




$$\frac{\partial}{\partial a} \ln J_{a, \sigma^2}(\xi_1) = \frac{(\xi_1 - a)}{\sigma^2} f_{a, \sigma^2}(\xi_1)$$

$$\int T(x) \cdot \frac{\partial}{\partial \theta} f(x, \theta) dx = M \left(T(\xi) \cdot \frac{\partial}{\partial \theta} \ln L(\theta) \right)$$

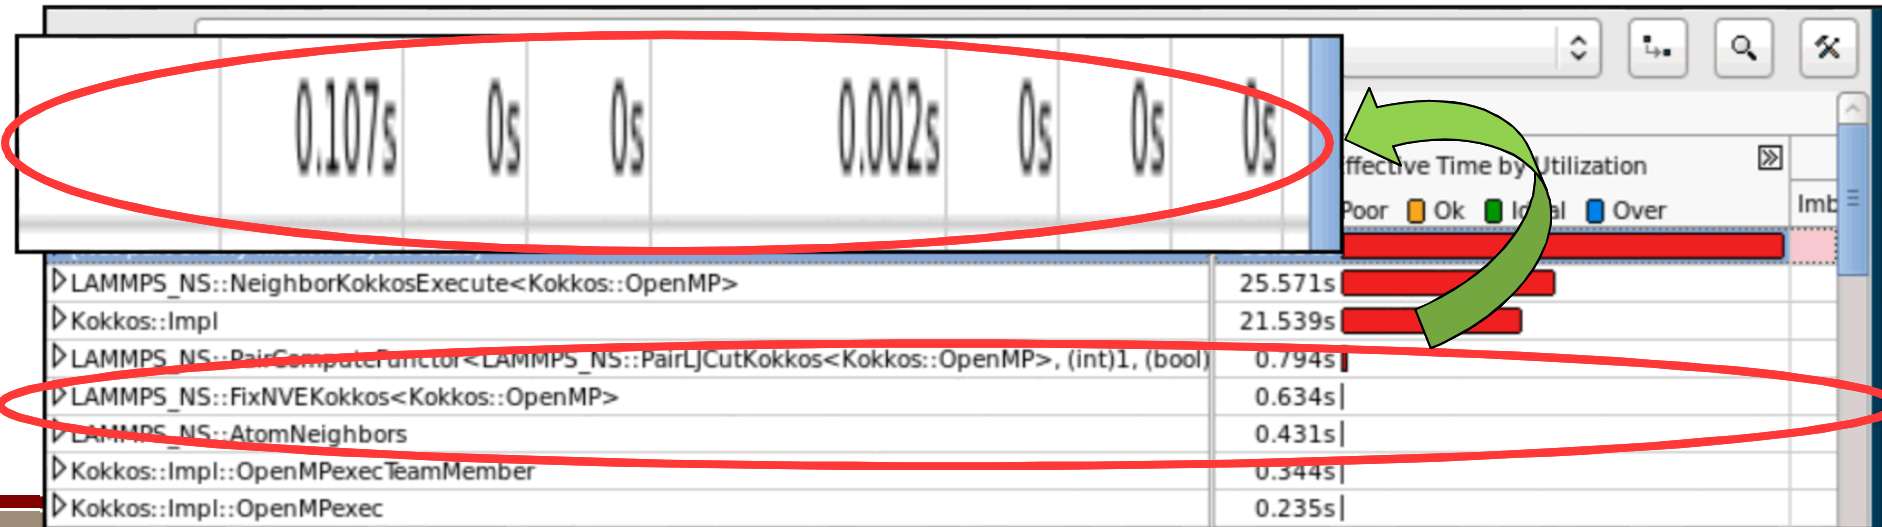


Profiling Kokkos Applications

S.D. Hammond (sdhammo@sandia.gov), the Kokkos Team and the LDMS Team
Center for Computing Research, Sandia National Laboratories/NM

Background

- Greater adoption of Kokkos within Sandia and ECP has driven stronger focus on performance analysis capabilities
- C++ abstractions of Kokkos, particularly with OpenMP challenge existing tools
- Desire for continuous monitoring of overnight testing/builds



History



- Lightweight performance instrumentation using dynamic tool loading was explored in a Sandia LDRD during 2012
 - Mechanism to instrument loops in preparation for vectorization analysis during Knights Landing exploration
- Initial exploration for Kokkos parallel patterns in 2014
- Profiling of parallel region and data allocations/copying requires **no code instrumentation** based on feedback from users
 - Want a single binary to be able to be profiled without recompilation

Kokkos Profiling Overview

User Application:

```
void mykernel() {  
    ..  
  
    Kokkos::parallel_for (N, KOKKOS_LAMBDA (const int i) {  
        printf ("Hello from i = %i\n", i);  
    });  
  
}
```


Profiling Tool:

```
START_KERNEL_PROFILE  
(e.g. START_TIMER)
```

```
END_KERNEL_PROFILER  
(e.g. END_TIMER)
```

Kokkos Runtime:

```
Kokkos::parallel_for (...) {  
    CALL_PROFILE_HOOK_START  
  
    PARALLEL BODY  
  
    CALL_PROFILE_HOOK_END  
}
```



Kokkos Profiling Overview

User Application:

```
void mykernel() {  
  ..  
  
  Kokkos::parallel_for (N, KOKKOS_LAMBDA (const int i) {  
    printf ("Hello from i = %i\n", i);  
  });  
  
}
```

Profiling Tool:

```
START_KERNEL_PROFILER  
(e.g. START_TIMER)
```

```
END_KERNEL_PROFILER  
(e.g. END_TIMER)
```

Kokkos Runtime:

```
Kokkos::parallel_for (...) {  
  CALL_PROFILE_HOOK_START  
  
  PARALLEL BODY  
  
  CALL_PROFILE_HOOK_END  
}
```

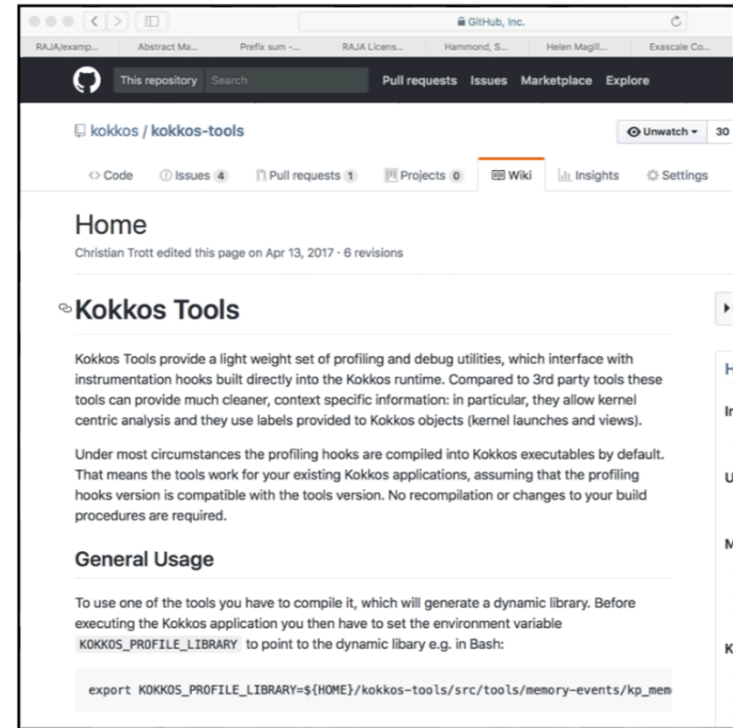
LDMS

LDMS and Performance Database in the next Talk

Updates in the Profiling Interface



- Added hooks to track allocations performed in Kokkos
 - Typically the largest consumers of memory in an application/HWM
 - Want to track NVLINK operations
- Unstructured profiling sections/regions
 - Used in Trilinos to provide better insights into solver performance
- Documentation



Example Use Cases (Memory)



#	Time	Ptr	Size	MemSpace	Op	Name
0.	002958	0x7f5c9dd030a0	8998912		Host Allocate	Vector
0.	004280	0x7f5c9d46d0a0	8998912		Host Allocate	Vector
0.	005055	0x7f5c9cbd70a0	8998912		Host Allocate	Vector
2.	008489	0x2c60470	4499464		Host Allocate	MatrixRowOffsets
2.	009254	0x7f5c40b0f0a0	119164000		Host Allocate	MatrixLocalIndicies
2.	017489	0x7f5c327c50a0	238328000		Host Allocate	MatrixValues
4.	638244	0x7f5c7f76a0a0	8998912		Host Allocate	no-label
4.	832563	0x373b0d0	562440		Host Allocate	MatrixRowOffsets
4.	832609	0x7f5c7c72e0a0	14609056		Host Allocate	MatrixLocalIndicies
4.	833970	0x7f5c6a4220a0	29218112		Host Allocate	MatrixValues
5.	154044	0x3a12ed0	1124864		Host Allocate	Vector
5.	154236	0x3b25970	1124864		Host Allocate	Vector
5.	154358	0x3c38410	8998912		Host Allocate	Vector
5.	155568	0x44cd4b0	1124864		Host Allocate	no-label
5.	178790	0x4fa7bf0	70312		Host Allocate	MatrixRowOffsets
5.	178849	0x4fba010	1755904		Host Allocate	MatrixLocalIndicies
5.	179045	0x5166bb0	3511808		Host Allocate	MatrixValues
5.	217818	0x4646f40	140608		Host Allocate	Vector
5.	218140	0x4669520	140608		Host Allocate	Vector
5.	218422	0x468bb00	1124864		Host Allocate	Vector
5.	220563	0x479e5a0	140608		Host Allocate	no-label
5.	223459	0x4903620	8792		Host Allocate	MatrixRowOffsets
5.	223489	0x49069f0	202616		Host Allocate	MatrixLocalIndicies

HPCG Kokkos Data Allocation Profile

Example Use Cases (Kernel Time)



Reference: ComputeSPMV (Region)	0.83942	205	0.00409	7.219	2.422
Main: Validation Testing Phase (Region)	0.73603	1	0.73603	6.330	2.123
Main: Reference SpMV+MG (Region)	0.31914	1	0.31914	2.744	0.921
Optimized: ComputeDotProduct (Region)	0.09877	474	0.00021	0.849	0.285
Optimized: ComputeWAXPBY (Region)	0.08671	468	0.00019	0.746	0.250
Reference: ComputeDotProduct (Region)	0.04568	151	0.00030	0.393	0.132
Main: Reference CG::OptimizeProblem (Region)	0.03189	1	0.03189	0.274	0.092
Reference: ComputeWAXPBY (Region)	0.03050	150	0.00020	0.262	0.088
Main: Clean up (Region)	0.01020	1	0.01020	0.088	0.029
Main: Report Results (Region)	0.00265	1	0.00265	0.023	0.008
Main: TestNorms (Region)	0.00000	1	0.00000	0.000	0.000

N11KokkosGraph19GraphColoringHandleIKiS1_S1_N6Kokkos60penMPENS2_9HostSpaceES4_E16ReduceMaxFunctorE (ParRed)	0.00013	4	0.00003	0.001	0.000
KokkosKernels::Impl::StridedCopy (ParFor)	0.00002	4	0.00000	0.000	0.000

Summary:

Total Execution Time (incl. Kokkos + Non-Kokkos): 34.66302 seconds
 Total Time in Kokkos kernels: 11.62839 seconds
 -> Time outside Kokkos kernels: 23.03464 seconds
 -> Percentage in Kokkos kernels: 33.55 %
 Total Calls to Kokkos Kernels: 31732

Current Work



- Want to expand Kokkos profiling to be useful for continuous integration benchmarking
 - Feedback from ATDM development teams that this is crucial for cross platform perspective
 - Intense focus on parallel kernel performance (natural alignment with Kokkos profiling hooks)
- Add historical analysis/trends capability to inform development teams
- Working with LDMS team for broader performance characterization
 - KokkosP provides strong application connectivity

Information and Tools



- Available on the Kokkos Github website
 - <http://github.com/kokkos>
- Integration into TAU is already complete, HPCToolkit on-going
- Close connection with Score-P and ExaPAPI ECP projects
 - Will soon be able to capture Kokkos events in Score-P tracing operations and view in Vampir
- Tools working on KNL, POWER8/9, ARM, Sky Lake and AMD systems

Kokkos Profiling Overview

User Application:

```
void mykernel() {  
    ..  
    Kokkos::parallel_for (N, KOKKOS_LAMBDA (const int i) {  
        printf ("Hello from i = %i\n", i);  
    });  
}
```

Profiling Tool:

```
START_KERNEL_PROFILER  
(e.g. START_TIMER)
```

```
END_KERNEL_PROFILER  
(e.g. END_TIMER)
```

Kokkos Runtime:

```
Kokkos::parallel_for (...) {  
    CALL_PROFILE_HOOK_START  
  
    PARALLEL BODY  
  
    CALL_PROFILE_HOOK_END  
}
```

LDMS

LDMS and Performance Database in the next Talk

THE UNIVERSITY OF CHICAGO
DIVISION OF THE PHYSICAL SCIENCES
DEPARTMENT OF CHEMISTRY
5780 SOUTH CAMPUS DRIVE
CHICAGO, ILLINOIS 60637
TEL: 773-936-3700
WWW.CHEM.UCHICAGO.EDU