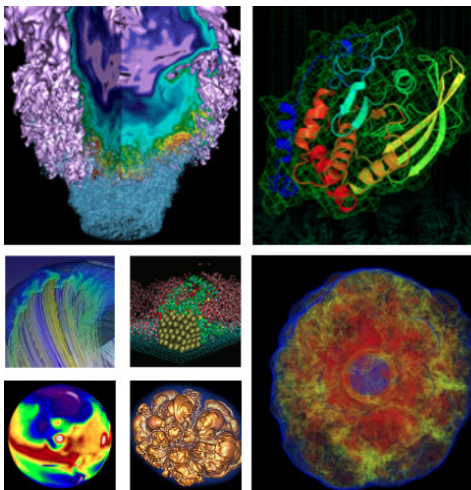


Geant4 + TiMemory

Automated Timing + Memory Analysis



National Energy Research
Scientific Computing Center



U.S. DEPARTMENT OF
ENERGY

Office of
Science



Jonathan R. Madsen

✉ jrmadsen@lbl.gov

National Energy Research Scientific Computing Center
Lawrence Berkeley National Laboratory

August 28, 2018

- **TiMemory** is a lightweight “in-situ” package for consistent and automated performance reports
 - Supplements profilers – profilers have overhead and that overhead skews/misrepresents run-time
- Cross-language \Rightarrow C, C++ , and Python implementations can be used independently or simultaneously
- Designed to be very easy to use
 - Uses macros e.g., `TIMEMORY_AUTO_TIMER(“”)` macro at beginning of function
 - When not compiled with **TiMemory** support this is an empty macro
 - insert “auto-timer” without parameter customization: function, file, and line-number automatically recorded
- Macro creates an object that starts measurements upon construction, stops measurements upon destruction, and accumulates results over multiple “laps”
- Records call-stack
- Support for PAPI hardware counters and CUDA event timers actively under development

- JSON serialization
- Built-in features (via Python)
 - Timing and memory plotting
 - Support for attaching analysis as `CTEST_NOTES` and attaching images as `<DartMeasurementFile>` to CDash dashboard
- MPI-compatible
 - No data merging occurs other than one combined report, each MPI process has independent data
- Thread-safe
 - Automated merging of thread-local data when report is requested or when thread exits
- Termination signal handling (SIGHUP, SIGQUIT, etc.) to dump “profile” when tracking down memory and/or timing issues

Log in | Dashboard | Dashboard | Back | Previous | Current | Project
Friday, April 13 2018 15:43:18

Site: [Home](#)
Build Name: [master] Linux x86_64 [Cdash 5.0.2][Helm][python 3.4][C++11][MPI]
Stamp: 2018-04-13-09:15:00 [Continuous]

- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_army_test.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/status_report.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_report.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_report.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_decorator.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_debug.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_toggle.out](#)
- [2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_context_manager.out](#)

2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/timing_army_test.out

```

> [pyc] main[array_test.py]:285      : 4.142 wall, 1.178 user + 2.970 system = 3.140 CPU [sec] [100.0%] | RSS (tot,self)_curr,peak | ( 0.1572,1 ) | ( 0.1572,0 ) [MB] [total # of laps: 13]
> [pyc] |_lump_array_test.py:174      : 3.824 wall, 1.168 user + 2.656 system = 3.920 CPU [sec] [99.9%] | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 390.7181,3 ) [MB] [total # of laps: 151]
> [pyc] |_l_created_array_test.py:177 : 4.169 wall, 0.328 user + 3.840 system = 3.536 CPU [sec] [98.0%] | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 390.7198,7 ) [MB] [total # of laps: 151]
> [pyc] |_array_finalize_test_disk_array[]array_test.py:156 : 0.081 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_init_lasts_array_wake[]array_test.py:125 : 0.084 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_create_wake_rollbackfacts_array_wake[]array_test.py:182 : 0.083 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_bind_finalizeLazrLazr_array_wake[]array_test.py:168 : 0.081 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_call_init_test_disk_array[]array_test.py:167 : 0.081 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_finalize_test_array_wake[]array_test.py:127 : 0.083 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_do_rollbackLazr_array_wake[]array_test.py:130 : 0.081 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_del_init_test_disk_array[]array_test.py:163 : 0.081 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 381.61572,1 ) | ( 0.21 0.0 ) [MB] [total # of laps: 151]
> [pyc] |_dump[]array_test.py:121 : 1.082 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 0.11572,1 ) | ( 0.21 0.0 ) [MB]
> [pyc] |_do_step[]array_test.py:128 : 1.081 wall, 0.000 user + 0.080 system = 0.000 CPU [sec] ( 0.0%) | RSS (tot,self)_curr,peak | ( 0.1572,1 ) | ( 0.21 0.0 ) [MB]
    
```

2018-04-13 09:15:00 -- Amp/TMemory/odash/Continuous/build-TMemory/hst_output/nested_report.out

```

> [pyc] main[AUTO_TENER_FOR_NESTED_TEST][10][10]nested_test.py:104 : 0.660 wall, 2.410 user + 0.230 system = 2.630 CPU [sec] [30.4%] | RSS (tot,self)_curr,peak | ( 5.2110,7 ) | ( 5.2110,7 ) [MB]
> [pyc] |_nested_func_init[15]nested_test.py:169 : 0.453 wall, 0.400 user + 0.040 system = 0.400 CPU [sec] [97.2%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.4 0.4 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci13[]nested_test.py:164 : 0.441 wall, 0.350 user + 0.090 system = 0.430 CPU [sec] [97.6%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.4 0.4 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci14[]nested_test.py:164 : 0.283 wall, 0.250 user + 0.040 system = 0.290 CPU [sec] [102.3%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.2 0.2 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci13[]nested_test.py:164 : 0.174 wall, 0.168 user + 0.010 system = 0.170 CPU [sec] [97.7%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.1 0.1 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci13[]nested_test.py:164 : 0.186 wall, 0.180 user + 0.000 system = 0.180 CPU [sec] [94.7%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.1 0.1 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci11[]nested_test.py:164 : 0.855 wall, 0.850 user + 0.000 system = 0.850 CPU [sec] [92.7%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci18[]nested_test.py:164 : 0.830 wall, 0.820 user + 0.000 system = 0.800 CPU [sec] [101.3%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci9[]nested_test.py:164 : 0.823 wall, 0.820 user + 0.000 system = 0.820 CPU [sec] [86.9%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci18[]nested_test.py:164 : 0.814 wall, 0.820 user + 0.000 system = 0.820 CPU [sec] [147.3%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci18[]nested_test.py:164 : 0.800 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [128.2%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci10[]nested_test.py:164 : 0.804 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [249.2%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci6[]nested_test.py:164 : 0.807 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [207.7%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci6[]nested_test.py:164 : 0.814 wall, 0.810 user + 0.000 system = 0.820 CPU [sec] [145.3%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci7[]nested_test.py:164 : 0.803 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [159.5%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci6[]nested_test.py:164 : 0.801 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [201.7%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci8[]nested_test.py:164 : 0.804 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [243.3%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci15[]nested_test.py:164 : 0.801 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [851.2%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci15[]nested_test.py:164 : 0.823 wall, 0.820 user + 0.000 system = 0.820 CPU [sec] [87.9%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci18[]nested_test.py:164 : 0.814 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [72.8%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci18[]nested_test.py:164 : 0.804 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [257.3%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci18[]nested_test.py:164 : 0.807 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [126.2%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
> [pyc] |_fibonacci10[]nested_test.py:164 : 0.803 wall, 0.810 user + 0.000 system = 0.810 CPU [sec] [229.2%] | RSS (tot,self)_curr,peak | ( 0.4 0.4 ) | ( 0.0 0.0 ) [MB] [total # of laps: 2]
    
```

Figure 2: ASCII output of timing + memory in CDash

```
#ifndef GEANT4_USE_TIMEMORY // compiled with TiMemory

#include <timemory/timemory.hpp>

typedef tim::auto_timer G4AutoTimer;

inline void InitializeTiMemory()
{
    tim::manager* instance = tim::manager::instance();
    instance->enable(true);
}

#else // not compiled with TiMemory

#define TIMEMORY_AUTO_TIMER(str)
#define TIMEMORY_AUTO_TIMER_OBJ(str) {}
#define TIMEMORY_BASIC_AUTO_TIMER(str)
#define TIMEMORY_BASIC_AUTO_TIMER_OBJ(str) {}
#define TIMEMORY_DEBUG_BASIC_AUTO_TIMER(str)
#define TIMEMORY_DEBUG_AUTO_TIMER(str)

inline void InitializeTiMemory() { }

#endif
```

- Insert “auto-timer”

```
void G4RunManager::BeamOn(G4int, const char*, G4int)
{
    TIMEMORY_AUTO_TIMER("");
    // ...
}
```

- Result

```
> [exe] total execution time
      : 25.622 wall, 163.840 user + 0.340 system
      = 164.180 CPU [sec] (640.8%)
      : RSS {tot,self}_{curr,peak}
      : (94.9|94.9) | (71.5|71.5) [MB] (x1 laps)
> [cxx] |_BeamOn@'G4RunManager.cc':327
      : 25.591 wall, 163.820 user + 0.340 system
      = 164.160 CPU [sec] (641.5%)
      : RSS {tot,self}_{curr,peak}
      : (94.9|94.9) | (58.2|58.2) [MB] (x2 laps)
> [cxx] |_RunInitialization@'G4RunManagerKernel.cc':666
      : 1.634 wall, 1.600 user + 0.030 system
      = 1.630 CPU [sec] ( 99.7%)
      : RSS {tot,self}_{curr,peak}
      : (92.6|92.6) | (22.5|22.5) [MB] (x2 laps)
```

```
|| |0> [pyc] main@'toast_ground_sim_simple.py'  
: 41.10 wall, 69.15 user + 4.69 system = 73.84 CPU [sec] (179.6%)
```

- Timing fields:
 - Real time *e.g.*, wall clock time
 - User time *e.g.*, time spent executing user code
 - System time *e.g.*, time spent executing system code
 - CPU time *e.g.*, user + system time
 - % CPU utilization *e.g.*, $\text{cpu} / \text{real} * 100$
- Total time from all the “laps”


```
| | 10> [pyc] main@'toast_ground_sim_simple.py'  
      : RSS {tot,self}_{curr,peak} : (114.5|223.7) | (107.4|215.6) [MB]
```

- Resident set size (RSS) fields: total-current, total-peak, self-current, self-peak
 - Total fields are measured when the timer was stopped
 - Self RSS fields are difference between timer start and timer stop
 - Current RSS is the measurement of the allocated pages in memory
 - Peak RSS is the peak RSS usage for the process by the OS
- Temporary memory allocation can be calculated via *peak* – *current*
 - This calculation is valid when the timing scope is setting the “high-water mark” of memory allocation – but this is where the temporary memory usage is of primary interest
 - Relevant example: application involving matrix calculations is throwing runtime OOM/bad_alloc error due to intermediate solution matrices within the matrix library.

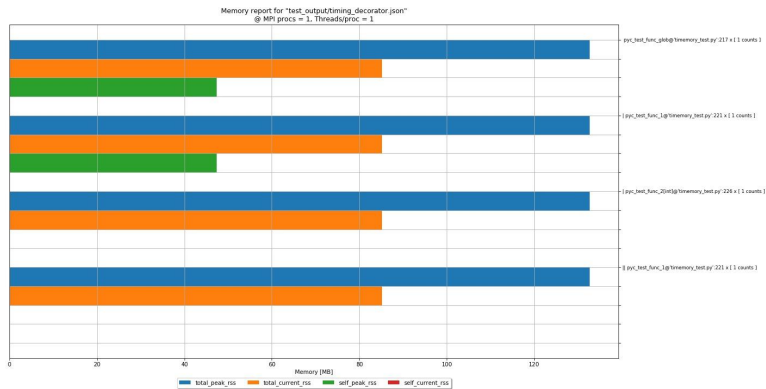


Figure 4: Memory Plot from TiMemory unit test

- **TiMemory** depends on two header-only libraries (automatically included):
 - **PyBind11** – creates Python bindings
 - **cereal** – produces JSON serialization
- Installation from Python package managers includes the C++ development files
- Install from source: github.com/jrmadsen/TiMemory
- Install from PyPi: `pip install -v timemory`
- Install from Anaconda: `conda install -c jrmadsen timemory`
- Documentation: jrmadsen.github.io/TiMemory
- System requirements
 - Operating system: Windows 10, Linux, macOS
 - Compilers: GCC 4.9+, Clang 4.0+, AppleClang, Intel, MSVC 14, MSVC 15
 - Packages: CMake ($\geq 2.8.12$)
 - Python: 2.7, ≥ 3.4

- All `rusage` fields
- GPU timers
 - CUDA: `cudaEvent_t` and `cudaEventRecord`
 - OpenCL: `clGetEventProfilingInfo`
- PAPI – hardware counters (Linux only)
 - PAPI_TOT_INS – total instructions issued
 - PAPI_VEC_DP – double precision vector/SIMD instructions
 - PAPI_L1_DCM – Level 1 data cache misses
 - PAPI_FP_OPS – floating-point operations executed
 - ... etc.