Intel Advisor on Cori

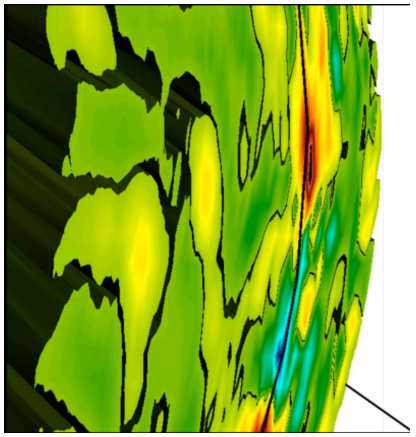

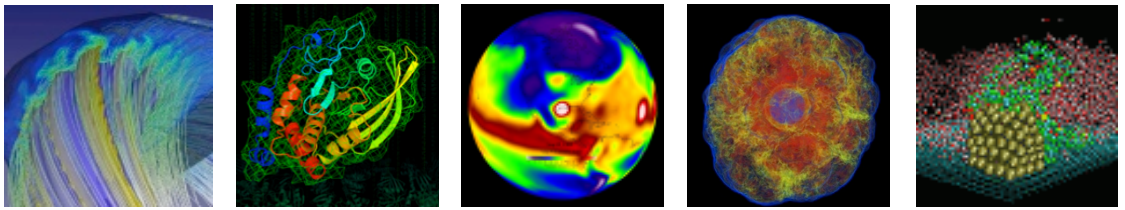

Charlene Yang

Application Performance Group
cjyang@lbl.gov

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Intel® Advisor: automatic and refined

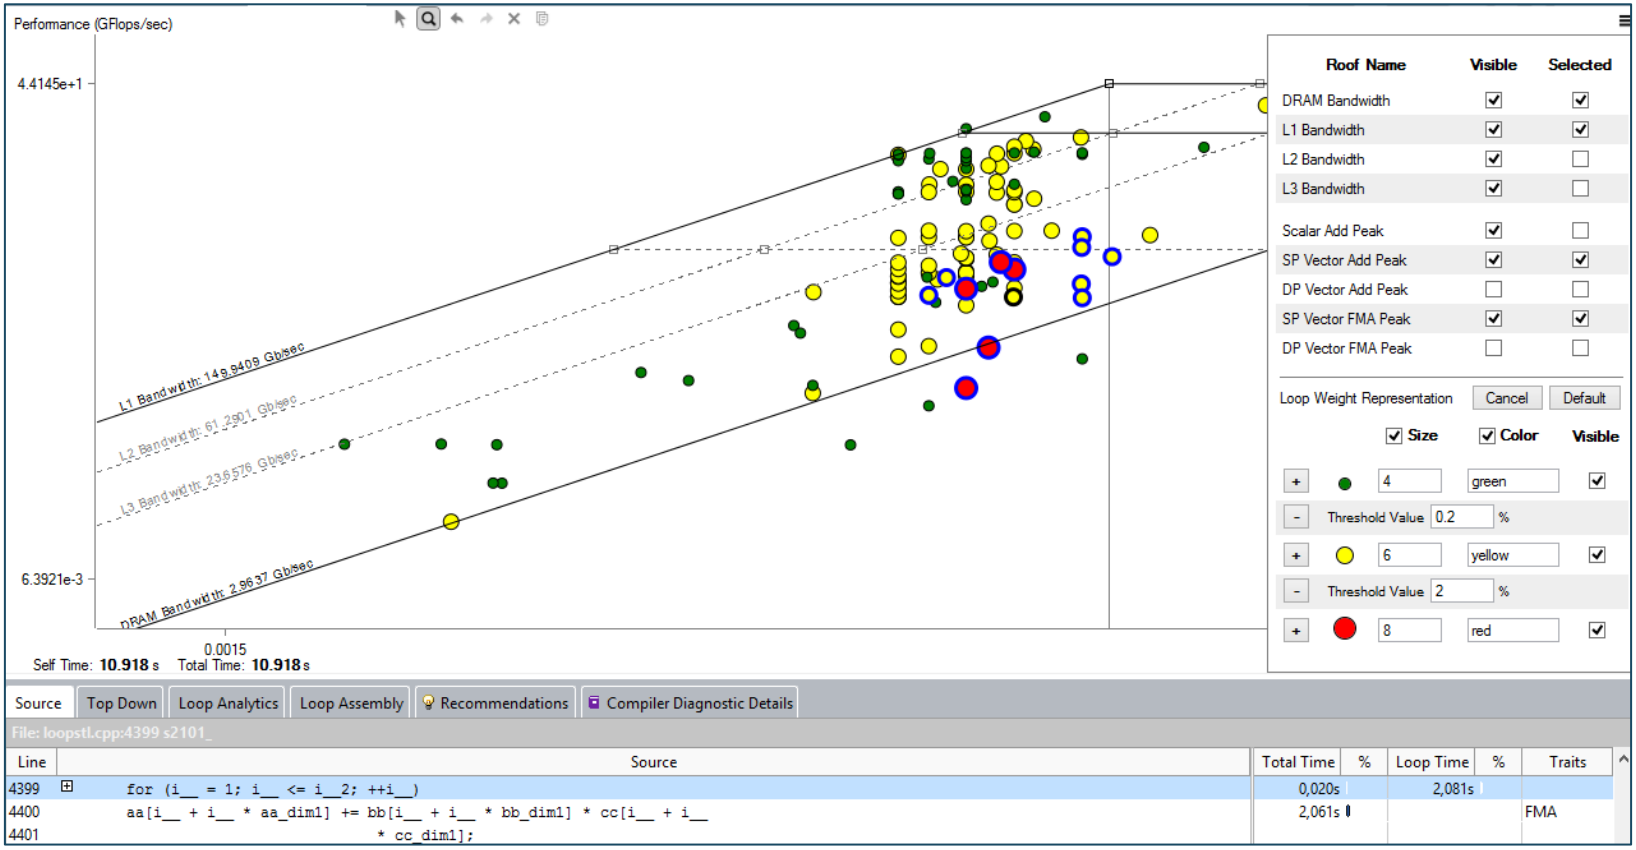

Intel® Advisor: code analytics

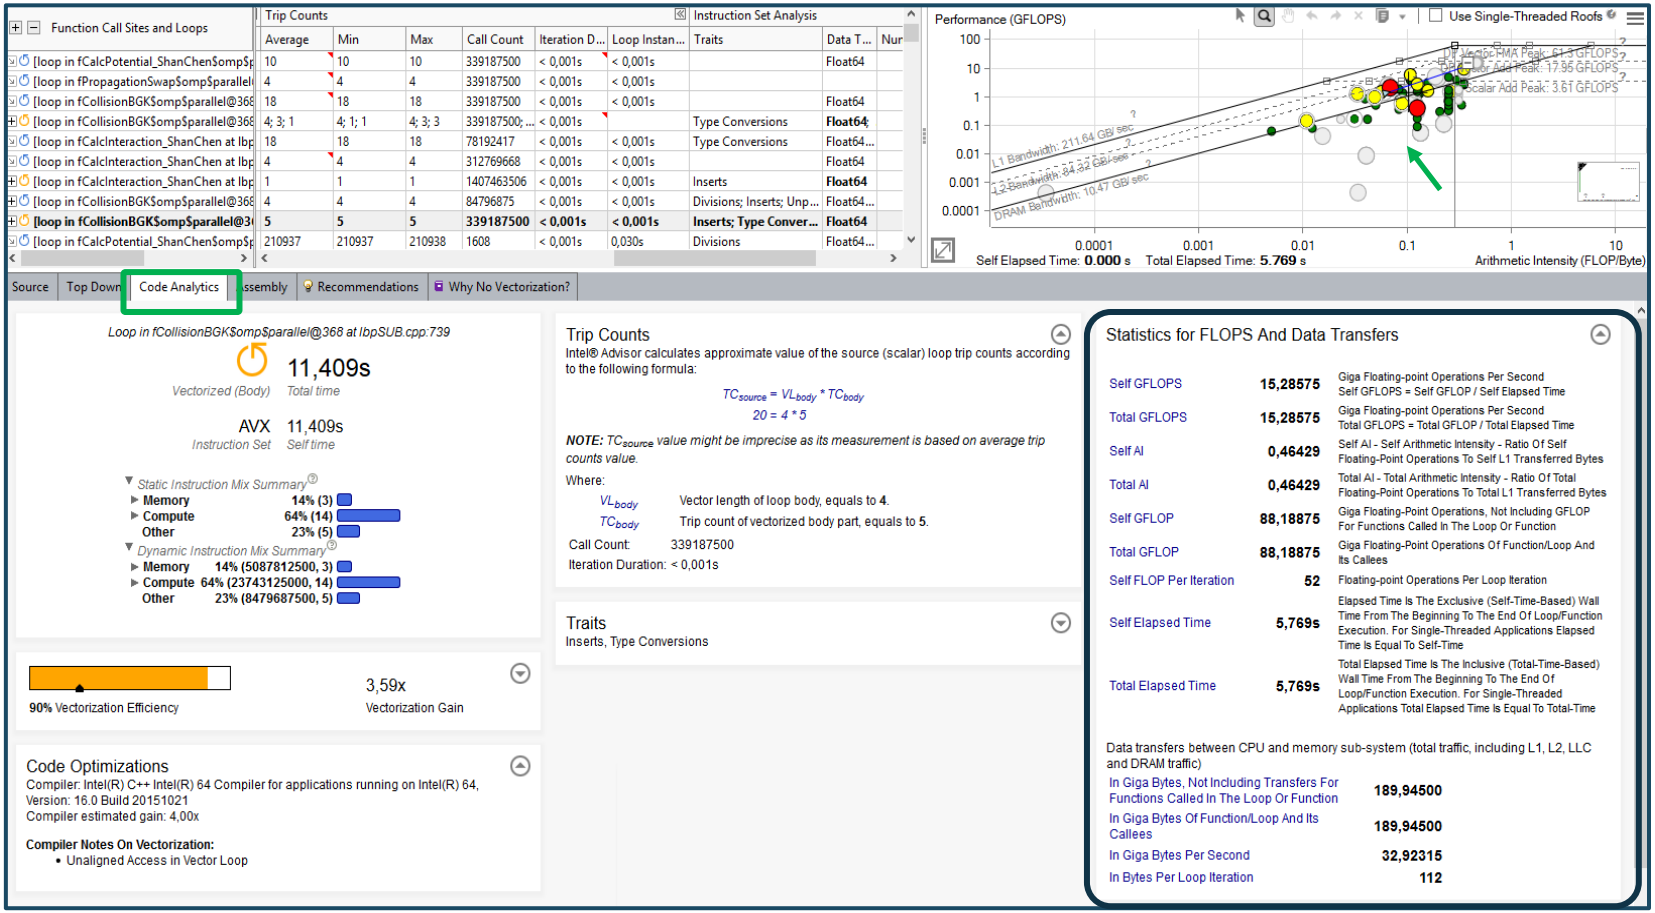

Intel® Advisor: 2-step data collection

Roofline : Axis X: **AI** = #**FLOP** / #**Bytes** Axis Y: **FLOP/S** = #**FLOP** (mask aware) / #**Seconds**	Overhead
Step 1: Survey (-collect survey) - Provide #**Seconds** - *Root access not needed* - User mode sampling, non-intrusive.	**1x**
Step 2: FLOPS *(-collect tripcounts –flops)* - Provide #**FLOP**, #**Bytes**, AVX-512 Mask - *Root access not needed* - Precise, instrumentation based, count number of instructions	**3-5x**

Setup on Cori

- **Compile with -g to get debugging info**

```
cc -g -dynamic -openmp -O2 -o mycode.exe mycode.c
```

- **Cache-Aware Roofline Model (CARM)**

```
module load advisor/2018.up1
```

- **Integrated Roofline Model (Cache Simulator)**

```
module load advisor/2018.integrated_roofline.up1
```

- **Incompatible GUI for regular and integrated Advisor**

Run Advisor on Cori

- **Start an interactive session on a KNL node**

```
salloc --qos=interactive -C knl -N 1 -t hh:mm:ss -A <your_account>
```

- **To collect data for roofline, do two collections: survey and tripcounts.**

```
srun -n <num-of-ranks> -c <num_of_cores_per_rank> advixe-cl -v
-collect survey -no-auto-finalize  -project-dir=<same_dir_name>
-data-limit=0 -- <your_executable>
```

```
srun -n <num-of-ranks> -c <num_of_cores_per_rank> advixe-cl -v
-collect tripcounts -flops -no-auto-finalize -project-dir=<same_dir_name>  -
data-limit=0 -- <your_executable>
```

- **Run on the Lustre filesystem $SCRATCH**
- **Finalization is expensive especially on KNL: do it offline!**

Pack/View results

- **Pack results/source file/binary (already packed in demo)**

```
advixe-cl --snapshot --project-dir <same_dir_name> --pack --cache-sources
--cache-binaries -- <target_file_name>
```

- **Load module**

```
module load advisor/2018.integrated_roofline.up1
module show advisor/2018.integrated_roofline.up1
```

```
setenv ADVISOR_XE_2018_DIR
/global/common/cori/software/intel/advisor_2018.0.2.537542
```

- **Copy file and view results**

```
cp -r $ADVISOR_XE_2018_DIR/ECP-meeting-tutorial/ .
cd ECP-meeting-tutorial/
```

```
advxe-gui stencil.advixeexpz
```

- **or use NX: https://nxcloud01.nersc.gov**

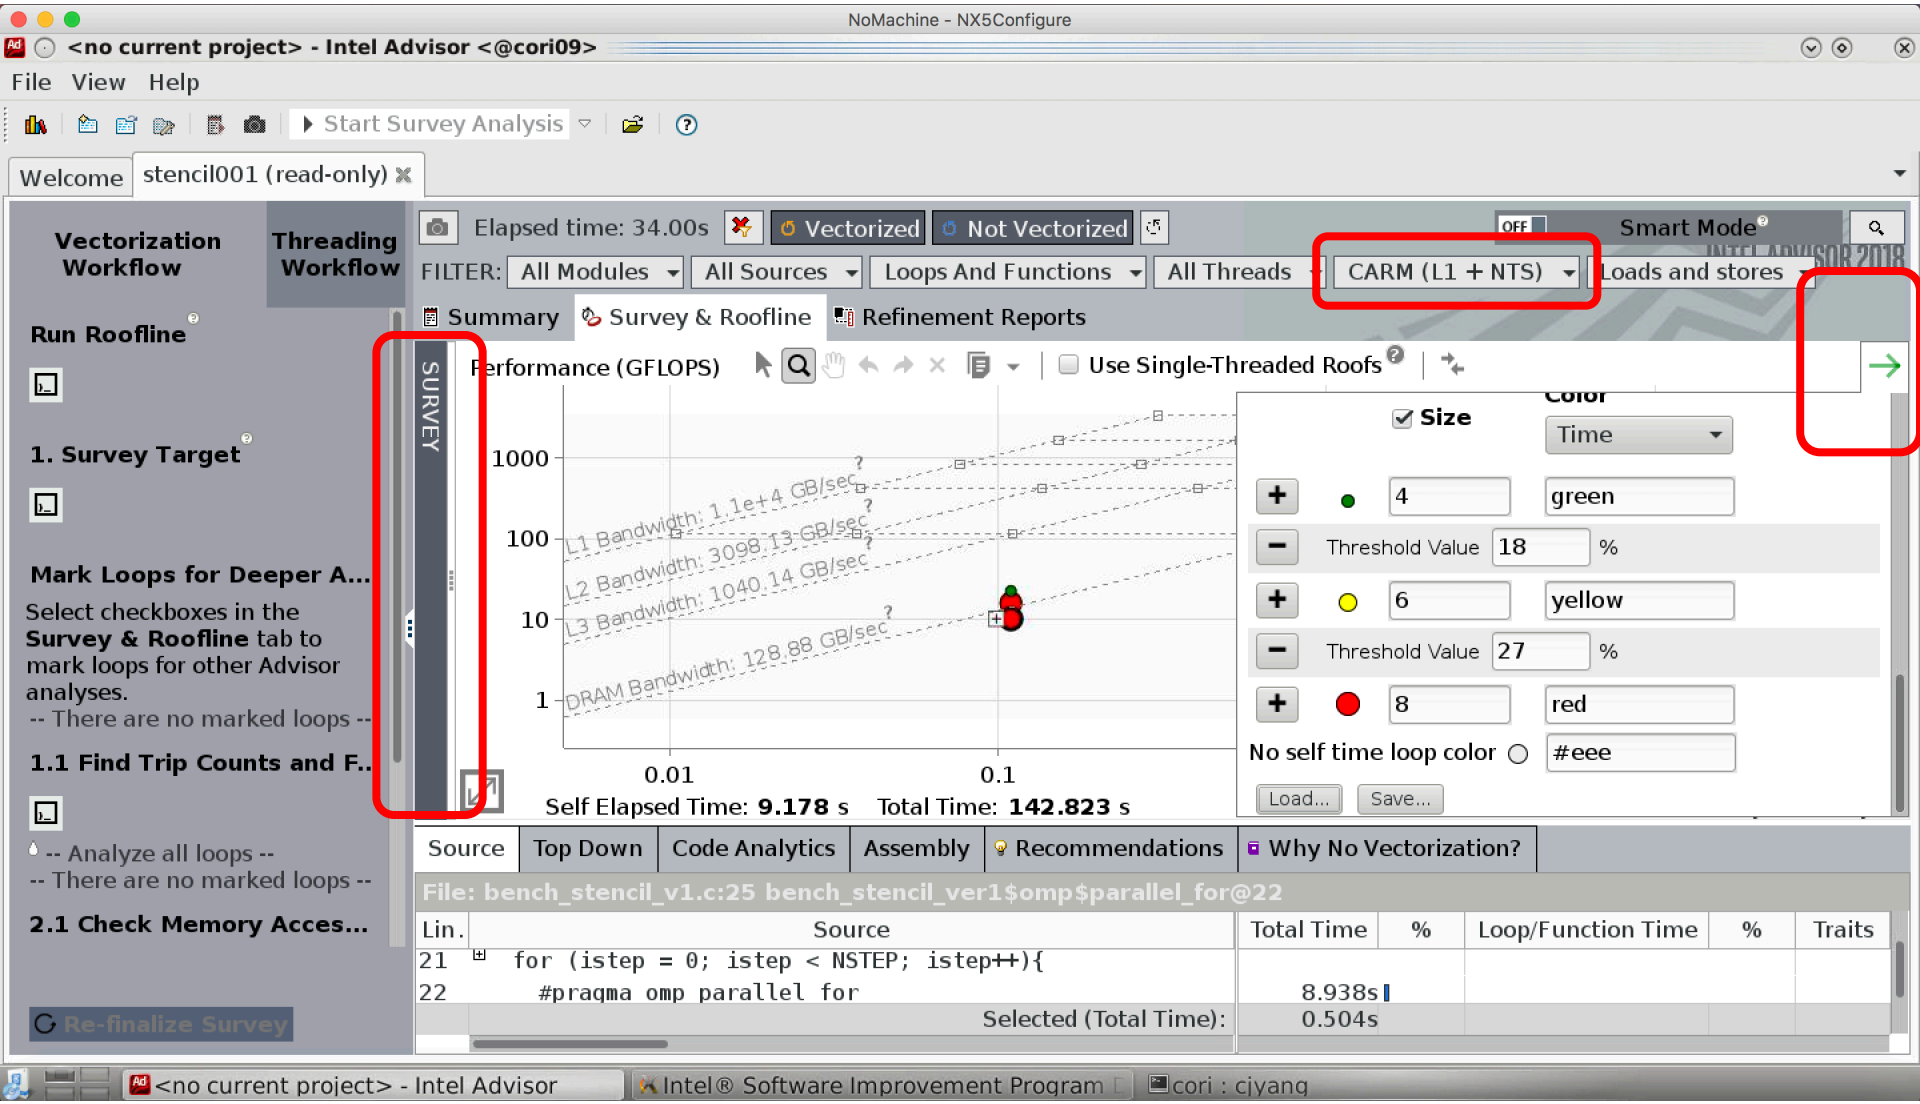
NoMachine - NX5Configure
<no current project> - Intel Advisor <@cori09>
File View Help
Start Survey Analysis
Welcome
stencil001 (read-only)
Vectorization Workflow
Threading Workflow
Run Roofline
1. Survey Target
Mark Loops for Deeper A...
Select checkboxes in the Survey & Roofline tab to mark loops for other Advisor analyses.
-- There are no marked loops --
1.1 Find Trip Counts and F...
-- Analyze all loops --
-- There are no marked loops --
2.1 Check Memory Acces...
Re-finalize Survey
Elapsed time: 34.00s
Vectorized
Not Vectorized
Smart Mode
FILTER: All Modules
All Sources
Loops And Functions
All Threads
CARM (L1 + NTS)
Loads and stores
Summary
Survey & Roofline
Refinement Reports
SURVEY
Performance (GFLOPS)
Use Single-Threaded Roofs
1000
100
10
1
L1 Bandwidth: 1.1e+4 GB/sec
L2 Bandwidth: 3098.13 GB/sec
L3 Bandwidth: 1040.14 GB/sec
DRAM Bandwidth: 128.88 GB/sec
0.01
0.1
Self Elapsed Time: 9.178 s
Total Time: 142.823 s
Size
Color
Time
4
green
Threshold Value 18 %
6
yellow
Threshold Value 27 %
8
red
No self time loop color
#eee
Load...
Save...
Source
Top Down
Code Analytics
Assembly
Recommendations
Why No Vectorization?
File: bench_stencil_v1.c:25 bench_stencil_ver1ompparallel_for@22
Lin.
Source
Total Time
%
Loop/Function Time
%
Traits
21 for (istep = 0; istep < NSTEP; istep++){
22 #pragma omp parallel for
8.938s
Selected (Total Time):
0.504s
<no current project> - Intel Advisor
Intel® Software Improvement Program
cori : cjyang

GPP optimization - cache blocking

Before

```
do my_igp = 1, ngpown ! OpenMP
 do iw = 1 , 3
     do ig = 1, igmax
         load wtilde_array(ig,my_igp) !512KB per row
         load aqsntemp(ig,n1) !512KB per row
         load eps(ig,my_igp) !512KB per row
```

After

```
do my_igp = 1, ngpown ! OpenMP
 do igbeg = 1, igmax, igblk
     do iw = 1 , 3
         do ig = igbeg, min(igbeg + igblk,igmax)
             load wtilde_array(ig,my_igp)
             load aqsntemp(ig,n1)
             load eps(ig,my_igp)
```

GPP optimization - changes in AI

- **L1 AI and L2 AI remain roughly the same on Haswell and KNL**
- **LLC AI improves by 3x due to fixed trip count 3 in loop *iw***
- **GFLOPS/s improved by 16-18%, less than 3x, because there are divide, shuffle and unpack instructions involved in the innermost loop**

	Haswell		KNL	
gppKernel.f90:303	Original	Cache-blocked	Original	Cache-blocked
L1+NTS AI	0.31	0.31	0.64	0.64
L2 AI	0.46	0.62	0.51	0.69
LLC AI	0.46	1.4	0.56	1.71
DRAM AI	3.31	3.44	26.45	26.83
GFLOPS/s	148.57	172.48	242.55	287.28

HSW (top) v.s. KNL (bottom)

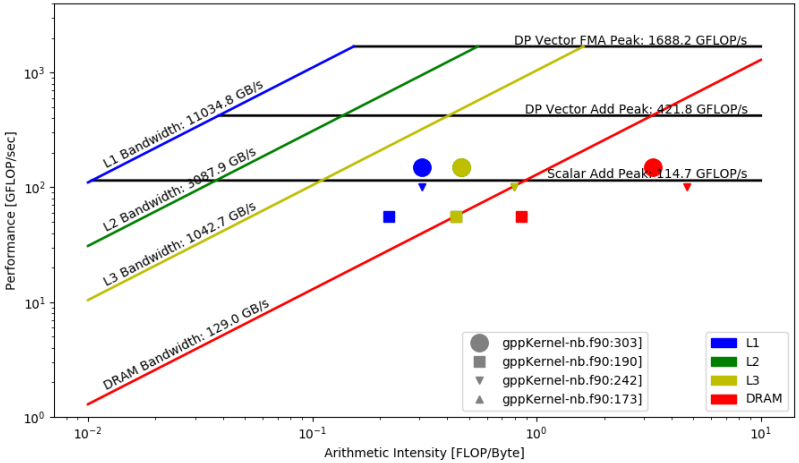

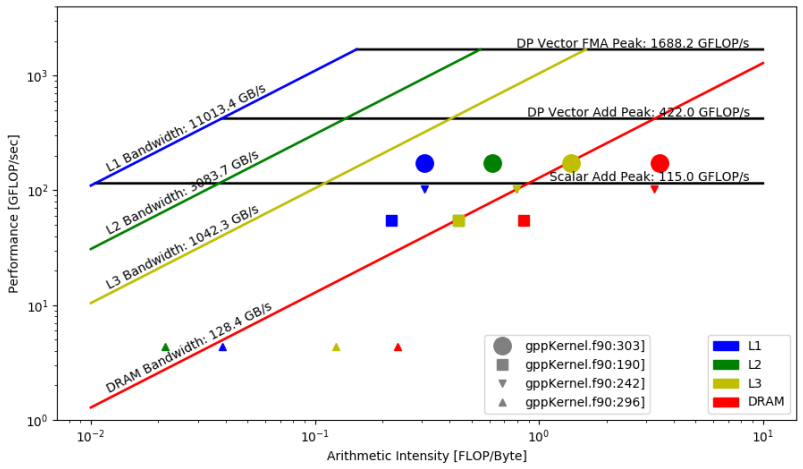

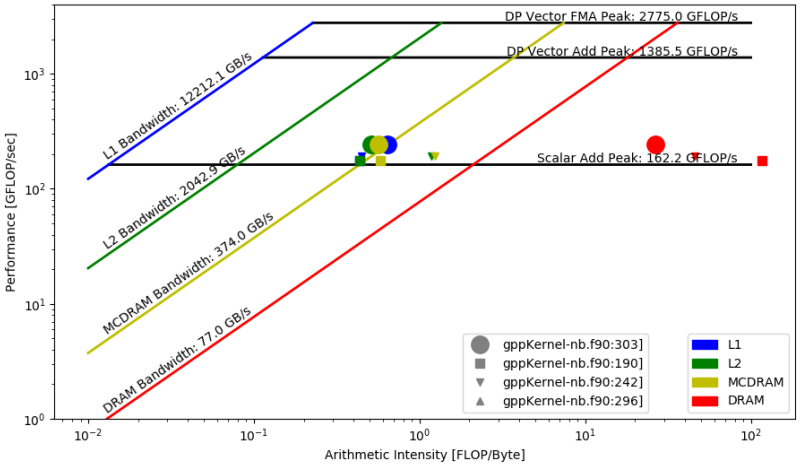

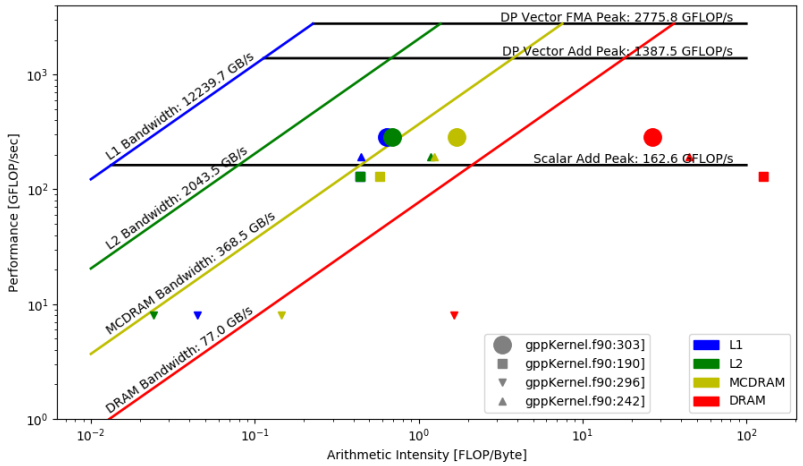


NERSC