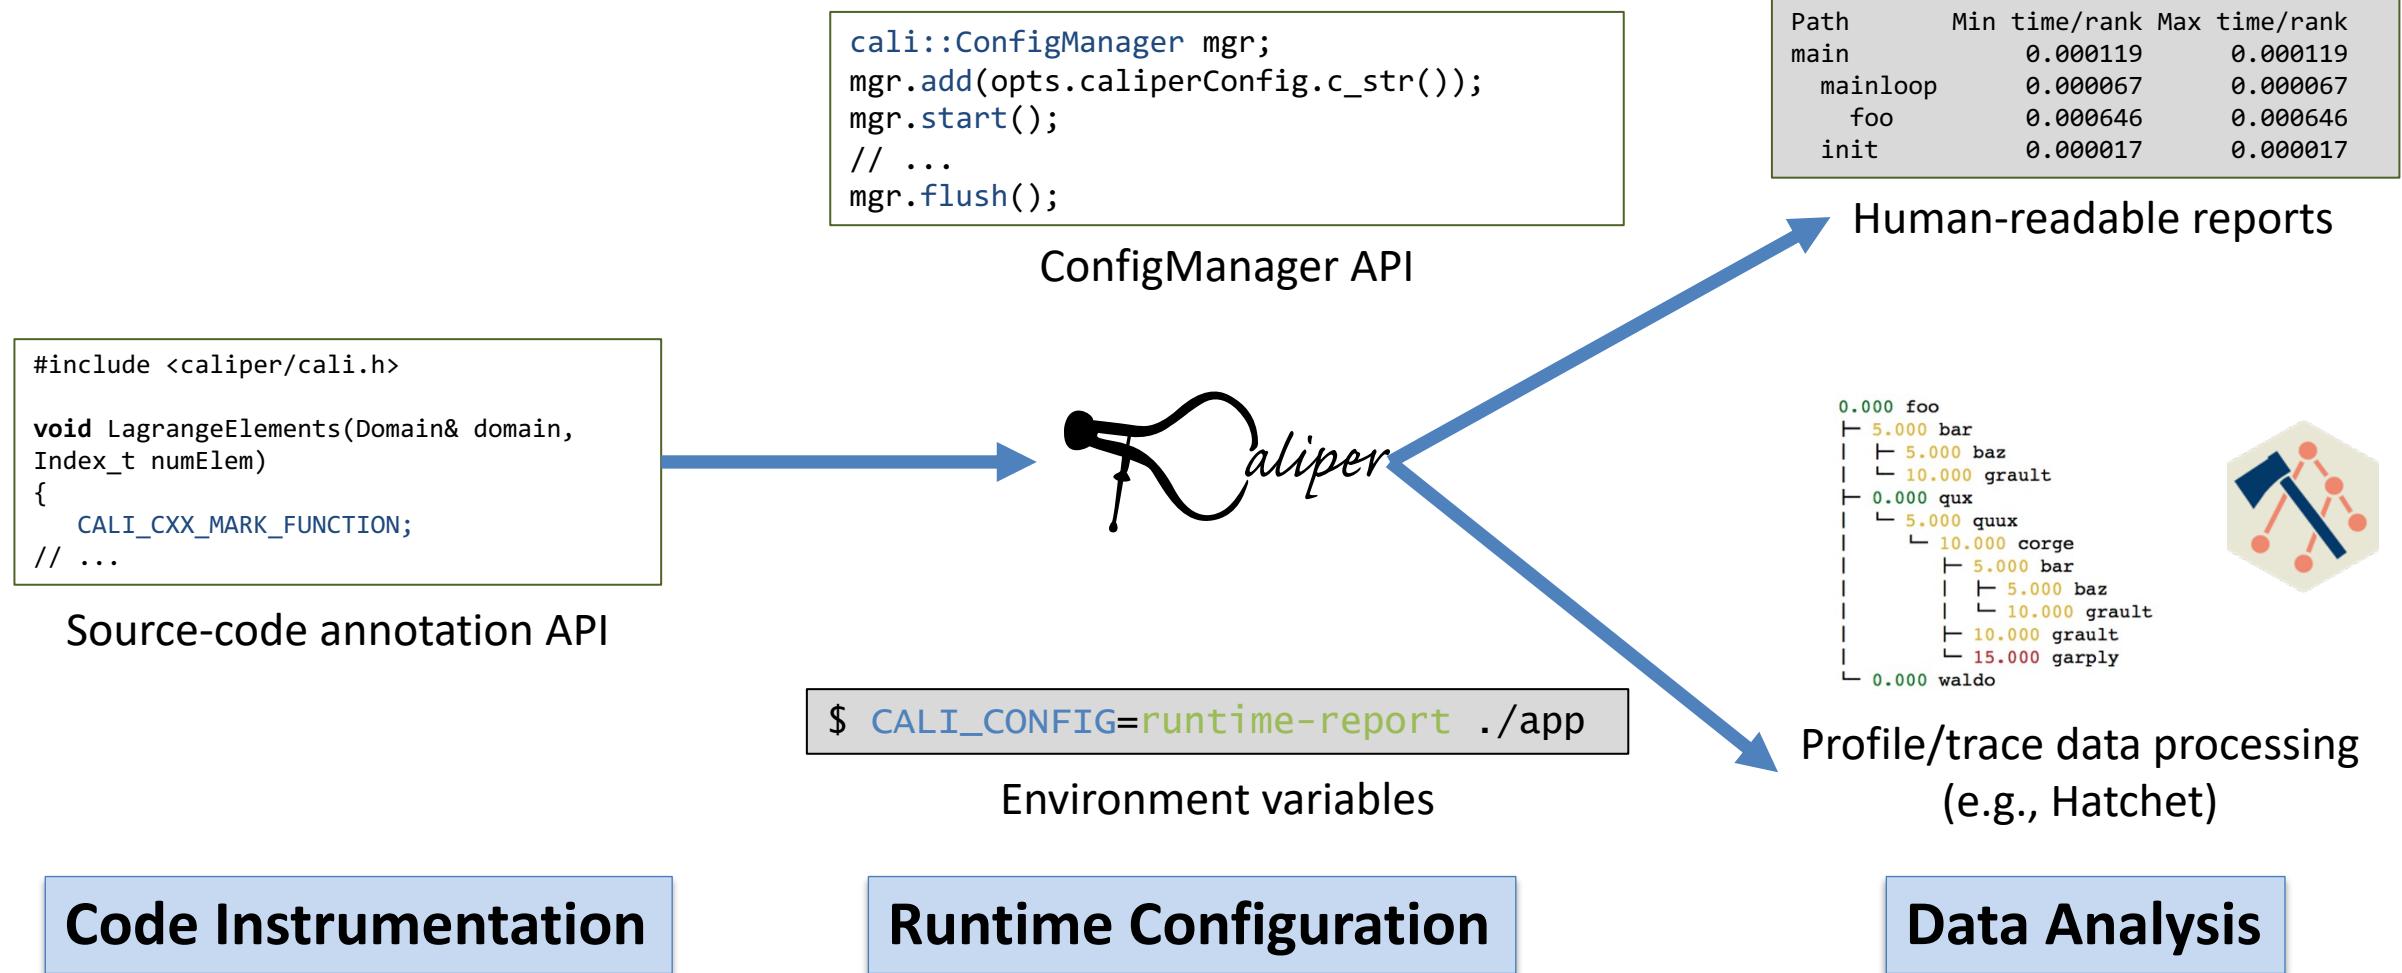


Comparing runs



Debugging

Using Caliper: Workflow



Region Profiling: Marking Code Regions

C/C++

```
#include <caliper/cali.h>

void main() {
    CALI_MARK_BEGIN("init");

    do_init();

    CALI_MARK_END("init");
}
```

Fortran

```
USE caliper_mod

CALL cali_begin_region('init')

CALL do_init()

CALL cali_end_region('init')
```

- Use annotation macros (C/C++) or functions to mark and name code regions

Region Profiling: Best Practices

- Be selective: Instrument high-level program subdivisions (kernels, phases, ...)
- Be clear: Choose meaningful names
- Start small: Add instrumentation incrementally

```
RAJA::ReduceSum<RAJA::omp_reduce, double> ompdot(0.0);

CALI_MARK_BEGIN("dotproduct");

RAJA::forall<RAJA::omp_parallel_for_exec>(RAJA::RangeSegment(0, N), [=] (int i) {
    ompdot += a[i] * b[i];
});
dot = ompdot.get();

CALI_MARK_END("dotproduct");
```

Caliper annotations give meaningful names to high-level program constructs

Region Profiling: Printing a Runtime Report

```
$ cd Caliper/build  
$ make cxx-example  
$ CALI_CONFIG=runtime-report ./examples/apps/cxx-example
```

Path	Min time/rank	Max time/rank	Avg time/rank	Time %
main	0.000119	0.000119	0.000119	7.079120
mainloop	0.000067	0.000067	0.000067	3.985723
foo	0.000646	0.000646	0.000646	38.429506
init	0.000017	0.000017	0.000017	1.011303

- Set the CALI_CONFIG environment variable to access Caliper's built-in profiling configurations
- “runtime-report” measures, aggregates, and prints time in annotated code regions

Control Profiling Programmatically: The ConfigManager API

```
#include <caliper/cali.h>
#include <caliper/cali-manager.h>

int main(int argc, char* argv[])
{
    cali::ConfigManager mgr;
    mgr.add(argv[1]);
    if (mgr.error())
        std::cerr << mgr.error_msg() << "\n";

    mgr.start();
    // ...
    mgr.flush();
}
```

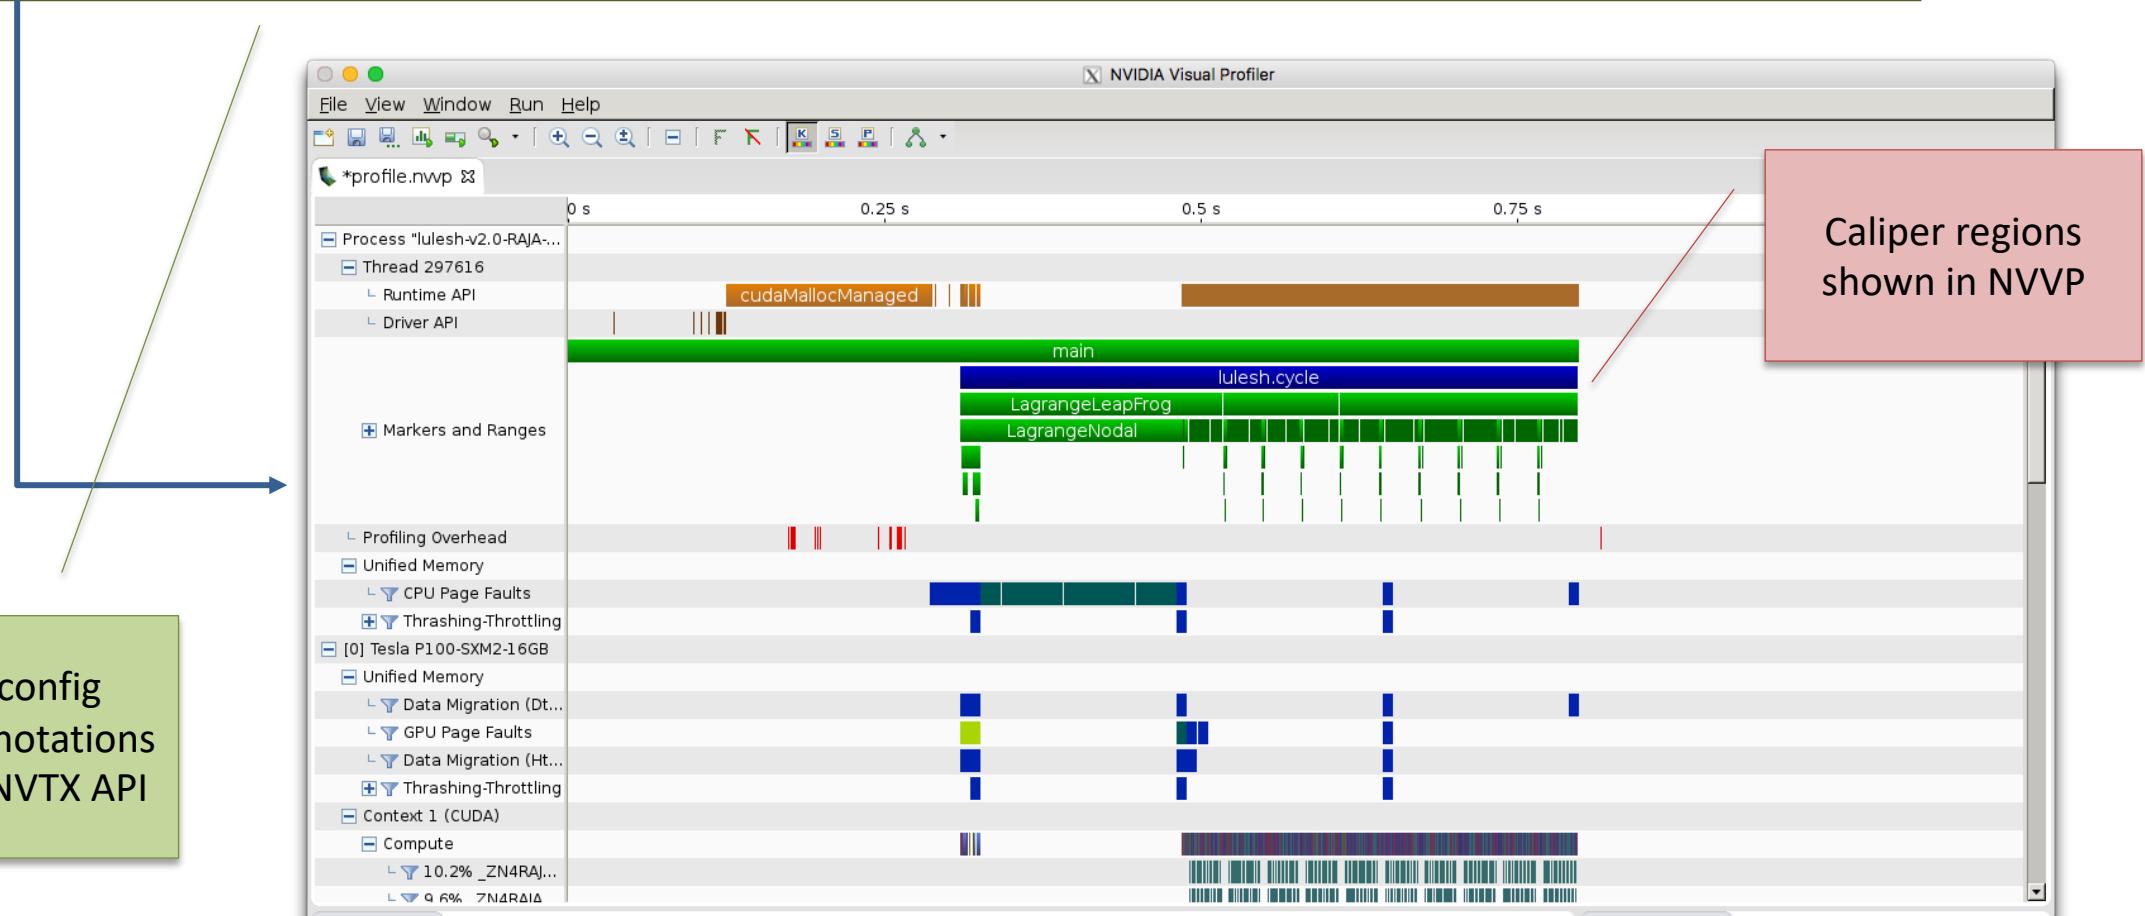
- Use ConfigManager to access Caliper's built-in profiling configurations
- Use add() to add profiling configurations (same config strings as CALI_CONFIG)
- Use start() to start profiling
- Use flush() to collect and write output

```
$ ./examples/apps/cxx-example -P runtime-report
```

- Now we can use command-line arguments or other program inputs to enable profiling

Forwarding Annotations to Third-Party Tools

```
$ CALI_CONFIG=nvtx nvprof <nvprof-opts> ./app
```



Call Graph Analysis with the Hatchet Python Library

- Caliper records machine-readable JSON data with hatchet-region-profile or hatchet-sample-profile

```
$ CALI_CONFIG=hatchet-sample-profile srun -n 8 ./lulesh2.0
```

Hatchet allows manipulation, computation, comparison, and visualization of call graph data

Manual Configuration Allows Custom Analyses

```
cali-query -q "select alloc.label#cuhti.fault.addr as Pool,  
cuhti.uvm.kind as UVM\ Event,  
scale(cuhti.uvm.bytes,1e-6) as MB,  
scale(cuhti.activity.duration,1e-9) as Time  
group by  
prop:nested,alloc.label#cuhti.fault.addr,cuhti.uvm.kind  
where cuhti.uvm.kind format tree" trace.cali
```

caliper.config

```
CALI_SERVICES_ENABLE=alloc,cuhti,cuhtitrace,mpi,trace,recorder  
CALI_ALLOC_RESOLVE_ADDRESSES=true  
CALI_CUHTI_CALLBACK_DOMAINS=sync  
CALI_CUHTITRACE_ACTIVITIES=uvm  
CALI_CUHTITRACE_CORRELATE_CONTEXT=false  
CALI_CUHTITRACE_FLUSH_ON_SNAPSHOT=true
```

Path				
main				
solve				
TIME_STEPPING				
enforceBC				
CURVI in EnforceBC				
CurviCartIC				
CurviCartIC::PART 3	Pool	UVM Event	MB	Time
curvilinear4sgwind	UM_pool	pagefaults.gpu		2.806946
curvilinear4sgwind	UM_pool	HtoD	7862.747136	0.232238
curvilinear4sgwind	UM_pool_temps	pagefaults.gpu		0.130167
curvilinear4sgwind	UM_pool	DtOH	9986.441216	0.378583
curvilinear4sgwind	UM_pool	pagefaults.cpu		

- Mapping CPU/GPU unified memory transfer events to Umpire memory pools in SW4

Recording Program Metadata: The Adiak Library

TeaLeaf_CUDA example [C++]

```
#include <adiak.hpp>

adiak::user();
adiak::launchdate();
adiak::jobsizes();

adiak::value("end_step", readInt(input, "end_step"));
adiak::value("halo_depth", readInt(input, "halo_depth"));

if (tl_use_ppcg) {
    adiak::value("solver", "PPCG");
// [...]
```

Use built-in Adiak functions to collect common metadata

Use key:value functions to collect program-specific data

- Use the [Adiak](#) C/C++ library to record program metadata
 - Environment info (user, launchdate, system name, ...)
 - Program configuration (input problem description, problem size, ...)
- Enables performance comparisons across runs

Contact & Links

- GitHub repository: <https://github.com/LLNL/Caliper>
- Documentation: <https://llnl.github.io/Caliper>
- GitHub Discussions: <https://github.com/LLNL/Caliper/discussions>
- Adiak: <https://github.com/LLNL/Adiak>
- Hatchet: <https://github.com/hatchet/hatchet>
- Contact: David Boehme (boehme3@llnl.gov)

Tutorial Environment on JUWELS

- Tutorial materials and example apps:

<https://github.com/daboehme/caliper-tutorial/tree/vihps-training-41>

- Build and setup on JUWELS:

```
$ module load CMake GCC CUDA ParaStationMPI
$ git clone --recursive https://github.com/daboehme/caliper-tutorial.git -b vihps-training-41
$ cd caliper-tutorial
$ . setup-env.sh juwels
```

- A pre-built installation with Caliper, Adiak, and the example apps is available:

```
$ export PATH=$PATH:/p/project/training2123/tools/caliper/bin
$ CALI_CONFIG=runtime-report lulesh2.0 -i 10
```

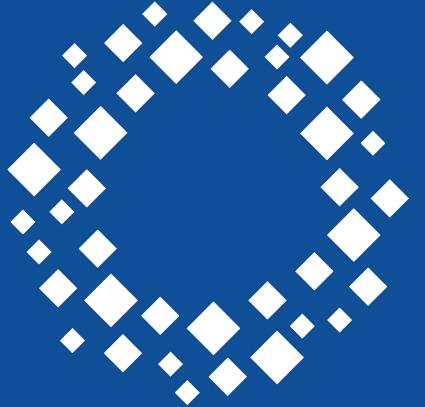
Using the training installation in your own Codes

- Note: Installation is built for GCC / ParaStationMPI / CUDA 11.3
- Classic build/make:

```
$ export CALI_DIR=/p/project/training2123/tools/caliper  
$ g++ -o app -L$(CALI_DIR)/lib -L$(CALI_DIR)/lib64 -ladiak -lcaliper -I$(CALI_DIR)/include *.c
```

- In CMake:

```
find_package(adiak)  
find_package(caliper)  
  
add_executable(myapp ${SOURCES})  
target_link_libraries(myapp PRIVATE adiak::adiak caliper)
```



CASC

Center for Applied
Scientific Computing



**Lawrence Livermore
National Laboratory**

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.