Guiding **Optimization on** with the Roofline Model















February, 2020







How to Enable NERSC's diverse community of 7,000 users, 750 projects, and 700 codes to run on advanced architectures like Cori, Perlmutter and beyond?







What was different about Cori?

Edison ("Ivy Bridge):

- 5576 nodes
- 24 physical cores per node
- 48 virtual cores per node
- 2.4 3.2 GHz
- 8 double precision ops/cycle
- 64 GB of DDR3 memory (2.5 GB per physical core)
- ~100 GB/s Memory Bandwidth

Cori ("Knights Landing"):

- 9304 nodes
- 68 physical cores per node
- 272 virtual cores per node
- 1.4 1.6 GHz
- 32 double precision ops/cycle
- 16 GB of fast memory 96GB of DDR4 memory
- Fast memory has 400 500 GB/s No L3 Cache





Science teams need a simple way to wrap their heads around performance when main focus is scientific productivity:

- 1. Need a sense of absolute performance when optimizing applications.
- How Do I know if My Performance is Good?
- Why am I not getting peak performance advertised -
- How Do I know when to stop?

2. Many potential optimization directions:

- How do I know which to apply? -
- What is the limiting factor in my app's performance? -
- Again, how do I know when to stop?







Optimizing Code For Cori is like:

A. A Staircase ?

A. A Labyrinth ?

A. A Space Elevator? 4





(More) Optimized Code







Roofline helps visualize this information! Guides optimizations

WARP Optimizations:

- Add tiling over grid targeting L2 cache on both Xeon-Phi Systems
- Add particle sorting to further improve locality and memory access pattern
- Apply vectorization over particles



(a) Haswell Roofline



(b) KNL Roofline









NESAP Example

















BerkeleyGW

- A massively parallel package • for GW calculations
- Sits on top of DFT codes
- **Computational motifs**
 - **FFTs**
 - Dense linear algebra ____
 - Large reductions _











Sigma-GPP

Pseudo Code

```
do n1 = 1, nbands n' e.g. 2763
  do igp = 1, ngpown G' e.g. 6633
     do ig = 1, ncouls G e.g. 26529
        do iw = 1, nw E e.g. 3
           compute: 1. mixed data types
                   e.g. complex double, double, integer
                2. various memory access patterns
                   e.g. (ig, igp) (ig, n1) (igp, n1) (iw, n1) (n1)
                3. complex number divisions
                4. nw is very small, will be unrolled
           reduction: 1. complex numbers
                2. all top 3 loops, billions of iterations
```







Optimization process for GPP Kernel

- 1. Add OpenMP
- Initial Vectorization (loop reordering, 2. conditional removal)
- Cache-Blocking 3.
- Improved Vectorization (Divides) 4.
- Hyper-threading 5.





Optimization Process











Vectorization





ngpown typically in 100's to 1000s. Good for many threads.

Original inner loop.

ncouls typically in 1000s - 10,000s.

Attempt to save work breaks vectorization



Change in Roofline

Office of

Science

J.S. DEPARTMENT OF



KNL Roofline Optimization Path



The loss of L3 on MIC makes locality more important.





```
!$OMP DO
do my_igp = 1, ngpown
                                                                     Required Cache size to reuse 3 times:
    do iw = 1, 3
       do ig = 1, igmax
                                                                     1536 KB
          load wtilde_array(ig,my_igp) 819 MB, 512KB per row
                                                                     L2 on KNL is 512 KB per core
          load aqsntemp(ig,n1) 256 MB, 512KB per row
                                                                     L2 on Has. is 256 KB per core
          load l_eps_array(ig,my_igp) 819 MB, 512KB per row
          do work (including divide)
                                                                     L3 on Has. is 3800 KB per core
```

Without blocking we spill out of L2 on catch us.





KNL and Haswell. But, Haswell has L3 to





!\$OMP DO	
do my_igp = 1, ngpown do igbeg = 1, igmax, igblk	Required Cache size
do iw = 1 , 3	1536 KB
do ig = igbeg, min(igbeg + igblk,igmax)	
load wtilde_array(ig,my_igp) 819 MB, 512KB per row	L2 on KNL is 512 KB p
load aqsntemp(ig,n1) 256 MB, 512KB per row	L2 on Has. is 256 KB
load I_eps_array(ig,my_igp) 819 MB, 512KB per row do work (including divide)	L3 on Has. is 3800 KB

Without blocking we spill out of L2 on catch us.





KNL and Haswell. But, Haswell has L3 to



to reuse 3 times:



Cache Blocking Optimization



KNL Roofline Optimization Path









Cache Blocking Optimization (Hierarchical Roofline)

Original Code

Cache-Blocking Code











Additional Speedups from Hyperthreading









GPP on GPUs in 8 Steps

- Collapse n', G', and G loops 1.
- Bring n' loop in; collapse only G' and G 2.
- 3. **Adjust threadblock size**
- **Reduce branching; pull iw loop outside** 4.
- 5. Swap indices to suite parallelisation
- Simplify code 6.
- **Replace div. with rcp. and mul.**
- **Replace abs with power of 2** 8.
- 9. Cache blocking









V1. Naïve Implementation

Collapse the first 3 loops to gain parallelism

!\$ACC PARALLEL LOOP COLLAPSE(3) REDUCTION(+:)
do n1 = 1, nbands
 do igp = 1, ngpown
 do ig = 1, ncouls
 do iw = 1, nw #unrolled
 compute and reduction

	TFLOPs	Time (sec)	TFLO
v1.collapse3	3.71	1.63	2.2







P/s



V2. More Compute Per Thread

Move n' loop in, and collapse the first 2 loops
 !\$ACC PARALLEL LOOP COLLAPSE(2) REDUCTION(+:)
 do igp = 1, ngpown
 do ig = 1, ncouls
 do n1 = 1, nbands #unrolled too!
 do iw = 1, nw #unrolled
 compute and reduction

	TFLOPs	Time	TFLOP/s
v1.collapse3	3.71	1.63 🔶	2.27
v2.collapse2	3.71	1.73	2.15











V2. More Compute Per Thread

- L2/HBM AI increases!
- **Register count at 186** •
 - Very low occupancy

Need more warps to hide la

8 warps per SM _____



Active Warps Per Scheduler [warp]	2.00 Instructions Per Active Issue Slot [inst/cycle]
Eligible Warps Per Scheduler [warp]	0.35 No Eligible [%]
Issued Warp Per Scheduler	0.30 One or More Eligible [%]



•







V3. Increase Threadblock Size

Force threadblack size to be 519 instead of the default 198 !\$ACC PARALLEL LOOP COLLAPSE(2) VECTOR LENGTH(512) REDUCTION(+:)

Register spills but performance may not be bad! •

0 bytes stack frame, 0 bytes spill stores, 0 bytes spill loads ptxas info : Used 186 registers, 624 bytes cmem[0], 32 bytes cmem[2]

104 bytes stack frame, 188 bytes spill stores, 168 bytes spill loads ptxas info : Used 128 registers, 624 bytes cmem[0], 32 bytes cmem[2]







V3. Increase Threadblock Size

• More bandwidth bound now bu latency hiding is successful!

	TFLOPs	Time	TFLOP/s
v2.collapse2	3.71	1.73	2.15
v3.vector512	3.71	1.40	2.65



Active Warps Per Scheduler [warp]	4.00 Instructions Per Active Issue Slot [inst/cycle]
Eligible Warps Per Scheduler [warp]	0.67 No Eligible [%]
Issued Warp Per Scheduler	0.42 One or More Eligible [%]









Bring iw loop outside of the kernel

#reduce branching do iw = 1, nw !\$ACC PARALLEL LOOP COLLAPSE(2) VECTOR LENGTH(512) REDUCTION(+:) do igp = 1, ngpown do ig = 1, ncouls do n1 = 1, nbands #unrolled compute and reduction

Fewer variables to be reduced -> lower register pressure 0 bytes stack frame, 0 bytes spill stores, 0 bytes spill loads ptxas info : Used 122 registers, 600 bytes cmem[0], 32 bytes cmem[2] Office of





V4. Reduce Branching

- Aggregated data for all kerne •
- **BRA** instruction court • 16% 14,278,897,053 5,975,051,812 x 2

	TFLOPs	Time	TFLOP/s
v3.vector512	3.71	1.40	2.65
v4.iwoutside	3.52	1.17	3.00







BERKELEY LAB

V5. Swap Indices

do iw = 1, nw
!\$ACC PARALLEL LOOP
do igp = 1, ngpown
<pre>do ig = 1, ncouls</pre>
do n1 = 1, nbands
<pre>wx_array(iw,n1) to (n1,iw)</pre>

	TFLOPs	Time	TFLOP/s
v4.iwoutside	3.52	1.17	3.00
v5.swapindices	3.52	1.16	3.03









V6. Simplify Code

- **Fewer instructions** -> less work
 - Pull repeated instructions outside the loop
 - Use temporary variables to hold intermediate values for reuse
- Less branches -> better programmin
 - 3 branches is more than 1 branc worse than 2 branches!



	TFLOPs	Time	TFLOP/
v5.swapindices	3.52	1.16	3.03
v6.simplify	3.30	1.10	3.00





V7. Replace Divides

- Replace (complex) div. with (double) rcp. and (complex) mul.
- **Lower instruction count: 40%**
- More bandwidth bound now!

	TFLOPs	Time	TFLOP/s
v6.simplify	3.30	1.10	3.00
v7.divs	2.09	0.66	3.18











V7. Replace Divides

Can be confirmed by Nsight Compute profiles









V8. Replace abs(x) with x**2



- sqrt(complex) vs power of 2 •
- **Causing pipeline to wait** ٠

	TFLOPs	Time	TFLOP/s
v7.divs	2.09	0.66	3.18
v8.abs	1.99	0.62	3.23









V8. Replace abs(x) with x**2

Before:

Wait: warp stalled waiting on a fixed latency execution dependency ____

#	Source	Sampling Data (All)	Sampling Data (Not Issued)	Instruction
308	rden = 1D0 / rden	31,172	17,748	857,
309	ssx = -Omega2 * conjg (cden) * rden * delw	44,670	27,041	1,200,
310	ssx = —Omega2 * delw / cden	0	0	
311	endif	0	0	
312	upd1 - 0.000	6	0	
313	<pre>if (abs(ssx) .le. ssxcutoff .or. wxt .ge. 0.0d0) then</pre>	467,330	225,479	10,351,
314	upul = ycoulxocc * ssx * matngmatmgp	Total Sample Coun	t: 467330 90, 336	2,300,
315	endif	Dispatch Stall: 102	96 (2.2%) Ø	
316 !	<pre>if (abs(ssx) .gt. ssxcutoff .and. wxt .lt. 0.0d0) then</pre>	Imc Miss: 59 (0.0% Math Pipe Throttle	6) : 92769 (19,9%) 0	
317 !	upd1 = 0.0d0	Misc: 518 (0.1%)	0	
318 !	else	No Instructions: 25 Not Selected: 4237	849 (5.5%) 78 (9.1%) 0	
319 !	upd1 = vcoulxocc * ssx * matngmatmgp	Selected: 6/328 (1	4,4%) 0	
320 !	end if	Wait: 217938 (46.6	5%) 0	
321	upd2 = vcoulx * sch * matngmatmgp * 0.5d0	186,963	98,916	2,684,
322	ssx_array_3 = ssx_array_3 + upd1	46,678	25,214	766,
323	sch_array_3 = sch_array_3 + upd2	47,767	24,797	766,
324	enddo ! loop over n1_loc	0	0	
325	enddo ! loop over g	125,175	109, <mark>449</mark>	149,









V8. Replace abs(x) with x**2

After:

Wait: 46.6% -> 23.7% _

#	Source	Sampling Data (All) Sampling	g Data (Not Issued)	Instructions Exe
309	wdiff = -Omega2 * conjg(cden)	16,491	10,128	342,907
310	ssx = rden * delw * wdiff	23,965	14,655	685,815
311 !	ssx = —Omega2 * delw / cden	0	0	
312	endif	0	0	
313	upd1 = 0.0d0	0	0	
314	rden – ssx * conjg(ssx)	50,010	28,089	766,944,
315	if (rden .le. ssxcutoff .or. wxt .ge. 0.0d0) then	45,380	25,259	1,150,416
316	upd1 - vcoulxocc * ssx * matngmatmop	Topi Sample Court 45380	111,372	2,300,832
317	endif	Dispatch Stall: 929 (2.0%)	0	
318 !	<pre>if (abs(ssx) .gt. ssxcutoff .and. wxt .lt. 0.0d0) then</pre>	Math Pipe Throttle: 19231 (4 Misc: 37 (0.1%)	(12.4%)	
319 !	upd1 = 0.0d0	Not Selected: 6591 (14.5%)	0	
320 !	else	Selected: 7853 (17.5%) Wait: 10739 (23.7%)	0	
321 !	upd1 = vcoulxocc * ssx * matngmatmgp	0	0	
322 !	end if	0	0	
323	upd2 = vcoulx * sch * matngmatmgp * 0.5d0	232,746	141, <mark>5</mark> 86	3,067,776
324	ssx_array_3 = ssx_array_3 + upd1	27,603	12,011	766,944
325	sch_array_3 = sch_array_3 + upd2	79,874	45,706	766,944
326	enddo ! loop over n1_loc	0	0	









V9. Cache Blocking

- Non-coalesced memory access for agsntemp
- **Causing Long Scoreboard Warp State**
 - Warp stalled waiting for L1TEX (local, global, surface, tex) memory operation ____

	#	Source	Sampling Data (All)	Sampling Data (Not Issued)	Instructions Executed	Predicated-On Thread 📥
	282	<pre>vcoulx = vcoul_loc(my_igp)</pre>	106	19	2,876,040	
	283	\$ACC LOOP SEQ	0	0		
	284	<pre>do n1_loc = 1, ntband_dist</pre>	55,775	16,601	3,464,669,589	
	285	<pre>aqsmconj = conjg(aqsmtemp_local(n1_loc,my_igp))</pre>	107,540	32,579	4,604,540,040	1
	286	<pre>matngmatmgp = aqsmconj * aqsntemp(ig,n1_loc)</pre>	1,261,414	907,564	3,067,776,000	
	287	vcoulxocc = vcoulx * occ_array(ni_loc)	86,741	Total Sample Count: 9075	4,304,000	
	288	<pre>wxt = wx_array_t(n1_loc,iw)</pre>	79,604	Dispatch Stall: 571 (0.1%	4,304,000	
	289	wdiff = wxt - wtilde	44,174	Imc Miss: 28 (0.0%)	0,416,000	
	290	wdiffr = wdiff*conjg(wdiff)	56,809	Long Scoreboard: 852075	5 (93.9%) 5,944,000	
	291	rden = 1d0 / wdiffr	239,665	Misc: 52 (0.0%)	8,192,000	1
	292	<pre>delw = wtilde*conjg(wdiff)*rden</pre>	196,6 <mark>8</mark> 1	No Instructions: 7 (0.0%)	4,304,000	
	293	! delw = wtilde / wdiff	0	Walt. 59717 (4.4%)		
	294	delwr = delw ∗conjg(delw)	64,322	42,701	766,944,000	
	295	sch = 0.0d0	0	0		
	296	ssx = 0.0d0	0	0		
	297	if (wdiffr.gt.limittwo .and. delwr.lt.limitone) the	n 64,705	43,075	766,944,000	
	298	sch = delw * epsa	63,513	37,943	1,533,87 <mark>9,720</mark>	
	299	cden = wxt**2 - wtilde2	22,067	10,628	767,419,200	
•						Þ







V9. Cache Blocking

- **Break loops into chunks and reuse data across threadblocks** •
- **Increase L2 hit rate** •

!\$ACC LOOP GANG VECTOR do ig blk = 1, ig blksize **!\$ACC LOOP SEQ** do ig = ig blk, ncouls,

ig blksize

	TFLOPs	Time	TFLOP/s
v8.abs	1.99	0.62	3.23
v9.block	2.00	0.57	3.50







V9. Cache Blocking

Less Long Scoreboard samples and higher L2/L1 hit rate

288 do n1_l	<pre>oc = n1loc_blk, ntband_dist, n1loc_blksize</pre>	109,008	50, <mark>326</mark>	3,015,710,080
289 aqsmco	onj = conjg(aqsmtemp_local(n1_loc,my_igp)	106, <mark>38</mark> 6	34,899	3,080,417,280
290 matngr	<pre>natmgp = aqsmconj * aqsntemp(ig,n1_loc)</pre>	464,140	301,148	3,075,808,000
291 vcoul	<pre>kocc = vcoulx * occ_array(n1_loc)</pre>	63,288	28 Total S	ample Count: 301148
292 wxt =	<pre>wx_array_t(n1_loc,iw)</pre>	127 <mark>,09</mark> 4	32 Dispato	h Stall: 1057 (0.4%)
293 wdiff	= wxt - wtilde	95,442	56 Imc Mis	s: 12 (0.0%) ttle: 194 (0.1%)
294 wdiff	r = wdiff*conjg(wdiff)	139 <mark>,2</mark> 57	79 Long S	coreboard: 222222 (73.8%)
295 rden =	= <mark>1</mark> d0 / wdiffr	307,221	170 Math	pe Throttle: 26215 (8.7%) ottle: 3 (0.0%)
296 delw =	= wtilde*conjg(wdiff)*rden	204,412	121 Misc: 7	5 (0.0%)
297 ! delw = wtilde /	wdiff	0	No Inst Wait: 5	ructions: 1 (0.0%) 1369 (17.1%)
298 delwr	= delw∗conjg(delw)	82,954	51 <mark>,381</mark>	768,952,000

Memory Workload Analysis

Detailed analysis of the memory resources of the GPU. Memory can become a limiting factor for the overall kernel performance when fully utilizing the involved hardware units (Mem Busy), exhausting the available communication bandwidth between those units (Max Bandwidth), or by reaching the maximum throughput of issuing memory instructions (Mem Pipes Busy).

Memory Throughput [Gbyte/second]	131.85 (-78.99%) Nem Busy [%]
L1 Hit Rate [%]	71.38 (+27.27%) Max Bandwidth [%]
L2 Hit Rate [%]	88.89 (+3272.88%) Men Pipes Busy [%]









Ω

31.10	(+29.96%)
28.44	(-59.32%)
12.84	(+59.44%)



Summary

8 Steps to Optimize Sigma-GPP

- Collapse n', G', and G loops 1.
- Bring n' loop in; collapse only G' and G 2.
- Adjust threadblock size 3.
- **Reduce branching; pull iw loop outside** 4.
- Swap indices to suite parallelisation 5.
- Simplify code 6.
- **Replace div. with rcp. and mul.** 7.
- **Replace abs with power of 2** 8.
- 9. Cache blocking

```
!$ACC PARALLEL LOOP REDUCTION (+:
do n1 = 1, nbands
   do igp = 1, ngpown
      do ig = 1, ncouls
         do iw = 1, nw
```

	TFLOPs
v1.collapse3	3.71
v9.block	2.00
	3x !





compute and reduction



Conclusion

- Code is still bandwidth and latency bound
 - shared memory
 - lower register count
 - improve FMA ratio

 Together with profilers, Roofline provides the complete solution for your performance analysis and optimization needs!









Roofline is a great way to Frame Conversation with Application Teams

Helps Motivate, Direct and Visualize the optimization process.

Coming soon to a tool near you!!

























BerkeleyGW Use Case

- Big systems require more memory. Cost scales as N_{atoms}^2 to store the data. \star
- In an MPI GW implementation, in practice, to avoid communication, data is duplicated and \star each MPI task has a memory overhead.
- ★ Users sometimes forced to use 1 of 24 available cores, in order to provide MPI tasks with enough memory. 90% of the computing capability is lost.









In house code (I'm one of main developers). Use as "prototype" for App Readiness.

Significant Bottleneck is large matrix reduction like operations. Turning arrays into numbers.

$$\langle n\mathbf{k} | \Sigma_{\rm CH}(E) | n'\mathbf{k} \rangle = \frac{1}{2} \sum_{n''} \sum_{\mathbf{q} \mathbf{G} \mathbf{G}'} M_{n''n}^*(\mathbf{k}, -\mathbf{q}, -\mathbf{G}) M_{n''n'}(\mathbf{k}, -\mathbf{q}, -\mathbf{G}') \\ \times \frac{\Omega_{\mathbf{G} \mathbf{G}'}^2(\mathbf{q}) \left(1 - i \tan \phi_{\mathbf{G} \mathbf{G}'}(\mathbf{q})\right)}{\tilde{\omega}_{\mathbf{G} \mathbf{G}'}(\mathbf{q}) \left(E - E_{n''\mathbf{k} - \mathbf{q}} - \tilde{\omega}_{\mathbf{G} \mathbf{G}'}(\mathbf{q})\right)} v(\mathbf{q} + \mathbf{G}')$$









The Ant Farm Flow Chart

Make Algorithm Changes

So, you are Memory Bandwidth Bound?



Identify the key arrays leading to high memory bandwidth usage and make sure they are/will-1. be allocated in HBM on Knights Landing.

Profit by getting ~ 4-5x more bandwidth GB/s.







What to do?

Make sure you have good OpenMP scalability. Look at VTune to see thread activity for major 1. **OpenMP** regions.



1. Make sure your code is vectorizing. Look at Cycles per Instruction (CPI) and VPU utilization in vtune.

See whether intel compiler vectorized loop using compiler flag: -qopt-report=5









You may be memory latency bound (or you may be spending all your time in IO and Communication).

If running with hyper-threading improves performance, you *might* be latency bound:

If you can, try to reduce the number of memory requests per flop by accessing contiguous and predictable segments of memory and reusing variables in cache as much as possible.

On Knights-Landing, each core will support up to 4 threads. Use them all.







Are you memory or compute bound? Or both?







ound



Cray XC40 system with 9,600+ Intel Knights Landing (KNL) nodes:

- 68 cores, 272 Hardware Threads
- Up to 32 FLOPs per Cycle, 1.2-1.4 GHz Clock Rate
- Wide (512 Bit) vector Units
- Multiple Memory Tiers: 96 GB DRAM / 16 GB HBM
- NVRAM Burst Buffer 1.5 PB, 1.5 TB/sec



