MuCoSim: Analysis of TeaLeaf

Inhalt

Project Description

Customer Info

- Name: Thomas Gruber
- E-Mail: thomas.gruber@fau.de

Application Info

- Code: TeaLeaf
- URL: https://github.com/UK-MAC/TeaLeaf

Tealeaf is a miniapp solving linear heat conduction equation on a spatially decomposed regularly grid using a 5 point stencil with implicit solvers. The benchmark was part of the Mantevo Benchmark Suite that contains selected benchmarks of the National Labratories in the USA. The TeaLeaf mini-app itself was created and is maintained by researchers from the UK (https://ieeexplore.ieee.org/document/8049027)

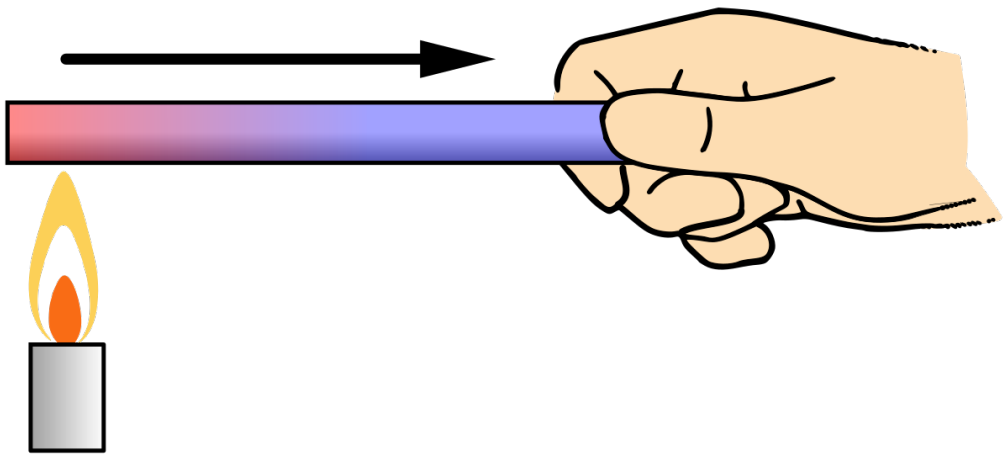

Symbolic picture of linear heat conduction. (source, CC license)

But the heat conduction calculations are only a placeholder. The idea of the TeaLeaf mini-app is to explore the design space of new, scalable iterative sparse linear solvers which can target next generation architectures.

Testsystem

- Host/Clustername: icx36
- Cluster Info URL: https://hpc.fau.de/systems-services/systems-documentation-instructions/clusters/test-cluster/
- CPU type: 2x Intel Xeon Platinum 8360Y @ 2.4 GHz
- Memory capacity: 256 GB
- Number of cores per node: 72/144
- Interconnect: None

Software Environment

Compiler:

- Compiler: Intel Fortran 19.0.5.281
- MPI: Intel MPI Library 2019 Update 5
- Operating System: Ubuntu 20.04.3 LTS
- Addition libraries:
  - LIKWID 5.2.0

How to build software

```
$ module load intel64/19.0up05
$ module load likwid/5.2-dev
$ git clone --recursive git@github.com:UK-MAC/TeaLeaf.git
$ cd TeaLeaf
$ cd TeaLeaf_ref
$ make COMPILER=INTEL MPI_COMPILER=mpiifort C_MPI_COMPILER=mpiicc
```

Although selecting a valid compiler `INTEL`, the OpenMP flags need be added manually, otherwise there is no `-qopenmp` on the compile lines. This can be fixed by changing `OOMP_INTEL` to `OMP_INTEL`.

Note: This logbook does not manipulate the optimization flags for the compilers.

Testcase description

The testcase can be defined by the file `tea.in`. The default content is:

```
*tea
state 1 density=100.0 energy=0.0001
state 2 density=0.1 energy=25.0 geometry=rectangle xmin=0.0 xmax=1.0 ymin=1.0 ymax=2.0
state 3 density=0.1 energy=0.1 geometry=rectangle xmin=1.0 xmax=6.0 ymin=1.0 ymax=2.0
state 4 density=0.1 energy=0.1 geometry=rectangle xmin=5.0 xmax=6.0 ymin=1.0 ymax=8.0
state 5 density=0.1 energy=0.1 geometry=rectangle xmin=5.0 xmax=10.0 ymin=7.0 ymax=8.0

x_cells=1000
y_cells=1000

xmin=0.0
ymin=0.0
xmax=10.0
ymax=10.0

initial_timestep=0.004
end_step=10

tl_max_iters=10000
test_problem 4
tl_use_ppcg
tl_eps=1.0e-15
tl_preconditioner_type=none
halo_depth=2

*endtea
```

The default test case uses `tl_use_ppcg` (Polynomially Preconditioned Conjugate Gradient method), thus the CG kernels internally. This can be changed to other type of kernels with:

- `tl_use_chebyshev`: Chebyshev method to solve the linear system.
- `tl_use_cg`: Conjugate Gradient method to solve the linear system.
- `tl_use_jacobi`: Jacobi method to solve the linear system. Note that this a very slowly converging method compared to other options. This is the default method is no method is explicitly selected.

We keep the default `tl_use_ppcg` for this analysis

How to run software

```
$ module load intel64/19.0up05
$ module load likwid/5.2-dev
$ cd TeaLeaf
$ cd TeaLeaf_ref
$ likwid-mpirun -np X -t Y ./tea_leaf
```

Task1: Scaling runs

For scaling runs, we find a runtime/performance number in the output of TeaLeaf (file `tea.out`). The last line always look like this:

```
 Wall clock    2.03252601623535
```

We use this value to plot the runtime of 10 iterations (see `end-step` in testcase definition) for 1 to 72 OpenMP threads inside a single MPI process (`for i in {1..72}; do likwid-mpirun -np 1 -t $i ./tea_leaf; tail -n 1 tea.out | awk '{print $3}'; done`).

We plot it with gnuplot:

```
set terminal png
set output 'openmp-scaling.png'
set title 'Weak scaling of TeaLeaf on Intel IcelakeSP 8360Y, Intel 19.0.5, default flags, ppcg method'
set xlabel '#Threads'
set xrange [1:72]
set ylabel 'Runtime [s]'
plot 'openmp-scaling.dat' w linespoints title 'Runtime'
```

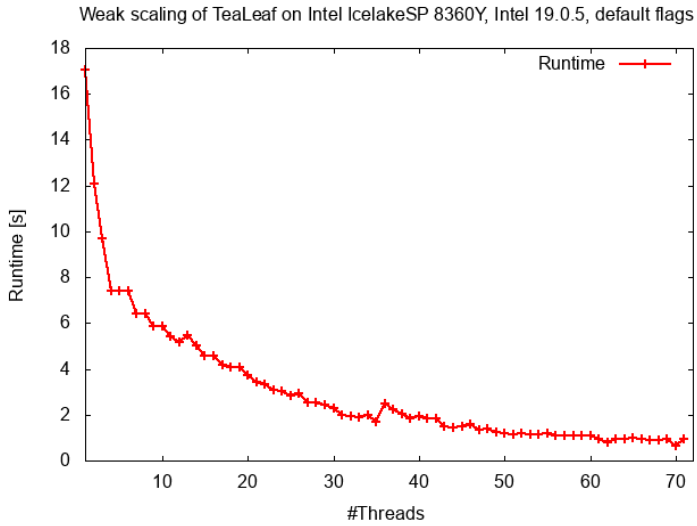

Since the system is not booted with SNC, each CPU socket (32 physical hardware threads) also forms a NUMA domain. So the data for 1 to 32 threads shows the scaling inside a NUMA domain. The data shows that the code is not limited by the memory bandwidth. The maximum speedup is 54 with 71 threads.

Task2: Whole application measurements

In order to get a first impression about the runtime behavior of the code, we use a basic set of LIKWID performance groups:

- DATA: Load-store-ratio
- FLOPS_DP/SP: Flops rate
- L2: L1 <-> L2 memory traffic
- L3: L2 <-> L3 memory traffic
- MEM: Memory traffic

Running: `likwid-mpirun -np X -t Y -g GROUP ./tea_leaf`

```
$ likwid-mpirun -np 1 -t 72 -g DATA ./tea_leaf
+---------------------------+-------------+--------------+-----------+-----------+--------
---+-----------+-----------+
|           Metric          |     Sum     |      Min     |    Max    |    Avg    |  %ile
25  |  %ile 50  |  %ile 75  |
+---------------------------+-------------+--------------+-----------+-----------+--------
---+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |    182.3832 |       2.5331 |    2.5331 |    2.5331 |
2.5331 |    2.5331 |    2.5331 |
| Runtime unhalted [s] STAT |    102.2424 | 3.318629e-05 |    1.8729 |    1.4200 |
1.4307 |    1.4331 |    1.4381 |
|      Clock [MHz] STAT     | 220207.5120 |     828.5050 | 3182.9980 | 3058.4377 |
3086.4877 | 3089.4152 | 3090.6367 |
|          CPI STAT         |     57.7787 |       0.6431 |    6.6610 |    0.8025 |
0.7086 |    0.7151 |    0.7288 |
|  Load to store ratio STAT |    349.1250 |       2.4059 |    5.5516 |    4.8490 |
4.7978 |    4.8452 |    4.8854 |
+---------------------------+-------------+--------------+-----------+-----------+--------
---+-----------+-----------+
```

```
$ likwid-mpirun -np 1 -t 72 -g FLOPS_DP ./tea_leaf
+---------------------------+-------------+--------------+-----------+-----------+--------
---+-----------+-----------+
|           Metric          |     Sum     |      Min     |    Max    |    Avg    |  %ile
25  |  %ile 50  |  %ile 75  |
+---------------------------+-------------+--------------+-----------+-----------+--------
---+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |    183.2400 |       2.5450 |    2.5450 |    2.5450 |
2.5450 |    2.5450 |    2.5450 |
| Runtime unhalted [s] STAT |     98.9696 | 4.598167e-05 |    1.8340 |    1.3746 |
1.3835 |    1.3855 |    1.3903 |
|      Clock [MHz] STAT     | 220489.3976 |    1005.4431 | 3175.0080 | 3062.3527 |
3089.4229 | 3090.1907 | 3090.6307 |
|          CPI STAT         |     56.3625 |       0.6312 |    6.2647 |    0.7828 |
0.6945 |    0.6994 |    0.7132 |
|     DP [MFLOP/s] STAT     |  74459.0943 | 5.108131e-06 | 1054.4745 | 1034.1541 |
1045.6021 | 1049.6541 | 1049.6547 |
|   AVX DP [MFLOP/s] STAT   |           0 |            0 |         0 |         0 |
0 |         0 |         0 |
|  AVX512 DP [MFLOP/s] STAT |           0 |            0 |         0 |         0 |
0 |         0 |         0 |
|   Packed [MUOPS/s] STAT   |  36566.0353 |            0 |  517.7326 |  507.8616 |
513.8562 |  515.3623 |  515.3623 |
|   Scalar [MUOPS/s] STAT   |   1327.0244 | 5.108131e-06 |   19.0093 |   18.4309 |
18.8398 |   18.9295 |   18.9298 |
|  Vectorization ratio STAT |   6851.3668 |            0 |   96.6357 |   95.1579 |
96.4570 |   96.4571 |   96.4585 |
+---------------------------+-------------+--------------+-----------+-----------+--------
---+-----------+-----------+

$ likwid-mpirun -np 1 -t 72 -g L2 ./tea_leaf
+-------------------------------------+-------------+--------------+-----------+----------
-+-----------+-----------+-----------+
|               Metric                |     Sum     |      Min     |    Max    |    Avg
|  %ile 25  |  %ile 50  |  %ile 75  |
+-------------------------------------+-------------+--------------+-----------+----------
-+-----------+-----------+-----------+
|       Runtime (RDTSC) [s] STAT      |    182.1240 |       2.5295 |    2.5295 |    2.5295
|    2.5295 |    2.5295 |    2.5295 |
|      Runtime unhalted [s] STAT      |    101.0706 | 2.927316e-05 |    1.8539 |    1.4038
|    1.4150 |    1.4172 |    1.4195 |
|           Clock [MHz] STAT          | 220296.7610 |     797.8414 | 3176.4964 | 3059.6772
| 3089.9822 | 3090.1244 | 3090.7659 |
|              CPI STAT               |     56.2838 |       0.6480 |    5.5915 |    0.7817
|    0.6929 |    0.7201 |    0.7272 |
|  L2D load bandwidth [MBytes/s] STAT | 333093.2221 |       0.0181 | 5114.7470 | 4626.2948
| 4689.5857 | 4705.2503 | 4717.5767 |
|  L2D load data volume [GBytes] STAT |    842.5502 | 4.576000e-05 |   12.9376 |   11.7021
|   11.8640 |   11.9019 |   11.9332 |
| L2D evict bandwidth [MBytes/s] STAT | 123716.7908 |       0.0030 | 2372.2682 | 1718.2888
| 1722.4463 | 1735.5784 | 1750.3171 |
| L2D evict data volume [GBytes] STAT |    312.9379 | 7.488000e-06 |    6.0006 |    4.3464
|    4.3606 |    4.3919 |    4.4311 |
|     L2 bandwidth [MBytes/s] STAT    | 463386.4874 |       0.1610 | 7367.1822 | 6435.9234
| 6512.0293 | 6540.5209 | 6555.5289 |
|     L2 data volume [GBytes] STAT    |   1172.1230 |       0.0004 |   18.6351 |   16.2795
|   16.4720 |   16.5441 |   16.5820 |
+-------------------------------------+-------------+--------------+-----------+----------
-+-----------+-----------+-----------+
```

```
$ likwid-mpirun -np 1 -t 72 -g L3 ./tea_leaf
+-----------------------------------------+-------------+-------------+-----------+-------
----+-----------+-----------+-----------+
|                Metric                   |     Sum     |      Min    |    Max    |    Avg
|  %ile 25  |  %ile 50  |  %ile 75  |
+-----------------------------------------+-------------+-------------+-----------+-------
----+-----------+-----------+-----------+
|        Runtime (RDTSC) [s] STAT         |    179.4456 |       2.4923 |    2.4923 |
2.4923 |    2.4923 |    2.4923 |    2.4923 |
|        Runtime unhalted [s] STAT        |     97.2075 | 3.300339e-05 |    1.8161 |
1.3501 |    1.3606 |    1.3625 |    1.3646 |
|            Clock [MHz] STAT             | 220257.8685 |     868.6019 | 3178.1382 |
3059.1371 | 3088.2684 | 3088.6590 | 3089.1183 |
|               CPI STAT                  |     55.6943 |       0.6687 |    6.0087 |
0.7735 |    0.6796 |    0.6954 |    0.7088 |
|    L3 load bandwidth [MBytes/s] STAT    |  25824.6451 |       0.1112 | 2117.9824 |
358.6756 |  274.3393 |  346.5871 |  376.0652 |
|    L3 load data volume [GBytes] STAT    |     64.3634 |       0.0003 |    5.2787 |
0.8939 |    0.7797 |    0.8647 |    0.9645 |
|   L3 evict bandwidth [MBytes/s] STAT    |   5293.9780 |       0.0022 | 1770.9849 |
73.5275 |   38.5921 |   41.8437 |   47.1217 |
|   L3 evict data volume [GBytes] STAT    |     13.1947 | 5.568000e-06 |    4.4139 |
0.1833 |    0.0968 |    0.1067 |    0.1219 |
| L3|MEM evict bandwidth [MBytes/s] STAT |   5292.8730 |       0.0022 | 1773.9910 |
73.5121 |   38.1772 |   41.4412 |   47.1217 |
| L3|MEM evict data volume [GBytes] STAT |     13.1916 | 5.568000e-06 |    4.4214 |
0.1832 |    0.0967 |    0.1058 |    0.1219 |
|  Dropped CLs bandwidth [MBytes/s] STAT |  20731.5238 |       0.1086 |  539.3441 |
287.9378 |  271.5942 |  303.4878 |  329.7788 |
|  Dropped CLs data volume [GBytes] STAT |     51.6701 |       0.0003 |    1.3442 |
0.7176 |    0.6769 |    0.7564 |    0.8219 |
|      L3 bandwidth [MBytes/s] STAT       |  31117.5179 |       0.1134 | 3891.9735 |
432.1877 |  347.5271 |  389.3839 |  427.9696 |
|      L3 data volume [GBytes] STAT       |     77.5552 |       0.0003 |    9.7001 |
1.0772 |    0.8701 |    0.9711 |    1.0735 |
+-----------------------------------------+-------------+-------------+-----------+-------
----+-----------+-----------+-----------+

$ likwid-mpirun -np 1 -t 72 -g MEM ./tea_leaf
+-----------------------------------------+-------------+-------------+-----------+-------
----+-----------+-----------+-----------+
|                Metric                   |     Sum     |      Min    |    Max    |    Avg
|  %ile 25  |  %ile 50  |  %ile 75  |
+-----------------------------------------+-------------+-------------+-----------+-------
----+-----------+-----------+-----------+
|        Runtime (RDTSC) [s] STAT         |    178.9056 |       2.4848 |    2.4848 |
2.4848 |    2.4848 |    2.4848 |    2.4848 |
|        Runtime unhalted [s] STAT        |     96.5925 | 2.904567e-05 |    1.7959 |
1.3416 |    1.3498 |    1.3535 |    1.3567 |
|            Clock [MHz] STAT             | 220204.6689 |     798.2143 | 3178.4431 |
3058.3982 | 3088.3971 | 3089.1634 | 3089.6046 |
|               CPI STAT                  |     53.9485 |       0.6265 |    4.4839 |
0.7493 |    0.6865 |    0.6917 |    0.7052 |
|  Memory read bandwidth [MBytes/s] STAT |   1825.9255 |            0 | 1502.9683 |
25.3601 |         0 |         0 |         0 |
|  Memory read data volume [GBytes] STAT |      4.5371 |            0 |    3.7346 |
0.0630 |         0 |         0 |         0 |
| Memory write bandwidth [MBytes/s] STAT |   2817.0067 |            0 | 2438.0808 |
39.1251 |         0 |         0 |         0 |
| Memory write data volume [GBytes] STAT |      6.9997 |            0 |    6.0581 |
0.0972 |         0 |         0 |         0 |
|    Memory bandwidth [MBytes/s] STAT     |   4642.9322 |            0 | 3941.0491 |
64.4852 |         0 |         0 |         0 |
|    Memory data volume [GBytes] STAT     |     11.5367 |            0 |    9.7927 |
0.1602 |         0 |         0 |         0 |
+-----------------------------------------+-------------+-------------+-----------+-------
----+-----------+-----------+-----------+
```

Main observations:

- One HW thread runs with reduced CPU frequency, the outputs show that it is the last HW thread (this causes also the difference in CPI)

- The load-store-ratio is mostly stable with around 4.8 but the maximum is 5.5 and the minimum 2.4. The 2.4 are caused by the last HW thread which is mostly idleing.
- With the current compile options, the code performs only SSE double-precision floating-point operations.
- Most data is served by the L2 cache and the data is mostly read once, most of it is dropped at the L3 but not reloaded from memory.

Note: Here we present only single measurements for each group but it is beneficial to run the measurements like the scaling tests and produce plots `for i in {1..72}; do likwid-mpirun -np 1 -t $i -g GROUP ./tea_leaf; done`.

Note: Until here is enough for the first talk. Put effort in the general description of the benchmark (where was it extracted, what is computed, ...) and derive the general behavior of the code through the measurements.

Task3: Runtime profile

For the runtime profile, we commonly use `gprof` due to the already available support by the compilers. In order to get the sampling data, we add the flag `-pg` to `FLAGS_INTEL` in `Makefile` and rebuild (`make clean && make`).

Afterwards, we run TeaLeaf with 2 MPI processes and 2 OpenMP threads each to cover all functions used in bigger runs: `likwid-mpirun -np 2 -t 2 ./tea_leaf`

In the end, we run `prof tea_leaf gmon.out` to get the profile:

```
Each sample counts as 0.01 seconds.
  %   cumulative   self              self     total
 time   seconds   seconds    calls   s/call   s/call  name
 94.14     32.58    32.58    22816     0.00     0.00
tea_leaf_ppcg_kernel_module_mp_tea_leaf_kernel_ppcg_inner_
  1.47     33.09     0.51      558     0.00     0.00
tea_leaf_cg_kernel_module_mp_tea_leaf_cg_calc_w_kernel_
  1.30     33.54     0.45      558     0.00     0.00
tea_leaf_cg_kernel_module_mp_tea_leaf_cg_calc_ur_kernel_
  0.75     33.80     0.26                             __intel_avx_rep_memcpy
  0.55     33.99     0.19      558     0.00     0.00
tea_leaf_cg_kernel_module_mp_tea_leaf_cg_calc_p_kernel_
  0.46     34.15     0.16      184     0.00     0.00
tea_leaf_ppcg_kernel_module_mp_tea_leaf_kernel_ppcg_init_sd_
  0.46     34.31     0.16      164     0.00     0.00
tea_leaf_ppcg_kernel_module_mp_tea_leaf_ppcg_calc_rrn_kernel_
  0.20     34.38     0.07       20     0.00     0.00
tea_leaf_common_kernel_module_mp_tea_leaf_common_init_kernel_
  0.20     34.45     0.07      164     0.00     0.00
tea_leaf_ppcg_kernel_module_mp_tea_leaf_ppcg_calc_zrnorm_kernel_
  0.14     34.50     0.05       20     0.00     0.00
tea_leaf_common_kernel_module_mp_tea_leaf_calc_residual_kernel_
```

After analysis of the code, we identify 3 main kernels in the `tl_use_ppcg` configuration:

- `tea_leaf_cg_kernel_module_mp_tea_leaf_cg_calc_w_kernel_`
- `tea_leaf_cg_kernel_module_mp_tea_leaf_cg_calc_ur_kernel_`
- `tea_leaf_cg_kernel_module_mp_tea_leaf_cg_calc_p_kernel_`

Task4: Instrument kernels with MarkerAPI

We identified the kernels in Task2 and now put MarkerAPI calls around it. But at first, we have to initialize and finalize the LIKWID MarkerAPI in a serial region.

In `tea_leaf.f90` we add the calls to the `PROGRAM tea_leaf`:

```
diff --git a/tea_leaf.f90 b/tea_leaf.f90
index 2efcd22..3f21e71 100644
--- a/tea_leaf.f90
+++ b/tea_leaf.f90
@@ -35,14 +35,19 @@
 PROGRAM tea_leaf

   USE tea_module
+  USE likwid

   IMPLICIT NONE

 !$ INTEGER :: OMP_GET_NUM_THREADS,OMP_GET_THREAD_NUM
+  call likwid_markerInit()

   CALL tea_init_comms()

 !$OMP PARALLEL
+  call likwid_markerRegisterRegion("w")
+  call likwid_markerRegisterRegion("ur")
+  call likwid_markerRegisterRegion("p")
   IF(parallel%boss)THEN
 !$  IF(OMP_GET_THREAD_NUM().EQ.0) THEN
       WRITE(*,*)
@@ -68,6 +73,7 @@ PROGRAM tea_leaf
   CALL diffuse

   ! Deallocate everything
+  call likwid_markerClose()

 END PROGRAM tea_leaf
```

With `USE likwid` we add the MarkerAPI module to our programm. The main programm opens a parallel region, we use that region to register the region names in the API for all threads to reduce the overhead in the main application. Somewhere in the beginning and end of the `PROGRAM` we add the initialization and finalization calls of the MarkerAPI.

```
diff --git a/kernels/tea_leaf_cg_kernel.f90 b/kernels/tea_leaf_cg_kernel.f90
index 86b7543..1ec8bf4 100644
--- a/kernels/tea_leaf_cg_kernel.f90
+++ b/kernels/tea_leaf_cg_kernel.f90
@@ -23,6 +23,7 @@ MODULE tea_leaf_cg_kernel_module
~
   USE definitions_module, only: tl_ppcg_active
   USE tea_leaf_common_kernel_module
+  USE likwid
~
   IMPLICIT NONE
~
@@ -139,6 +140,7 @@ SUBROUTINE tea_leaf_cg_calc_w_kernel(x_min,             &
   pw = 0.0_8
~
 !$OMP PARALLEL REDUCTION(+:pw)
+    call likwid_markerStartRegion("w")
 !$OMP DO
     DO k=y_min,y_max
        DO j=x_min,x_max
@@ -151,6 +153,7 @@ SUBROUTINE tea_leaf_cg_calc_w_kernel(x_min,             &
        ENDDO
     ENDDO
 !$OMP END DO NOWAIT
+    call likwid_markerStopRegion("w")
 !$OMP END PARALLEL
~
 END SUBROUTINE tea_leaf_cg_calc_w_kernel
@@ -232,6 +235,7 @@ SUBROUTINE tea_leaf_cg_calc_ur_kernel(x_min,             &
   rrn = 0.0_8
~
 !$OMP PARALLEL REDUCTION(+:rrn)
+  call likwid_markerStartRegion("ur")
   IF (preconditioner_type .NE. TL_PREC_NONE) THEN
~
     IF (preconditioner_type .EQ. TL_PREC_JAC_DIAG) THEN
@@ -299,6 +303,7 @@ SUBROUTINE tea_leaf_cg_calc_ur_kernel(x_min,             &
     ENDDO
 !$OMP END DO NOWAIT
   ENDIF
+  call likwid_markerStopRegion("ur")
 !$OMP END PARALLEL
~
 END SUBROUTINE tea_leaf_cg_calc_ur_kernel
@@ -325,6 +330,7 @@ SUBROUTINE tea_leaf_cg_calc_p_kernel(x_min,             &
   REAL(kind=8) :: beta
~
 !$OMP PARALLEL
+  call likwid_markerStartRegion("p")
   IF (preconditioner_type .NE. TL_PREC_NONE .or. tl_ppcg_active) THEN
 !$OMP DO
     DO k=y_min,y_max
@@ -342,6 +348,7 @@ SUBROUTINE tea_leaf_cg_calc_p_kernel(x_min,             &
     ENDDO
 !$OMP END DO NOWAIT
   ENDIF
+  call likwid_markerStopRegion("p")
 !$OMP END PARALLEL
~
 END SUBROUTINE tea_leaf_cg_calc_p_kernel
```

All three kernels of interest are in the same file `kernels/tea_leaf_cg_kernel.f90`. Similar to the other file, we have to add `USE LIKWID` to use LIKWID's functions. Each of the kernels contains a parallel region which is enriched with the start and stop calls.

In order to build the application, we have to add `-I $LIKWID_INCDIR` to the include paths and `-L $LIKWID_LIBDIR` to the library paths. The environment variables are defined by the likwid modules on the system. Finally, we add the LIKWID library for linking `-llikwid`. Since it is Fortran90 code which commonly does not provide macros, we can omit the `-DLIKWID_PERFMON` but it also does not hurt.

Task5: Measurements of the selected hot spots

Afterwards we can run in similarly to the Task2 but add the `-m` CLI switch to activate the MarkerAPI.

```
$ likwid-mpirun -s 0x0 -n 1 -t 72 -m -g FLOPS_DP ./tea_leaf
Region: w
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|           Metric           |     Sum     |    Min    |    Max    |    Avg    |  %ile 25
|  %ile 50  |  %ile 75  |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |      0.3951 |    0.0051 |    0.0060 |    0.0055 |    0.0053
|    0.0054 |    0.0056 |
| Runtime unhalted [s] STAT |      0.5629 |    0.0074 |    0.0086 |    0.0078 |    0.0076
|    0.0077 |    0.0080 |
|      Clock [MHz] STAT     | 222664.8705 | 3090.8785 | 3094.0863 | 3092.5676 | 3092.0583
| 3092.5209 | 3093.2281 |
|          CPI STAT         |     23.8137 |    0.3113 |    0.3617 |    0.3307 |    0.3220
|    0.3263 |    0.3352 |
|     DP [MFLOP/s] STAT     | 662900.3614 | 8341.7151 | 9780.3329 | 9206.9495 | 8988.6078
| 9330.9124 | 9469.5250 |
|   AVX DP [MFLOP/s] STAT   |           0 |         0 |         0 |         0 |         0
|         0 |         0 |
|  AVX512 DP [MFLOP/s] STAT |           0 |         0 |         0 |         0 |         0
|         0 |         0 |
|   Packed [MUOPS/s] STAT   | 328574.5409 | 4134.6712 | 4847.7394 | 4563.5353 | 4455.3140
| 4624.9787 | 4693.6837 |
|   Scalar [MUOPS/s] STAT   |   5751.2792 |   72.3727 |   84.8541 |   79.8789 |   77.9798
|   80.9550 |   82.1576 |
|  Vectorization ratio STAT |   7076.1392 |   98.2797 |   98.2798 |   98.2797 |   98.2797
|   98.2797 |   98.2797 |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
Region: ur
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|           Metric           |     Sum     |    Min    |    Max    |    Avg    |  %ile 25
|  %ile 50  |  %ile 75  |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |      0.2496 |    0.0033 |    0.0038 |    0.0035 |    0.0034
|    0.0035 |    0.0035 |
| Runtime unhalted [s] STAT |      0.3787 |    0.0050 |    0.0056 |    0.0053 |    0.0052
|    0.0052 |    0.0054 |
|      Clock [MHz] STAT     | 222668.4328 | 3088.7923 | 3095.3406 | 3092.6171 | 3091.7109
| 3092.5563 | 3093.5345 |
|          CPI STAT         |     29.1717 |    0.3850 |    0.4311 |    0.4052 |    0.3961
|    0.4045 |    0.4124 |
|     DP [MFLOP/s] STAT     | 484880.7811 | 6179.9097 | 7106.3397 | 6734.4553 | 6573.5645
| 6754.4185 | 6877.9142 |
|   AVX DP [MFLOP/s] STAT   |           0 |         0 |         0 |         0 |         0
|         0 |         0 |
|  AVX512 DP [MFLOP/s] STAT |           0 |         0 |         0 |         0 |         0
|         0 |         0 |
|   Packed [MUOPS/s] STAT   | 232426.1095 | 2962.3241 | 3406.4016 | 3228.1404 | 3151.0175
| 3237.7092 | 3296.9065 |
|   Scalar [MUOPS/s] STAT   |  20028.5627 |  255.2616 |  293.5366 |  278.1745 |  271.5296
|       279 |  284.1012 |
|  Vectorization ratio STAT |   6628.7832 |   92.0664 |   92.0667 |   92.0664 |   92.0664
|   92.0664 |   92.0664 |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
Region: p
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|           Metric           |     Sum     |    Min    |    Max    |    Avg    |  %ile 25
|  %ile 50  |  %ile 75  |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |      0.1319 |    0.0016 |    0.0028 |    0.0018 |    0.0017
|    0.0017 |    0.0017 |
| Runtime unhalted [s] STAT |      0.2260 |    0.0028 |    0.0044 |    0.0031 |    0.0029
|    0.0030 |    0.0031 |
|      Clock [MHz] STAT     | 222662.0068 | 3087.4521 | 3097.5647 | 3092.5279 | 3091.0457
| 3092.3075 | 3093.6306 |
|          CPI STAT         |     41.3623 |    0.5172 |    0.8132 |    0.5745 |    0.5334
```

```
|    0.5484 |    0.5579 |
|     DP [MFLOP/s] STAT     | 312131.1050 | 2740.8344 | 4834.9551 | 4335.1542 | 4459.1813
| 4524.2979 | 4597.2635 |
|   AVX DP [MFLOP/s] STAT   |           0 |         0 |         0 |         0 |         0
|         0 |         0 |
|  AVX512 DP [MFLOP/s] STAT |           0 |         0 |         0 |         0 |         0
|         0 |         0 |
|   Packed [MUOPS/s] STAT   | 149758.2348 | 1315.0318 | 2319.7753 | 2079.9755 | 2139.4819
| 2170.7243 | 2205.7412 |
|   Scalar [MUOPS/s] STAT   |  12614.6354 |  110.7707 |  195.4045 |  175.2033 |  180.2176
|  182.8493 |  185.7812 |
|  Vectorization ratio STAT |   6640.6376 |   92.2310 |   92.2317 |   92.2311 |   92.2310
|   92.2310 |   92.2310 |
+---------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
```

So the first region `w` has the highest FLOPS rate and vectorization ratio. All regions have a load imbalance visible by comparing the min/max values. But we also see again that **no** AVX and **no** AVX512 vectorization is used.

```
$ likwid-mpirun -s 0x0 -n 1 -t 72 -m -g FLOPS_DP ./tea_leaf
Region: w
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|           Metric           |     Sum     |    Min    |    Max    |    Avg    |  %ile 25
|  %ile 50  |  %ile 75  |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |      0.3944 |    0.0052 |    0.0060 |    0.0055 |    0.0053
|    0.0054 |    0.0056 |
| Runtime unhalted [s] STAT |      0.5596 |    0.0074 |    0.0084 |    0.0078 |    0.0075
|    0.0077 |    0.0079 |
|      Clock [MHz] STAT     | 222679.0232 | 3090.9384 | 3094.9578 | 3092.7642 | 3092.3092
| 3092.6911 | 3093.2651 |
|          CPI STAT         |     23.7030 |    0.3092 |    0.3581 |    0.3292 |    0.3202
|    0.3273 |    0.3356 |
|  Load to store ratio STAT |    656.3725 |    9.0900 |    9.1379 |    9.1163 |    9.1086
|    9.1164 |    9.1238 |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
Region: ur
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|           Metric           |     Sum     |    Min    |    Max    |    Avg    |  %ile 25
|  %ile 50  |  %ile 75  |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |      0.2487 |    0.0033 |    0.0037 |    0.0035 |    0.0034
|    0.0035 |    0.0035 |
| Runtime unhalted [s] STAT |      0.3718 |    0.0049 |    0.0055 |    0.0052 |    0.0051
|    0.0052 |    0.0052 |
|      Clock [MHz] STAT     | 222680.9948 | 3090.4554 | 3095.5358 | 3092.7916 | 3092.2490
| 3092.6694 | 3093.3182 |
|          CPI STAT         |     28.6990 |    0.3792 |    0.4281 |    0.3986 |    0.3892
|    0.3997 |    0.4048 |
|  Load to store ratio STAT |    152.7532 |    2.1097 |    2.1333 |    2.1216 |    2.1173
|    2.1214 |    2.1256 |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
Region: p
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|           Metric           |     Sum     |    Min    |    Max    |    Avg    |  %ile 25
|  %ile 50  |  %ile 75  |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
|  Runtime (RDTSC) [s] STAT |      0.1342 |    0.0016 |    0.0028 |    0.0019 |    0.0017
|    0.0017 |    0.0018 |
| Runtime unhalted [s] STAT |      0.2250 |    0.0028 |    0.0043 |    0.0031 |    0.0029
|    0.0030 |    0.0030 |
|      Clock [MHz] STAT     | 222638.2538 | 3088.0150 | 3096.7584 | 3092.1980 | 3091.0701
| 3091.9533 | 3093.2619 |
|          CPI STAT         |     41.3154 |    0.5148 |    0.8012 |    0.5738 |    0.5329
|    0.5429 |    0.5587 |
|  Load to store ratio STAT |    154.1515 |    2.1179 |    2.1638 |    2.1410 |    2.1326
|    2.1407 |    2.1488 |
+----------------------------+-------------+-----------+-----------+-----------+-----------
+-----------+-----------+
```

The w region loads most data for a single store with a load-store-ratio 9:1.

```
$ likwid-mpirun -s 0x0 -n 1 -t 72 -m -g L2 ./tea_leaf
Region: w
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+
|               Metric                |      Sum    |     Min    |     Max    |     Avg
|   %ile 25  |   %ile 50  |   %ile 75  |
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+
|       Runtime (RDTSC) [s] STAT      |       0.3911 |     0.0051 |     0.0059 |
0.0054 |     0.0053 |     0.0054 |     0.0056 |
|      Runtime unhalted [s] STAT      |       0.5585 |     0.0074 |     0.0084 |
0.0078 |     0.0076 |     0.0077 |     0.0079 |
|           Clock [MHz] STAT          |  222674.3524 |  3091.0259 |  3094.3102 |
3092.6993 |  3092.1627 |  3092.7084 |  3093.2592 |
|              CPI STAT               |      23.6515 |     0.3104 |     0.3522 |
0.3285 |     0.3206 |     0.3274 |     0.3350 |
|  L2D load bandwidth [MBytes/s] STAT | 2.187759e+06 | 28168.1829 | 32183.6812 |
30385.5477 | 29696.2059 | 30503.9852 | 31089.0747 |
|  L2D load data volume [GBytes] STAT |      11.8777 |     0.1645 |     0.1664 |
0.1650 |     0.1647 |     0.1648 |     0.1649 |
| L2D evict bandwidth [MBytes/s] STAT |  474795.8580 |  6106.4541 |  6970.5408 |
6594.3869 |  6454.5156 |  6624.8844 |  6731.4596 |
| L2D evict data volume [GBytes] STAT |       2.5779 |     0.0356 |     0.0363 |
0.0358 |     0.0357 |     0.0358 |     0.0358 |
|     L2 bandwidth [MBytes/s] STAT    | 2.823101e+06 | 36248.2886 | 41555.5146 |
39209.7337 | 38373.9808 | 39366.7075 | 40129.6413 |
|     L2 data volume [GBytes] STAT    |      15.3269 |     0.2118 |     0.2152 |
0.2129 |     0.2123 |     0.2128 |     0.2131 |
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+

Region: ur
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+
|               Metric                |      Sum    |     Min    |     Max    |     Avg
|   %ile 25  |   %ile 50  |   %ile 75  |
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+
|       Runtime (RDTSC) [s] STAT      |       0.2470 |     0.0033 |     0.0036 |
0.0034 |     0.0034 |     0.0034 |     0.0035 |
|      Runtime unhalted [s] STAT      |       0.3741 |     0.0050 |     0.0055 |
0.0052 |     0.0051 |     0.0052 |     0.0053 |
|           Clock [MHz] STAT          |  222670.8948 |  3089.3231 |  3095.7052 |
3092.6513 |  3091.8668 |  3092.7182 |  3093.2810 |
|              CPI STAT               |      28.8774 |     0.3839 |     0.4241 |
0.4011 |     0.3925 |     0.4006 |     0.4072 |
|  L2D load bandwidth [MBytes/s] STAT | 2.785533e+06 | 36650.1579 | 40741.6806 |
38687.9545 | 38040.1631 | 38699.5999 | 39295.1659 |
|  L2D load data volume [GBytes] STAT |       9.5598 |     0.1323 |     0.1340 |
0.1328 |     0.1326 |     0.1327 |     0.1327 |
| L2D evict bandwidth [MBytes/s] STAT | 1.338589e+06 | 17660.6316 | 19580.9533 |
18591.5107 | 18283.0646 | 18598.6165 | 18873.0047 |
| L2D evict data volume [GBytes] STAT |       4.5944 |     0.0635 |     0.0645 |
0.0638 |     0.0637 |     0.0638 |     0.0638 |
|     L2 bandwidth [MBytes/s] STAT    | 4.391513e+06 | 57953.1404 | 64016.6093 |
60993.2383 | 60079.6459 | 61093.7766 | 62084.8399 |
|     L2 data volume [GBytes] STAT    |      15.0718 |     0.2083 |     0.2119 |
0.2093 |     0.2089 |     0.2092 |     0.2096 |
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+

Region: p
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+
|               Metric                |      Sum    |     Min    |     Max    |     Avg
|   %ile 25  |   %ile 50  |   %ile 75  |
+-------------------------------------+-------------+------------+------------+----------
--+------------+------------+------------+
|       Runtime (RDTSC) [s] STAT      |       0.1338 |     0.0016 |     0.0029 |
0.0019 |     0.0017 |     0.0017 |     0.0018 |
|      Runtime unhalted [s] STAT      |       0.2269 |     0.0029 |     0.0044 |
0.0032 |     0.0029 |     0.0030 |     0.0031 |
|           Clock [MHz] STAT          |  222689.7797 |  3088.1223 |  3097.5048 |
```

```
3092.9136 |  3091.8890 |  3092.6854 |  3093.9027 |
|               CPI STAT              |      41.5487 |     0.5189 |     0.8024 |
0.5771 |     0.5361 |     0.5511 |     0.5631 |
|  L2D load bandwidth [MBytes/s] STAT | 2.690730e+06 | 24050.8747 | 43159.9249 |
37371.2501 | 37660.7877 | 39084.4932 | 40764.6936 |
|  L2D load data volume [GBytes] STAT |       4.9410 |     0.0681 |     0.0696 |
0.0686 |     0.0684 |     0.0686 |     0.0687 |
| L2D evict bandwidth [MBytes/s] STAT | 1.276865e+06 | 11135.3983 | 20283.3914 |
17734.2384 | 17585.9803 | 18365.8703 | 19063.1122 |
| L2D evict data volume [GBytes] STAT |       2.3111 |     0.0318 |     0.0325 |
0.0321 |     0.0320 |     0.0321 |     0.0322 |
|     L2 bandwidth [MBytes/s] STAT    | 4.489000e+06 | 39640.4234 | 71142.1217 |
62347.2238 | 61846.2077 | 64314.6028 | 67060.3150 |
|     L2 data volume [GBytes] STAT    |       8.1283 |     0.1119 |     0.1145 |
0.1129 |     0.1125 |     0.1128 |     0.1133 |
+-------------------------------------+--------------+------------+------------+----------
--+------------+------------+------------+
```

```
$ likwid-mpirun -s 0x0 -n 1 -t 72 -m -g L3 ./tea_leaf
Region: w
+-----------------------------------------+-------------+------------+------------+--------
----+------------+------------+------------+
|                 Metric                  |     Sum     |    Min     |     Max    |     Avg
|   %ile 25  |   %ile 50  |   %ile 75  |
+-----------------------------------------+-------------+------------+------------+--------
----+------------+------------+------------+
|        Runtime (RDTSC) [s] STAT         |      0.3942 |     0.0052 |     0.0061 |
0.0055 |     0.0053 |     0.0054 |     0.0056 |
|        Runtime unhalted [s] STAT        |      0.5618 |     0.0074 |     0.0086 |
0.0078 |     0.0076 |     0.0077 |     0.0079 |
|            Clock [MHz] STAT             | 222673.7604 |  3091.1645 |  3094.3836 |
3092.6911 |  3092.2116 |  3092.6674 |  3093.0306 |
|                CPI STAT                 |     23.7804 |     0.3146 |     0.3651 |
0.3303 |     0.3211 |     0.3279 |     0.3330 |
|    L3 load bandwidth [MBytes/s] STAT    | 547763.1061 |  6093.6275 | 10032.7011 |
7607.8209 |  7012.5762 |  7554.7574 |  8100.1629 |
|    L3 load data volume [GBytes] STAT    |      3.0001 |     0.0318 |     0.0557 |
0.0417 |     0.0384 |     0.0412 |     0.0443 |
|   L3 evict bandwidth [MBytes/s] STAT    | 383278.5168 |  4220.2546 |  7367.1619 |
5323.3127 |  4954.9146 |  5286.2647 |  5743.4389 |
|   L3 evict data volume [GBytes] STAT    |      2.0987 |     0.0222 |     0.0392 |
0.0291 |     0.0269 |     0.0288 |     0.0318 |
| L3|MEM evict bandwidth [MBytes/s] STAT | 383781.8171 |  4220.9802 |  7367.7036 |
5330.3030 |  4957.5438 |  5287.2976 |  5744.0787 |
| L3|MEM evict data volume [GBytes] STAT |      2.1014 |     0.0222 |     0.0392 |
0.0292 |     0.0269 |     0.0288 |     0.0318 |
|  Dropped CLs bandwidth [MBytes/s] STAT | 164058.0066 |  1607.6004 |  2959.2598 |
2278.5834 |  2100.6162 |  2290.1839 |  2490.5356 |
|  Dropped CLs data volume [GBytes] STAT |      0.8984 |     0.0085 |     0.0165 |
0.0125 |     0.0113 |     0.0124 |     0.0135 |
|      L3 bandwidth [MBytes/s] STAT       | 931544.9223 | 10326.5545 | 17400.4047 |
12938.1239 | 11968.5625 | 12906.7665 | 13834.4516 |
|      L3 data volume [GBytes] STAT       |      5.1019 |     0.0540 |     0.0949 |
0.0709 |     0.0641 |     0.0697 |     0.0762 |
+-----------------------------------------+-------------+------------+------------+--------
----+------------+------------+------------+
Region: ur
+-----------------------------------------+-------------+-----------+------------+---------
---+------------+-----------+------------+
|                 Metric                  |     Sum     |    Min    |     Max    |     Avg
|   %ile 25  |  %ile 50  |   %ile 75  |
+-----------------------------------------+-------------+-----------+------------+---------
---+------------+-----------+------------+
|        Runtime (RDTSC) [s] STAT         |      0.2498 |    0.0033 |     0.0037 |
0.0035 |     0.0034 |    0.0035 |     0.0035 |
|        Runtime unhalted [s] STAT        |      0.3766 |    0.0050 |     0.0056 |
0.0052 |     0.0051 |    0.0052 |     0.0053 |
|            Clock [MHz] STAT             | 222657.6534 | 3090.0062 |  3094.8950 |
3092.4674 |  3091.8772 | 3092.2480 |  3093.0151 |
|                CPI STAT                 |     29.0267 |    0.3845 |     0.4327 |
0.4031 |     0.3941 |    0.4031 |     0.4083 |
|    L3 load bandwidth [MBytes/s] STAT    | 529451.4046 | 5582.9426 |  9318.9494 |
7353.4917 |  6773.1815 | 7400.0644 |  7896.2399 |
|    L3 load data volume [GBytes] STAT    |      1.8374 |    0.0191 |     0.0335 |
0.0255 |     0.0234 |    0.0252 |     0.0278 |
|   L3 evict bandwidth [MBytes/s] STAT    | 389961.5531 | 3882.2973 |  7313.7207 |
5416.1327 |  4885.2790 | 5400.2119 |  5798.9192 |
|   L3 evict data volume [GBytes] STAT    |      1.3535 |    0.0135 |     0.0269 |
0.0188 |     0.0170 |    0.0185 |     0.0204 |
| L3|MEM evict bandwidth [MBytes/s] STAT | 390504.3299 | 3883.2392 |  7317.6934 |
5423.6712 |  4908.4795 | 5422.2550 |  5799.7052 |
| L3|MEM evict data volume [GBytes] STAT |      1.3553 |    0.0135 |     0.0269 |
0.0188 |     0.0170 |    0.0185 |     0.0204 |
|  Dropped CLs bandwidth [MBytes/s] STAT | 139005.5038 | 1587.1651 |  2474.6792 |
1930.6320 |  1805.2347 | 1915.8014 |  2072.0529 |
|  Dropped CLs data volume [GBytes] STAT |      0.4817 |    0.0052 |     0.0087 |
0.0067 |     0.0062 |    0.0066 |     0.0071 |
|      L3 bandwidth [MBytes/s] STAT       | 919955.7344 | 9466.1818 | 16539.1931 |
12777.1630 | 11769.9133 |     12965 | 14005.5416 |
|      L3 data volume [GBytes] STAT       |      3.1932 |    0.0328 |     0.0604 |
0.0444 |     0.0404 |    0.0438 |     0.0481 |
```

```
+----------------------------------------+-------------+-----------+------------+---------
---+------------+-----------+------------+

Region: p
+----------------------------------------+-------------+-----------+------------+---------
---+------------+------------+------------+
|                 Metric                 |     Sum     |    Min    |     Max    |     Avg
|   %ile 25  |   %ile 50  |   %ile 75  |
+----------------------------------------+-------------+-----------+------------+---------
---+------------+------------+------------+
|        Runtime (RDTSC) [s] STAT        |      0.1334 |    0.0016 |     0.0029 |
0.0019 |     0.0017 |     0.0017 |     0.0018 |
|        Runtime unhalted [s] STAT       |      0.2282 |    0.0029 |     0.0044 |
0.0032 |     0.0030 |     0.0030 |     0.0031 |
|            Clock [MHz] STAT            | 222662.5467 | 3088.3580 |  3097.4498 |
3092.5354 |  3091.0188 |  3092.5239 |  3093.6255 |
|                CPI STAT                |     41.6914 |    0.5153 |     0.8152 |
0.5790 |     0.5387 |     0.5510 |     0.5631 |
|    L3 load bandwidth [MBytes/s] STAT   | 447897.4427 | 4609.6530 |  7479.4691 |
6220.7978 |  6106.5536 |  6362.6569 |  6559.6174 |
|    L3 load data volume [GBytes] STAT   |      0.8191 |    0.0088 |     0.0145 |
0.0114 |     0.0108 |     0.0112 |     0.0116 |
|   L3 evict bandwidth [MBytes/s] STAT   | 303649.1016 | 2688.5355 |  5081.6161 |
4217.3486 |  4096.4846 |  4459.4881 |  4617.8252 |
|   L3 evict data volume [GBytes] STAT   |      0.5574 |    0.0061 |     0.0091 |
0.0077 |     0.0074 |     0.0077 |     0.0081 |
| L3|MEM evict bandwidth [MBytes/s] STAT | 307169.2154 | 2688.8614 |  5164.4851 |
4266.2391 |  4090.4781 |  4439.8411 |  4605.4490 |
| L3|MEM evict data volume [GBytes] STAT |      0.5578 |    0.0061 |     0.0091 |
0.0077 |     0.0074 |     0.0077 |     0.0081 |
|  Dropped CLs bandwidth [MBytes/s] STAT | 140841.8846 | 1476.5916 |  2461.3722 |
1956.1373 |  1803.0280 |  1963.3892 |  2103.9631 |
|  Dropped CLs data volume [GBytes] STAT |      0.2618 |    0.0023 |     0.0061 |
0.0036 |     0.0030 |     0.0035 |     0.0038 |
|      L3 bandwidth [MBytes/s] STAT      | 755066.6579 | 7298.5144 | 12643.9542 |
10487.0369 | 10744.3701 | 11099.2378 | 11641.3543 |
|      L3 data volume [GBytes] STAT      |      1.3772 |    0.0148 |     0.0237 |
0.0191 |     0.0181 |     0.0190 |     0.0196 |
+----------------------------------------+-------------+-----------+------------+---------
---+------------+------------+------------+
```

```
$ likwid-mpirun -s 0x0 -n 1 -t 72 -m -g MEM ./tea_leaf
Region: w
+----------------------------------------+-------------+-----------+-----------+----------
-+-----------+-----------+-----------+
|                 Metric                 |     Sum     |    Min    |    Max    |    Avg
|  %ile 25  |  %ile 50  |  %ile 75  |
+----------------------------------------+-------------+-----------+-----------+----------
-+-----------+-----------+-----------+
|        Runtime (RDTSC) [s] STAT        |      0.3954 |    0.0051 |    0.0062 |    0.0055
|    0.0053 |    0.0054 |    0.0055 |
|        Runtime unhalted [s] STAT       |      0.5703 |    0.0073 |    0.0088 |    0.0079
|    0.0077 |    0.0078 |    0.0080 |
|           Clock [MHz] STAT             | 222665.2092 | 3091.1517 | 3094.1924 | 3092.5724
| 3092.0383 | 3092.5545 | 3093.1067 |
|               CPI STAT                 |     23.9898 |    0.3079 |    0.3715 |    0.3332
|    0.3241 |    0.3302 |    0.3368 |
|  Memory read bandwidth [MBytes/s] STAT |   1759.2025 |         0 | 1578.5205 |   24.4334
|         0 |         0 |         0 |
|  Memory read data volume [GBytes] STAT |      0.0093 |         0 |    0.0083 |    0.0001
|         0 |         0 |         0 |
| Memory write bandwidth [MBytes/s] STAT |   2987.1178 |         0 | 2251.2698 |   41.4877
|         0 |         0 |         0 |
| Memory write data volume [GBytes] STAT |      0.0160 |         0 |    0.0118 |    0.0002
|         0 |         0 |         0 |
|    Memory bandwidth [MBytes/s] STAT    |   4746.3202 |         0 | 3829.7902 |   65.9211
|         0 |         0 |         0 |
|    Memory data volume [GBytes] STAT    |      0.0253 |         0 |    0.0201 |    0.0004
|         0 |         0 |         0 |
+----------------------------------------+-------------+-----------+-----------+----------
-+-----------+-----------+-----------+
Region: ur
+----------------------------------------+-------------+-----------+-----------+----------
-+-----------+-----------+-----------+
|                 Metric                 |     Sum     |    Min    |    Max    |    Avg
|  %ile 25  |  %ile 50  |  %ile 75  |
+----------------------------------------+-------------+-----------+-----------+----------
-+-----------+-----------+-----------+
|        Runtime (RDTSC) [s] STAT        |      0.2492 |    0.0033 |    0.0037 |    0.0035
|    0.0034 |    0.0034 |    0.0035 |
|        Runtime unhalted [s] STAT       |      0.3852 |    0.0051 |    0.0060 |    0.0054
|    0.0052 |    0.0053 |    0.0054 |
|           Clock [MHz] STAT             | 222657.9436 | 3089.2719 | 3094.5525 | 3092.4714
| 3091.6544 | 3092.3975 | 3093.3116 |
|               CPI STAT                 |     29.3795 |    0.3905 |    0.4610 |    0.4080
|    0.4001 |    0.4056 |    0.4128 |
|  Memory read bandwidth [MBytes/s] STAT |   2582.7665 |         0 | 2442.6024 |   35.8718
|         0 |         0 |         0 |
|  Memory read data volume [GBytes] STAT |      0.0090 |         0 |    0.0085 |    0.0001
|         0 |         0 |         0 |
| Memory write bandwidth [MBytes/s] STAT |   3443.2911 |         0 | 3222.7850 |   47.8235
|         0 |         0 |         0 |
| Memory write data volume [GBytes] STAT |      0.0120 |         0 |    0.0112 |    0.0002
|         0 |         0 |         0 |
|    Memory bandwidth [MBytes/s] STAT    |   6026.0574 |         0 | 5665.3873 |   83.6952
|         0 |         0 |         0 |
|    Memory data volume [GBytes] STAT    |      0.0209 |         0 |    0.0197 |    0.0003
|         0 |         0 |         0 |
+----------------------------------------+-------------+-----------+-----------+----------
-+-----------+-----------+-----------+
Region: p
+----------------------------------------+-------------+-----------+------------+---------
--+-----------+-----------+-----------+
|                 Metric                 |     Sum     |    Min    |     Max    |    Avg
|  %ile 25  |  %ile 50  |  %ile 75  |
+----------------------------------------+-------------+-----------+------------+---------
--+-----------+-----------+-----------+
|        Runtime (RDTSC) [s] STAT        |      0.1407 |    0.0016 |     0.0030 |
0.0020 |    0.0018 |    0.0019 |    0.0019 |
|        Runtime unhalted [s] STAT       |      0.2351 |    0.0030 |     0.0045 |
0.0033 |    0.0031 |    0.0031 |    0.0032 |
|           Clock [MHz] STAT             | 222676.9137 | 3088.2191 |  3098.0323 |
3092.7349 | 3091.4886 | 3092.8811 | 3093.6885 |
|               CPI STAT                 |     41.9745 |    0.5229 |     0.8042 |
```

```
0.5830 |    0.5435 |    0.5563 |    0.5662 |
|  Memory read bandwidth [MBytes/s] STAT |   5671.5134 |         0 |  4543.8039 |
78.7710 |         0 |         0 |         0 |
|  Memory read data volume [GBytes] STAT |      0.0103 |         0 |     0.0074 |
0.0001 |         0 |         0 |         0 |
| Memory write bandwidth [MBytes/s] STAT |   7017.0959 |         0 |  5887.1373 |
97.4597 |         0 |         0 |         0 |
| Memory write data volume [GBytes] STAT |      0.0125 |         0 |     0.0096 |
0.0002 |         0 |         0 |         0 |
|    Memory bandwidth [MBytes/s] STAT    |  12688.6093 |         0 | 10430.9412 |
176.2307 |         0 |         0 |         0 |
|    Memory data volume [GBytes] STAT    |      0.0228 |         0 |     0.0170 |
0.0003 |         0 |         0 |         0 |
+----------------------------------------+-------------+-----------+------------+---------
--+-----------+-----------+-----------+
```

Note: Here we present only single measurements for each group but it is beneficial to run the measurements like the scaling tests and produce plots for each region `for i in {1..72}; do likwid-mpirun -np 1 -t $i -g GROUP -m ./tea_leaf; done`.

Task6: Discussion of hot spot measurements

We can see that the default test case fits mostly into the cache hierarchy. Access to memory are rare.

The measurements show that for optimizing the code we should check:

- Higher vectorization. The test system supports AVX512. We should adjust the compiler optimization flags in the beginning!
- Fix load imbalance

Note: While we used short sentences and sparse explaination to the findings, you can be more expressive to use this document already as final report.